UNIVERSITY OF OKLAHOMA

GRADUATE COLLEGE

A TELEOLOGICAL CONCEPTION OF FINANCIAL MARKETS

A DISSERTATION

SUBMITTED TO THE GRADUATE FACULTY

in partial fulfillment of the requirements for the

Degree of

DOCTOR OF PHILOSOPHY

By

VIOLET IVEY VICTORIA

Norman, Oklahoma

2026

A TELEOLOGICAL CONCEPTION OF FINANCIAL MARKETS

A DISSERTATION APPROVED FOR

THE SHYAM DEV PATWARDHAN DEPARTMENT OF PHILOSOPHY

BY THE COMMITTEE CONSISTING OF

Dr. Stephen Ellis, Chair

Dr. Nancy Snow

Dr. Wayne Riggs

Dr. Zev Trachtenberg

Dr. Lubomir Litov

Acknowledgements

One of the many benefits of not being a movie star or musician at an awards show is that I hold no one hostage spending a long time giving acknowledgements. Fueled by my expectation of future guilt lest a formative story go untold, I hereby disclose that this section is rather long and absolve you, dear reader, from muddling through any of it. Suffice it to say that I am one of the lucky few burdened with an exorbitant number of things to be thankful for. You have been warned.

My mother was 19 years and 359 days old when she had me, an undergraduate sophomore at this university. She met my father at Pepe Delgado's, where he was a line cook. Though their union was ill-fated, I feel confident that my preferences for Tex Mex were formed in utero thanks to Sir Delgado's foresight to locate his restaurant on Campus Corner. In this sense, I was Sooner Born (thanks, mom).

My mom loves to tell the story of breast feeding me in her Latin final, as the perfect encapsulation of the suffering of a young mother fighting to earn her degree. She also loves to tell the equally harrowing stories about my early propensity for scribbling in all her very expensive college textbooks. Although too early in my life to recount from a first-person perspective, I am told my writings were prolific in the sense that once I had a book in my grasp, scarcely a page was left unscribbled upon. With equal heroism, my grandmother tells the story of how, when getting her master's degree at OU (while parenting four children and working at 'the Pepe Delgado's of Pauls Valley,' Tio's), how after getting a parking ticket in Norman on her way to class, she hauled her kids with her to the courthouse to petition her ticket. I am not at liberty to say whether this specific appeal was successful, but Bonita Ivey never received another parking ticket in the City of Norman. As of this writing, my work on the dissertation has not

yielded any brushes with the law or motherhood, but the stubbornness and perseverance which I suspect is in our blood and which certainly has been cultivated by the women in my family has served me well in this process. One might say that in this sense, I was Sooner Bred.

Bonita married a man from Pauls Valley, Oklahoma, who despite never taking a single class at OU was perhaps their biggest football fan. One Christmas, someone in the family purchased him a bottle opener that was designed to play the Sooner fight song when used to open a beer. He didn't drink in the manner required to get one's money's worth from such an object, so instead, every time the football team would score a touchdown, he'd wedge the tip of a #2 pencil in just right to elicit the opener's tinny, garbled rendition of the fight song. Those in the immediate vicinity were expected to sing along. He was a devoted fan, and I am so lucky that he was a fan of mine. The man served on the rescue mission that recovered the Apollo 11 astronauts from their first walk on the moon but still came to and cheered me on for every small, insignificant in retrospect, achievement or competition of mine as though it was as important as his mission on the USS Hornet.

A few decades earlier in a town far away, a little boy moved from Puebla, Mexico to Los Angeles with his single mother. She became a seamstress, and eventually owned her own studio where she designed gowns for glitzy Hollywood starlets of the mid-century. Once the little boy learned English trial-by-fire style in the LA public school system (in a manner not dissimilar to a dissertation defense), he worked as a cabana boy at the Beverly Hills Hotel. He was a diligent worker and was always on time, although he was once reprimanded for eating a guest's leftover ice cream in the walk-in freezer after a shift. He trained as a draftsman in high school, and his teacher encouraged him to go to community college, which he did. He joined the air force as an airplane mechanic, travelling the world with his wife and family. After a stint in the air force

where he earned an Airman's Medal for rescuing two pilots from a burning plane during the Vietnam War, he moved with my grandma to Norman, Oklahoma. They often faced incidents of racism that only plagued them abroad when they were on military bases, enduring insults and slurs on at least one occasion by young men who were presumably students at this University.

A boy that began his life mixing dirt in with flour in order to have enough tortillas to eat, came to America where he served his country, had four kids, owned a house, and was able to retire. If anything resembling the American Dream ever existed, he surely achieved it. Still, he had an outsider's perspective, and in many situations was seen to be an outsider. What was surely at first uncomfortable for him became the way he engaged with the world– as the odd man out who could see and criticize the things he encountered in ways that are perhaps unnatural to those who fit in or inherit their surroundings. This ability to thrive as the one thing that doesn't look like the others has enabled me to do the same, and I thank him for making me see my outsider status in almost every situation I encounter as a strength. During our last lucid conversation with each other, he asked me what I wanted out of life. I told him that I really wanted to write a book. This dissertation is not quite a book, but now this commitment will be written somewhere that will haunt me if I don't do it. Here I am obliged to mention that when I die, I'll be Sooner Dead.

To Wayne Riggs and the graduate student admissions committee of 2019, thank you for tolerating my vaguely threatening application materials, and for taking a chance on me despite my never having taken a traditional philosophy class. Because of your risk appetite, I was able to spend the last years of my grandfather's life with him, and that will forever be a cherished time I wouldn't have had otherwise.

I am thankful to Dan Press, and Dr. Karl Jacoby, for making me feel like I had good ideas, and not being too disappointed with me for choosing finance over graduate school. Perhaps they were okay with it because they knew it wouldn't take. To Dr. Luke "Doc" Mayville, for being my first philosophy teacher, although I didn't realize Contemporary Civilization was ostensibly a philosophy course, or that philosophy still existed as a discipline at the time. To Dr. John Paul Broussard and Jane Baber, for trusting me to teach an M.S. course in Finance Ethics as a PhD student. To Dr. Max Latona, for hiring me for a job I didn't apply to.

I would also like to thank the Zicklin Center Business Ethics Workshop, and the department of Legal Studies and Business Ethics at Wharton for first exposing me to the subdiscipline I am becoming a part of and for graciously hosting me in their department for the Spring 2026 semester. I would particularly like to thank Brian Berkey – who likely organizes more conferences in an academic year than most attend – for being an exemplar of someone who does philosophy as a collective social practice.

I'd like to thank a boatload of other people: Matt Levine, for giving me something enjoyable to read every day in both my old career and new one. I'd like to thank Kandyce Cudjo and Ashley Barnett from the spin studio formerly known as Cyclebar Classen Curve for helping maintain whatever semblance of sanity I have left. An anesthetized thanks to my surgeon Dr. McGinnis, for stapling what remained of my shoulder tendons back together, and Natalie Wansick and the rest of Goddard's Physical Therapy team for making me hurt more for the sake of hurting less.

Charlotte McMeekin, for providing me with homemade pastries I would have bankrupted myself purchasing otherwise. Joshua Faytinger, for enabling my consumption of food and drink at several of Google's spectacularly designed corporate offices while working on this project.

Orly Morgan, for tolerating my minimalism phases, infrequent use of the cell phone, and general unseriousness. Abigail “Gail” Andrews, for getting us lost in the Middle East and being Oklahoma City’s second-biggest fan. Jacob Andrews, for laughing at my jokes. Kate Nathan, for insisting that I Uber to Moynihan Hall. Stephanie Holt, for reciprocating the transformation of mundane stories into scandalous tales of gossip and intrigue. Chase McCumber, for interpreting “carpool” to mean "chauffeur service.” Max Geisler, for the frequent distractions that, while admonished in the moment, were secretly appreciated. Justis Huddleston, for making a mean dirty martini and being the sole purveyor of a well-oiled morning coffee delivery service. And Luca Brasi, for reminding me every day that nothing is better than a walk.

To Steve Ellis, I issue an apology for the many hours of additional work I’ve saddled you with over the years. My apologies to Cindy and the wider Norman community, for distracting one of its best do-gooders from doing good.

Finally, inspired by David Edmond’s innovation of concluding acknowledgements with a petty complaint, I would like to *sincerely* thank Google Docs for their commitment to single-minded and wholly indecipherable aesthetic choices. Thank you, for your insistence that everyone view only one page at a time, lest they confuse themselves with the ability to see multiple pages at once. And thank you in addition, for the innovative cursor design, that lags unpredictably and sometimes looks as though it is chopping the last letter typed in half. It is so refreshing to be held on the tip of one's toes, not knowing whether the next character one intends to type will go at the end of the last written word or be spliced in between.

TABLE OF CONTENTS

Abstract

This dissertation, *A Teleological Conception of Financial Markets*, argues that financial markets have a specific purpose—allocating risk to enable growth—and that this purpose shapes how we should evaluate their function and ethics. Existing criticisms of financial markets often overlook their potential role in mitigating risk for society's benefit. When functioning properly, financial markets distribute risk broadly, helping individuals and firms hedge against uncertainties in essential sectors like agriculture, energy, and housing. Stability as a social good includes preventing price swings or shortages in life-sustaining goods. I argue that understanding financial markets requires normative judgments about how much financial risk a society should bear to promote economic development and maintain stability.

A simple example illustrates this: A wheat farmer cannot control a variety of factors that will contribute to the price their crop will ultimately fetch on the market. The weather might be good, yielding them lots of high-quality product, or there might be a drought, in which case the farmer faces the risk of not making enough money from that year's yield to continue being a farmer in the future. Financial systems, when functioning properly, help mitigate this kind of uncertainty by providing mechanisms to manage risk and ensure greater stability. Futures contracts enable two parties- the farmer who wants to guarantee a certain price for their crop, and a financier who takes on risk for the opportunity to make some financial return- to enter into a market transaction which distributes risk (from the farmer to the financier) in such a way as to make society better off. The farmer is essentially purchasing an insurance policy that guarantees a specific price (or a narrow range of prices) at which they can sell their wheat. This means that in a worst-case scenario (a drought, let us say), the farmer will still live to farm another season. This is because the farmer has guaranteed a price that will allow them to continue farming. In a

best-case scenario (ideal weather conditions), the farmer will not make *as much* money as they would have made without the 'insurance policy.' The reduction in their maximum reward in the best-case scenario is worth it to them (and to us, more generally, as a society that wants to have wheat readily available and at stable prices), because it prevents them from going bankrupt in a worst-case scenario. The financier on the other side of the transaction is presumably also entering into other bets, some of which will pay off and some of which won't.

Their self-interested actions, taken with a motivation to generate wealth on part of the financier and in order to sustain their livelihood in the farmer's case, result in a productive distribution of risk. This ability to smooth out the peaks and valleys of potential risk is not always properly appreciated—either by finance professionals or by those in academia.

My dissertation has three aims: First, to argue that financial markets differ from other markets commonly analyzed by philosophers in ways that existing theories fail to account for– specifically, the aforementioned propensity for distributing risk, along with the diachronic function that financial markets have in determining what product and service markets will look like in the future. Second, I argue that the purpose of financial markets (in contrast to other markets) is to distribute risk in such a way as to enable growth in the economy while managing volatility. Given that finance is an institution which impacts broader society, I make the claim that the risk taken on by financial markets is worth collective, democratic deliberation. Finally, I demonstrate the usefulness of this theoretical framework, making sense of things in modern financial markets that current theories cannot help us make sense of. I discuss each of these tasks now in turn:

In Chapter II, I establish a fundamental distinction between financial markets and markets for goods and services. Markets for goods and services certainly have a robust literature around them: there are philosophers who talk about the virtues and vices of markets for controversial goods like organs, drugs, and weapons, and services like surrogacy and prostitution. Both types of markets facilitate exchanges that reflect participant preferences, but financial markets serve a different role: they are designed to manage risk and direct capital toward productive opportunities. This future-oriented nature means they cannot be assessed using the same ethical and economic principles that apply to ordinary markets.

In addition to providing a traditional literature review in Chapter III, I want to establish that the literature on markets and finance largely fails to address the components of futurity and risk that my project draws out. The philosophical and economic literature on markets focuses primarily on the conditions and results of exchange. This, I argue, leaves the systemic role of risk in society undertheorized.

The fourth Chapter develops an ethical framework for evaluating financial markets. While many existing approaches rely on utilitarian calculations of overall economic benefit (that is, what will maximize the greatest good for the greatest number) or rule-based (or 'deontological' in philosophy's parlance) principles of fairness and duty, I argue that a virtue ethics perspective—one that focuses on the character and purpose of financial actors and institutions—provides a better foundation. I defend a specific purpose (*telos*) for financial markets: they should facilitate the distribution of risk in a way that enables the broader economy to grow. This framework allows for a more nuanced assessment of market behavior, distinguishing between financial practices that contribute to productive investment and those that merely extract value without creating real economic benefits.

Finally, Chapter V applies this framework to persistent problems in financial markets, such as speculative bubbles (e.g., cryptocurrencies, meme stocks), financial crises (such as the housing crisis in 2008), and systemic risks (such as our reliance on the dominance of U.S. dollar-denominated assets) that arise despite regulatory interventions. Existing theories struggle to explain why these problems persist even in well-regulated markets. By grounding financial markets in their proper purpose, my framework offers new insights into how and why these failures occur—and how they might be mitigated.

By clarifying the role that financial markets should play in society, and framing risk-allocation as a subject for deliberative democracy, I offer a means for us to evaluate finance in a manner that takes into account both its current flaws and future potential to benefit society.

.

Chapter I: Dissertation Introduction

In this dissertation, *A Teleological Conception of Financial Markets*, I argue that financial markets have a specific purpose: ensuring the existence of appropriate future markets for products and services, including their capacity to handle human needs and desires, and by managing the risks inherent in such a project. This account of the purpose of financial markets shapes how we should evaluate their functioning.

This approach will seem odd to many. The burden of this dissertation is to make it make sense. I will borrow from philosophers who argue that markets as a whole are best viewed with a teleological lens. But I see the financial sector of the economy as a distinct kind of human enterprise. I argue that financial markets differ in a morally relevant way from markets for products and services. The former focus on a trio of considerations— sustainability, futurity, and risk— that the latter does not provide for themselves. This is the first novel claim of my project.

But first things first. To make the appeal to teleology plausible, I start by noting that rule-based ethical theories are often inadequate for fully understanding the ethical issues involved in complicated social structures. One needs, rather, to evaluate actions in such settings in light of the ends or purposes of the institutions in which they take place, i.e., to take a teleological perspective. To get an initial grasp on this kind of reasoning, consider an analogy. Imagine a world where K-12 schools (as we know them now) are understood merely as an emergent arrangement of childcare. Parents need time to work and rest, so they see the need for an institution that relieves them of constant supervision. Some parents may also want to feel that their money is "doing something," so a vague expectation arises that students should acquire a few skills to help them get jobs. But in this world, the purpose of education is poorly theorized. No one thinks much about civic formation or intellectual development. This does not mean that

schools do not play an important role in these areas. Children still spend considerable time in school and school dynamics still have a big influence on their development.

In this scenario, teachers will still play a central role in children's lives, and so we may want to ensure that they are kind, patient, and approachable. Suppose Carl is a teacher who is morally upstanding and personally virtuous yet lacks any pedagogical skill. But is Carl really a *good teacher*? In the real world, we would be inclined to say "no." We see a larger purpose for educational institutions. The point of school is not merely childcare or vocational training. Children need to be guided toward judgment and understanding, if for no other reason than to become good citizens. Schools are well-suited for the task of providing a knowledge and skill baseline for these future members of society. Those who see schools merely as childcare would be missing an opportunity to do some good for all.

Without a teleological account of education, we would struggle to adequately assess Carl's adequacy as a teacher or the institution's decision to employ him. That he is not personally vicious seems like a necessary but not sufficient qualification, not an endorsement of our actual conception of education. Noting that he is "nice" to the children matters a lot if the goal of schooling is daycare. It does not go very far, however, if the goal is to cultivate intellectual and moral development. Conversely, if we (mis)identify the purpose of education (say, as "brainwashing"), we will adopt practices that lead us further from some of the valuable outcomes of education (arguably, its proper ends).

The point of this analogy is to make plausible the idea that we cannot adequately assess behavior within an institution without reference to the institution's role, or at least possible roles, in a society. This lesson will guide my account of financial markets: their ethical evaluation

depends not merely on adherence to general rules of behavior, but on how actions contribute to or detract from the proper purpose of the markets themselves.[1]

Even apart from teleology, existing criticisms and defenses of financial markets tend to overlook their unique nature. Markets for goods and services are well theorized: they facilitate exchanges that register participants' preferences and are typically assessed through normative frameworks concerned with autonomy, fairness, and efficiency. Financial markets, however, are organized around distinct functions. They allocate capital toward productive opportunities and manage forms of risk that cannot be eliminated. We often extrapolate our views about issues in finance from our views about issues in markets for products and services, despite these two kinds of markets being different in key ways.[2] Philosophical theorizing about markets for products and services focuses primarily on how well markets satisfy our (perhaps idiosyncratic) wants or needs, and so how they impact our (current) well-being. Note that the basic story about exchange in markets for products and services hangs on preference heterogeneity. Where preferences differ, market exchanges can truly be win-win.

The literature on markets provides sophisticated analyses of the moral status of commodities such as organs, of services such as prostitution, and of role-based duties associated with market participation. Yet these frameworks offer limited traction on financial phenomena—such as the rise of retail trading—in which individuals engage with financial markets in ways that are not easily explained by standard models of economic rationality. Traditional normative tools direct us to ask whether rights are violated or whether harmful manipulation occurs; absent such violations, they interpret the activity as a mere expression of preference. Such approaches fail to capture what is ethically distinctive about financial markets themselves. They do not

[1] There are also issues about how to correctly identify the moral purposes of institutions, which I will return to later.
[2] I mean this from both the perspective of everyday views, and academic ones.

necessarily help us with the issues of sustainability, futurity, and risk that are crucial for our society, both near- and long-term. Sustainability is often understood in the guise of *growth* because handling the human needs and desires of future populations has historically been understood to entail a need for the future economy to produce more goods than the current one. When we describe markets for products and services, we usually do so in terms of actions that occur within a single period of time, e.g., shopping for the best car within our preference range. But financial markets, especially when functioning well, have a diachronic dimension. They shape - via investment - the future market for products and services. At the same time, the level of investment ensures that future markets for products and services will be adequate to the level of demand. They also distribute risk (e.g., by means of intertemporal contracts) in ways designed to help individuals and firms hedge against uncertainties. Appropriate technology is clearly a social good. Growth— ensuring the continued adequacy of markets for products and services as a social good— includes developing the economy at a rate to at least account for changes in population, and perhaps increasing quality of life.[3] Risk mitigation serves as a social good by preventing price swings or spot shortages in life-sustaining goods.[4]

I argue that understanding financial markets requires normative judgments about both the future direction of markets for products and services (content and size) and about the degree and distribution of financial risk a society should bear (see Figure 1).[5]

[3] Although it should be noted that populations can decrease over time, and there is likely an upper threshold to consumption goods and their improvement of wellbeing
[4] See the COVID baby formula shortage as one recent example
[5] As I'll describe in further detail later, there is a latent view that it's permissible to maximize growth (and thus risk), so long as one stays within the bounds set by laws and regulation.

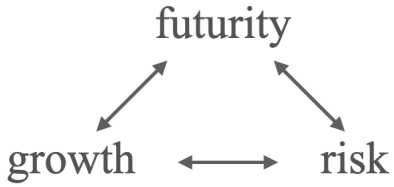

Figure 1: Focus of Financial Markets

The management of risk presents some of the most conceptually challenging issues in understanding financial markets, but a simple example illustrates how such markets can function to mitigate irreducible uncertainty. A wheat farmer cannot control a variety of factors that will determine the price their crop will ultimately fetch on the market. The weather might be good, yielding them lots of high-quality product. There might be a drought, in which case the farmer faces the risk of not making enough money from that year's yield to continue being a farmer in the future. Financial systems, when functioning properly, help mitigate this kind of uncertainty by limiting the range of prices a farmer can expect to get for their crop. A commodities futures contract is a financial product (a *derivative*) which specifies a time and price at which a fungible good can be traded. A buyer of such a contract agrees to purchase the commodity at the specified future date and price, while the seller agrees to deliver it under those terms.[6] The farmer is essentially purchasing an insurance policy that guarantees a specific price (or a narrow range of prices) at which they can sell their wheat. This means that in a worst-case scenario (a drought, let us say), the farmer will still live to farm another season. In a best-case scenario (ideal weather conditions), the farmer will not make *as much* money as they would have made without the 'insurance policy'. The reduction in their maximum reward in the best-case scenario is worth it

[6] These contracts are binding obligations, not mere options. The delivery component of futures contracts can be interesting. In most cases, financial institutions trading these derivatives will 'roll forward' their positions (closing out expiring contracts and opening new ones with later delivery dates) so they don't end up having to figure out how to store barrels of oil in their downtown Manhattan offices.

to them (and to us, more generally, as a society that wants to have wheat readily available and at stable prices), because it prevents them from going bankrupt in a worst-case scenario. The financier on the other side of the transaction is presumably also entering into other contracts, some of which will pay off and some of which will not. We do not care as much if the financier loses money from this transaction in the event of a drought, because we expect them to have diversified their risk. The self-interested actions of each counterparty in a financial transaction, taken with a motivation to generate wealth on part of the financier and to sustain their livelihood in the farmer's case, result in a productive distribution of risk.[7]

These sorts of financial interactions also serve to determine what the landscape of agriculture will look like in the future. The farmer need not grow wheat. The range of commodities futures available and their price are also important for a farmer's decision to plant some crops over others. They likely make an all-things-considered judgement about what to grow, considering a variety of factors such as the cost of farming a specific crop, the historical range of prices and demand, and the cost of 'insuring' some price range. If, for example, the insurance for barley is a lot cheaper and consumer demand is increasing, even though barley does not get as high of a price, the farmer might want to switch to barley. What the financial market environment looks like is also dependent on the state – what sanctions, subsidies, regulations, and laws exist that pertain to the market in question.[8] So finance can help a farmer

[7] This example is meant to illustrate a limited functional role that financial markets can play in redistributing certain forms of risk. It is not intended as an invisible-hand argument that such outcomes are sufficient to justify financial markets as promoting social good overall. Indeed, I argue that treating this risk-mitigating function as self-justifying has obscured important structural and systemic failures of contemporary finance. Here, the role of failure is important: we would rather the farmer take actions to prevent themselves from going bankrupt (in the event of a drought, by distributing their risks to others), and we are perhaps unflinching at the idea that the financier on the other side might lose money from this particular trade. For the sake of this simplified illustration, we will ignore systemic risks for now.

[8] The U.S. subsidizes the production of specific crops, which serves to shape our (over)weighting of corn production relative to its market price, for example. At risk of stating the obvious, these government interventions needn't be constrained by traditional economic rationality. We might have national security, cultural, or many other non-

both protect their livelihood, and nudge (or shove) them to farm some things over others. It should also be obvious that although finance can be used in this way, it both 1) does not necessarily do that and 2) the net-value of a positive financial transaction or relationship can be eroded or cancelled out by negative effects in other areas.[9]

The foregoing serves as a start on the first aim of my dissertation, to argue that financial markets differ from markets for products and services in ways that existing normative accounts fail to recognize. Second, I argue that the purpose of financial markets is closely tied to what makes them special. Given the role that finance plays in the broader society, I claim that what counts as the functioning of financial markets (with respect to growth, futurity, and risk) is worthy of collective, democratic deliberation.[10] The final aim of my dissertation is to demonstrate the usefulness of this teleological framework in making sense of judgments about modern financial markets that standard rule-based theories can't help us with. Overall, I intend for this dissertation to spur more accurate philosophical analysis of financial markets and so contribute to broader debates about financial regulation and economic justice. I will, for instance, make an argument for regulations that fluctuate with macroeconomic factors in the same manner that central bank-dictated interest rates do. This project will also connect risk allocation to democratic decision-making and economic justice, examining how finance can better serve society.

---

economic reasons to want to bolster certain industries or goods within industries. This paper won't address the validity of those practices, other than to say that they obviously change the market landscape compared to what it would look like in the absence of state intervention.

[9] For example, using futures contracts as insurance may make the farmer better off, but their 401k going bust because of price drops in public markets might have a bigger negative impact, on a consequentialist basis.

[10] 'Growth' has become a politicized term because there has been little debate as to *how much* growth there ought to be given that with growth will necessarily come risks– growth is taken to be either an altar to worship under or something to chagrin as responsible for economic ills.

The overall structure of the work tracks my aims. Chapter II establishes by way of explanation a fundamental distinction between financial markets and markets for goods and services. Chapter III surveys existing literature on both the ethics and political economy of markets, demonstrating that although these discussions are extensive, they remain incomplete in their treatment of finance. Chapter IV develops my positive account: a teleological framework for evaluating financial markets. Existing approaches to the morality of markets are, broadly, 'preference satisfaction' views, or 'rule-based' (or 'deontological' in philosophy's parlance) principles of fairness and duty. The former relies on utilitarian calculations of overall economic benefit (that is, what will maximize the greatest good for the greatest number). In contrast, I argue for a virtue ethical perspective that focuses on the purpose or *telos* of finance as a social institution. This approach is needed to thoroughly evaluate market practices in terms of their contribution to the system's constitutive ends rather than in terms of aggregate welfare or individual compliance with moral rules. This teleological perspective not only clarifies the conditions under which financial markets function well but also provides a principled basis for identifying when and how they fail to achieve or corrupt their own purpose. I defend a specific purpose (*telos*) for financial markets: they should distribute capital and risk in such a way as to bring about the future markets for goods and services, while maintaining stability and appropriate growth in the economy to meet future people's wants and needs.[11] This framework allows for a more nuanced assessment of market behavior, and will entail that even those agents working in purely self-interested ways may contribute usefully to the overall picture.[12]

[11] What is considered 'appropriate,' I'll argue, should be based on circumstances like anticipated medium and long-term population growth and an element of democratic deliberation.

[12] This is also a solution to the problem that it's impossible to police actions in financial markets based on something like 'the self-interested intention of the financier to speculate.' An agent's subjective intention to profit from, let's say, price movements rather than hedge an existing exposure to risk is virtually impossible to discern. So, trying to 'prevent speculation' in a market is a tall order for any regulatory approach. What my teleological framework

My argument does not depend on any controversial metaphysics of institutions. In its weakest version, my argument can be understood as simply saying that it is useful to think of institutions as if they have a *telos*. Political philosophers spend a lot of time talking about the proper orientation and construction of the state to engender justice and prevent injustice. Setting aside whether the state or anything else has a proper *telos*, it is justified to spend time thinking about the point of having a state because states are institutions that impact a lot of people. Virtually everyone operates under some regime or other and few, if any, of us are free from the impact that a just or unjust state can have on our lives. Financial markets impact almost the same number of people, and yet when we talk about finance, philosophers tend to take for granted the way it operates. The market just is, and no one asks what it is for. But once we see financial markets as institutions like states, it will be unsatisfactory to take them simply as features of the world.

This is how we treat other major institutions. Even those who favor a minimal state or who emphasize individual liberty still debate which state structures are preferable, which powers are legitimate, and which orientations promote or undermine important things like justice. We should treat financial markets similarly: as institutional configurations we can assess, criticize, and improve, rather than as unalterable fixtures whose pathologies we merely tolerate and deal with the repercussions of.

Chapter V applies my teleological framework to persistent problems in financial markets, such as speculative bubbles (e.g., cryptocurrencies, meme stocks), financial crises (such as the housing crisis in 2008), and systemic risks (such as our reliance on the dominance of U.S. dollar-denominated assets) that arise despite regulatory interventions. Existing theories struggle to

suggests instead is to evaluate actions in terms of their overall possible contribution to the constitutive ends of the financial system, independent of individual motivations.

explain why these problems persist even in well-regulated markets. By grounding financial markets in their proper purpose, my framework offers new insights into how and why these failures occur—and how they might be mitigated. Regulatory solutions are discussed where relevant. I should briefly note a challenge to identifying 'good' regulation: Most regulatory solutions operate with what I will call 'regulatory determinism.' This ontological view assumes that the appropriate regulations are in the sky, ready to be enacted with enough political will. Once we discern what the right lines to draw are, like what level of assets under management should denote a bank's being worthy of increased regulatory requirements, what level of personal income and wealth should denote a 'qualified investor' who is allowed to invest in private markets, or how much cash a bank should keep on hand as 'capital requirements,' our attitude tends to be that those are the regulations which should exist in perpetuity.

My approach will suggest we consider an adaptive attitude towards regulation, which takes into account environmental circumstances in the same manner that national interest rates fluctuate with the ebb and flow of external factors like rates of unemployment and inflation. What informs what we think the 'right regulation' is might come from a consequentialist, deontological, or behavioralist approach to finance, but in my view, none of these theories will be able to come up with a set of regulations that should exist in perpetuity, because the nature of financial regulation should allow for more flexibility based on environmental conditions.

By clarifying the role that financial markets should play in society, and framing risk-allocation as a subject for deliberative democracy, I offer a means for us to evaluate finance in a manner that takes into account both its current flaws and prognosis for how we might restructure finance as an institution for the benefit of society.

Chapter II: Overview of Financial Markets

The purpose of this chapter is to demonstrate the morally salient differences between financial markets and markets for products and services. I take it that most readers are decently familiar with the relevant features of markets for products and services, so I will touch on them only briefly. I proceed primarily by describing how financial markets work (or at least how participants in those markets think they are supposed to work). My narrative focuses on descriptions of financial assets, institutions, and agents most relevant to my project. This chapter has two aims for the sake of my overall argument. First, it enables the reader to understand how financial instruments function, especially the forms of value they generate independent of wealth accumulation or the satisfaction of individual preferences. Second, it illustrates why the analytical tools appropriate for thinking about markets for products and services are inadequate for explaining the structure and operation of financial markets.

Although this project ultimately develops a teleological account of financial markets, nothing in the present chapter presupposes that framework. Here, I simply present finance in the terms financiers use, to show that the practice already offers conceptually rich materials for philosophical reflection. If some of these practitioner terms sound teleological, the reader should understand the resemblance as coincidental rather than intentional. No normative commitments are built into the exposition here.

A popular conception of markets in general asserts that, given a certain set of market conditions, the pursuit of self-interest on part of market participants will yield a socially optimal allocation of scarce resources. The conditions in question are known collectively as "perfect competition": no interdependencies amongst agents' utility functions (e.g., my utility doesn't go down just because yours goes up); no missing markets (trading options are complete, such that

all utility issues can be addressed by market transactions); no barriers to entry into or exit from a given market, so that no one is 'excluded' or 'trapped' in a market; and no one has market power because there are a sufficient number of traders such that no single agent may influence price. Under conditions of perfect competition, market transactions among instrumentally rational agents will be *Pareto improvements*: the results of a transaction will be strictly preferred by at least one person, and no one will strictly prefer the pre-transaction situation (or else they would not have participated). As a result, the equilibrium allocation that results from a market under conditions of perfect competition will be *Pareto optimal* (allocatively efficient): no further Pareto improvements will be possible. This is the content of the First Fundamental Theorem of Welfare Economics.[13]

However, the terminology itself can invite a conflation between Pareto improvement and welfare improvement. Because Pareto improvements are defined in terms of preference satisfaction, it is easy to slide from the formal notion ("at least one agent prefers the new state") to a substantive conclusion about well-being ("the agent who preferred the new state has been made better off"). But the fact that a preference is satisfied does not, in itself, establish that anyone's welfare has genuinely increased. If we treat "pareto improvement" as synonymous with "welfare improvement," we make the inference that satisfying someone's preferences increases their actual well-being. This inference is highly implausible. People routinely want and voluntarily purchase things that are not good for them, even by their own lights. There are plentiful markets in which one can satisfy their preferences for various drugs and substances, which most purchasers readily acknowledge the consumption of runs counter to their welfare

[13] Daniel M. Hausman and Michael S. McPherson, *Economic Analysis, Moral Philosophy, and Public Policy*, 3rd ed. Cambridge University Press, 2016.

improvement.[14] More dangerous still, would be a conflation between the term 'pareto efficient' and a situation in which a change could be implemented which would make *everyone* better off. This would completely muddle the term's technical meaning. A situation is Pareto efficient precisely when no such change is possible.

There is a *lot* to be said about this highly idealized characterization of the market, but the most important reactions boil down to explaining 1) the ways in which our reality is non-ideal and 2) how the market mechanism can and has broken down. The conditions of perfect competition are clearly idealizations: preferences are sometimes entangled, markets are not always complete (you can't buy salvation!), and transactors can have market power to various degrees. But the basic idea that many market transactions mutually satisfy preferences points to a way in which markets generally help people pursue their goals.

For the purposes of this dissertation, it is important to see that this description of 'how markets work' oscillates between describing markets for products and services and financial markets. The key features of both sorts of markets are colloquially understood to be the same: given a set of conditions, the pursuit of self-interest will yield an optimal allocation of scarce resources. This approach obfuscates a number of important issues. In some ways, financial markets are closer approximations to the idealized market mechanism: in the market for publicly traded securities, for example, there are a lot of buyers and sellers. As agents, professional traders (at least) act as rational agents, carefully and coolly tracking information and calculating the value of their options[15], in a way that closer approximates the ideal rational agent in comparison to average consumers. On the other hand, the basic Pareto-improvement story seems

[14] Hayden, Grant M., and Stephen E. Ellis. "Law and economics after behavioral economics." *U. Kan. L. Rev.* 55 (2006): 629.

[15] Here I mean 'options' in the sense of 'a thing that may be chosen,' not the derivative instrument.

much more plausible in markets for products and services. Those markets are characterized by a high degree of preference diversity: the transactors in the market want different things. It makes sense to see market trading as *win-win* then– each person gets what they like (at least to the extent they can afford it). What makes trading possible in these markets are different preference orderings: I have some cash but would prefer to have a loaf of bread, and the baker has a loaf of bread but would prefer to have the amount of cash I have.

Financial markets, on the other hand, are characterized by preference homogeneity (most transactors seek money for future uses) but belief heterogeneity (they disagree about what current trades will pay off). What makes trading possible in financial markets is almost *never* a difference in preferences but is rather a difference in *beliefs*: everyone (with a few caveats discussed and criticized in later chapters) wants to make money, but people's beliefs about the future differ. In a simple, stereotypical financial transaction, for example, one person sells a stock because they think it will not increase in price and another person buys the very same stock because they think the price will go up. This transaction would still be a Pareto improvement because both sides want to transact, but at least one person will not get what they actually want (more money).[16] That suggests that even mutually agreeable transactions in financial markets might not generally help people pursue their own goals. The description of products within finance connects to the key features of financial markets that do not exist within traditional markets for products and services: the way in which they serve to enable growth, determine features of future product and service markets, and provide a mechanism for distributing risk.[17]

[16] Stephen Ellis. "Market Hegemony and Economic Theory." Philosophy of the Social Sciences 28, no. 4 (December 1998): 513-23.

[17] What the reader will notice is *missing*, however, is any set of normative bounds on how much risk (and accordingly, how much growth) ought to be pursued by any party to these transactions. This is the project taken up in earnest in Chapter IV.

Markets in general are supposed to achieve desirable results because they are, among other things, signaling mechanisms. Prices are traditionally viewed as signals of value and suggestions for further action– increasing price indicates increased demand or falling supply for a given good; decreasing price indicates falling demand or rising supply for such products. It is by this mechanism that 'true' prices are 'discovered.' They are the amalgamation of what market participants are willing to buy or sell with respect to a given good. This mechanism is given the name 'price discovery,' and there is theoretically no end to this process, as the world constantly changes in ways that beget reevaluation.[18] All of this is theoretically true for both markets for products and services and financial markets, but *what* the prices tell us in the two situations might be distinguishable.

According to economist Friedrich Hayek, prices capture all the relevant (action-guiding) aspects of the information about markets, at least as far as behavior is concerned. In a decentralized market with many independent participants, price movements dynamically reflect the changes in preferences or beliefs of its constituent participants. Consumers who are willing to pay a given price receive a good or service; products which are most in demand generate the highest incomes for their producers. Once prices make the capacities and desires of market agents clear enough, those desires will motivate mutually satisfactory transactions.[19] Or at least that is what is supposed to happen in a market with rational agents– they engage in Pareto-improving behavior.[20]

---

[18] None of this entails that prices are ever 'correct' or 'perfect.' The market price of a product, service, or financial instrument is an equilibrium result of the preferences, beliefs, and actions of a range of people.

[19] This is an idealized mechanism often criticized on the basis of its reliance on normative Pareto improvements - see Ellis, "Pareto Optimality". It will suffice here for mutually satisfactory transactions to be regarded as generally good, even if they are not invariably good.

[20] On Hayek's view, the cumulative understanding of a given population will always exceed that of any individual agent or body. It is central to Hayek's view that any form of centralized planning will necessarily fail to have access to all relevant information. "Planning" itself is the real problem then, not just centralization. Hayek, "The use of knowledge in society," *The American Economic Review*, Vol. 35, No. 4. (Sep. 1945): 519-530.

In *Information in Society,* Hayek illustrates this dynamic using a commodity price.[21] In the story of the commodity tin, the idea is that prices incorporate relevant information. If I am a consumer of tin, and the price of tin goes up, this is supposed to signal to me that something has changed. I should conserve tin, find a substitute for tin, or consider going into the tin business myself. Importantly, I need not understand *why* this change has occurred: the information provided by the price mechanism helps me allocate resources without having to allocate resources to understand what is causing the change within the price mechanism. This mechanism is of normative importance because it facilitates investment decisions, which influences the structure of production. The role of entrepreneurs includes managing this structure, which produces other (future) goods.

In this market-pricing narrative, financial markets serve as a mechanism to efficiently and productively distribute *capital*— resources for the production of other products and services. Investors use their limited capital to invest in businesses they think will be successful. The prices of individual securities are seen as indications of a company's value, an indicator that reflects the judgment of the market's participants. Companies with innovative ideas are given capital by investors to streamline production, create products, expand business lines, acquire other businesses, etc. Those companies succeed or fail in realizing their ideas. When they succeed, consumers with the highest demand acquire the product, and investors get their money back plus some extra— some of which is presumably used to continue finding innovative ideas to invest in. This straightforward account of financial markets sees them as allocating monetary resources to uses perceived as high value. This rewards and encourages consumer-market innovation in both processes and products and services themselves. At the same time, consumers are supposed to

[21] A commodity is a unique product which has a set of financial products associated with it.

benefit by having access to ever-improving new technologies and products. The market enriches successful investors, providing incentives to stay in the finance 'game,' as well as deterring those who do not succeed.[22] No matter the economic system or method of distribution, there are tradeoffs to be made between growth and risk. Any potential avenue for new ideas, products, or services carries with it innumerable possibilities, some characterizable as successes, others as failures. After the fact, it is easy to point to projects that were "too risky" or "worth it," because after the fact, the project has either worked out or not.

Financial markets are diachronic, whereas markets for products and services are mostly synchronic. Perhaps what I buy (or do not buy) at the store today can shape which products continue to be produced or are discontinued, but there are famous examples to the contrary (where popular products are discontinued for other reasons).

Finance, much more clearly, shapes the future of what markets for products and services will look like. The simple example of the farmer I've used doesn't quite capture this element: the financial transaction described in which the farmer guarantees a certain price range for their crop is effectively an insurance policy taken out on a predetermined behavior (which we can assume to lack any new innovation; they will farm their crop the way it is usually farmed). That example is useful in its simplicity, but is also synchronic: we take for granted what consumers want and therefore there will be markets of a certain size in corn, wheat, barley, etc. The options trade is a distribution of risk but does not directly contribute to any 'pie changes' (changing features in what the economic pie looks like or how big it is).

[22] I'll extend the use of the 'game' metaphor for reasons I make clear below. Regarding the use of this term, I aim for the reader to import their background connotations of *both* board- or video games *and* game theory.

On the other hand, let us exemplify the future-changing aspect of finance: We can imagine an entrepreneur who has a great idea for revolutionizing some product market (instead of standing desks, they want to make laying desks[23]) but does not have enough capital to start this business. Ideally, other entrepreneurs exist too. Financial markets enable entrepreneurs (likely only some of them) to create something new– something that currently does not have data on how well it sells, shaping what markets for products and services will look like tomorrow. Parties with money (investors, funds, etc.) have the opportunity to select, according to some set of preferences and beliefs, and following some set of conventions, which entrepreneurial ideas seem worth investing in. At least some parties will make their selections in accordance with at least some non-quantitative values. This is, in part, due to the epistemically limited environment they are in. With a publicly traded stock, or even a longtime private company, there is more information to go off of: what revenue has looked like over some past time period, whether there is stable demand, how competent those running the operation seem to be, and what the rest of the competitive landscape looks like for a particular business. For early-stage companies though, there is less to go off of. There was no consumer demand for planes or cars before they were invented, and their introduction changed and induced additional demand for travel. Innovative products are difficult to predict the success of for this reason, and that is one major rationale for why early-stage investment is traditionally thought to be better done by private investors and firms who theoretically can handle the downside risk better than the average person who has fewer tools and resources with which to do their own due diligence. And this is the diachronic

[23] I attribute this fantastic idea to a frustration vocalized by Gary Klar on the trading floor of Goldman Sachs in 2018.

role that finance plays: judgements, some of which will necessarily be subjective, are used to make decisions about what will exist in future product and service markets.

There is a fundamental misconception underpinning the efficient market hypothesis, which is that efficient markets will uncover what an asset is truly worth, as opposed to what it really does, which is uncover what humans, at that point in time, truly believe an asset to be worth. Although philosophers can see the analog to epistemological literature clearly in the distinctions between truth, justification, and belief, a trader's strategy does not require true belief. It merely requires justified belief relative to the trader's model of the future world and their prediction of what others believe.

There are a couple of ways to think about risk distribution here. First, there is a distributed risk of sudden changes in the company's valuation.[24] This is a temporally interesting arrangement: investors give up money now in exchange for a claim on an uncertain future, and financial markets help distribute that future-oriented risk across many people rather than concentrating it in the hands of the farmer alone. In effect, finance converts long, uncertain futures into present, actionable resources. Such dispersion is also supposed to limit the concentration of negative consequences of failure: no single household or small cluster of families bears the entire weight of the collapse, and the shock is sufficiently spread out that it does not immediately destabilize the broader economy.

Functional Properties of Finance

Modern financial systems depend on both individual and institutional trust. For instance,

[24] This is instrumentally valuable to the company for several reasons relating to the nature of the secondary market, which will be covered in a later section.

people must believe, with a high degree of credence, that banks will reliably safeguard deposits (or that there is some backstop on their not doing so), that currency retains value (with limited volatility), and that market rules are enforced consistently (to avoid uncertainty and default risk in the future). Belief in limits to individual conduct, confidence in the function of various mechanisms, and the assumption that participants will either adhere to established rules or face penalties for wrongdoing are all essential.[25] Without these foundational commitments, key mechanisms like lending, investing, and payment systems would break down because participants would be fearful of loss or unfair treatment. Volatility, then, not just in the actual economy, but also in the psychological sense of distrust by the populace toward the financial system, undermines a society's economic system. Any consideration that poses a legitimate risk of eroding public trust in our capital allocation mechanism, regardless of whether it is currently functioning poorly or not, is worth considerable attention.

The following sections outline important aspects of how financial markets work. My streamlined depiction in this chapter is subject to several critiques, but this paper will mostly avoid complicating the simple narrative. Even this idealized mechanism can go wrong in many ways. However, it will be useful for demonstrating the ways in which financial markets do things that product and service markets cannot do: manage diachronic issues and distribute risk. The reader can understand it as a kind of glossary of financial products, agents, and institutions and as a primer on how the pieces fit together.

Later sections of this chapter describe the kinds of agents that interact in financial markets, the instruments of finance, and the institutional framework within which they are employed. Most explanations of finance start with those parts, but philosophers might benefit

[25] Markets for products and services also depend on trust, of course - will the customer pay the invoice? Will the supplier deliver on time? - but financial markets face these issues and more.

from starting with what I call functional properties. There are several attributes of the financial system that financiers view as valuable because their existence facilitates the normal functioning of financial markets, although they are not usually grouped together as a type. These functional properties are usually understood to include speed, liquidity, transparency, and price accuracy. These properties are often understood as things to be maximized: no financial system could be too fast, too liquid, or too transparent,[26] and are usually understood to be interrelated and mutually supporting: e.g., speed helps facilitate price accuracy; price accuracy facilitates liquidity. These properties can be understood in two ways. First, as constitutive features of financial markets, whose absence inhibits functioning, and secondly, as features that either do not exist or are less central in markets for goods and services.

Financial markets allow people with surplus money (e.g., capital beyond what they require for immediate consumption or known future liabilities) to allocate those resources to others who seek to use them for potentially productive purposes. Those who supply capital usually do so in expectation of a monetary return on their investment. At the same time, financial markets allow economic agents engaged in productive activity to manage and limit their exposure to risk by distributing some of it to others. Agents willing to supply capital and agents willing to assume another's risk (who can sometimes be the same agents) are incentivized to do so: they are rewarded when they allocate capital and price risk effectively (with profit) and penalized when they do this poorly (with losses in money or value of their ownership claims).

Speed reduces the time it takes for capital to flow to productive (and unproductive) uses, and so for risks to be realized or mitigated. Faster markets allow market participants to adjust positions in response to new information and so lower the friction that delays trading decisions.

[26] I think this is wrong, but it is a popular view that is often a motivating force for both commercial innovation as well as regulatory guidance.

Speed enables markets to act nearly in real-time, and for people to trade on newly gained information. A trader can instantiate their views rapidly, which has two effects. On the positive side, prices adjust quickly to new information, moving toward a state that better reflects the assets' fundamental values (i.e., what the assets are objectively worth based on available information). On the downside, the same rapid adjustment means that the value of the assets a trader holds may fluctuate quickly, increasing exposure to short-term volatility.

Liquidity measures how easily an asset can be converted to cash.[27] It is related to speed because ease is related to quickness. Highly liquid markets allow investors to enter or exit positions quickly, transferring risk from those who wish to avoid it to those willing to take it on. This is critical for distributing risk over time: if an asset were illiquid, an investor would be forced to bear risk until a buyer appears, potentially for a long period.[28] Liquidity is influenced by factors such as the number of market participants, asset fungibility, and regulatory frictions. For example, barrels of oil are highly liquid because they are standardized, fungible, and widely traded; houses are far less liquid because they are heterogeneous, regulatory-intensive, and involve fewer buyers on any given day.[29]

Transparency is an interesting value, as it is the one I think is most under-theorized. Transparency is important because investors should be able to know what they are investing in.

---

[27] Many people use this word to imply that 'liquid' assets are also those that one can convert to cash without significantly affecting its price, but that is dependent on other factors which do not hold universally true.

[28] One concern people raise is with corresponding assets that have a liquidity mismatch. Bond ETFs are one such case: A bond ETF is an exchange-traded (e.g., tradable at any point during any day when markets are open), which is 'liquid' in the sense that someone can sell a share of it to someone willing to buy it at any time. On the other side, this is an ETF that bundles *bonds,* and bonds are traditionally an asset that is less liquid than exchange-traded stocks. During a period of bond selloffs, then, the concern is that those with bond ETFs in their portfolio might overwhelm the market's capacity to sell bonds accordingly, which will lead to sudden price movements and volatility (often called "price gapping") and risks extreme downward price pressure.

[29] Interestingly, the COVID-19 pandemic caused an interesting situation that was widely reported where oil futures contracts were briefly trading in negative dollars. Storage for physical oil was in short supply because consumers were using less than expected amounts of oil. Those holding futures contracts upon expiry are obligated to take delivery of actual oil barrels. To avoid taking physical delivery at a loss (because they would need to pay for a place to put the oil), traders were willing to pay others to take the contracts from them.

There are obvious dangers to a lack of transparency, including the potential for outright fraud. This value receives a lot of attention in the legal literature, as it is of obvious interest to regulators, from run-of-the-mill consumer protection reasons to questions about where to draw lines among the regulatory burdens for different actors, e.g., small-to-medium versus large banks.

Price stability, by itself, tells us nothing about whether a price is accurate, if accuracy is understood as reflecting the underlying economic reality of the firm or asset being priced. A stable price may simply indicate that market participants agree with one another, not that their shared valuation tracks anything objectively true about the asset's productive capacity or future prospects. Stability denotes consensus rather than closeness to truth.

This misunderstanding underlies a similar misconception between what people want to consider normatively good 'investment' and bad 'speculation.' This misunderstanding imagines legitimate investment as being disciplined and guided by objective facts (e.g., EBITDA, quarterly earnings, etc.). 'Legitimate investment' is done within reason, then, and speculation is not. Speculation is seen to be motivated by greed, sometimes done with ignorance or in stupidity, and is seen as corrosive for those other legitimate investors. But both activities look similar in a certain light: they both necessarily involve beliefs about uncertain future outcomes, formed under conditions of incomplete information, and it is the shift in those beliefs (not what motivates that shift) which moves prices.

Next, I will describe the things in which financial markets deal ("assets"), the people who do things with the things ("financiers" and "market participants"), and the places in which they do them ("financial institutions").

Money and Banking

Money is a thing that facilitates transactions and serves as a store of value. Holding money, as opposed to buying other goods or assets, has advantages and disadvantages for agents. Money is useful because it is liquid, fungible, divisible, transportable, and durable. The way money moves through an economy depends on the total amount of money in circulation, and on the interest rate environment, and inflation. The interest rate environment can be understood as the price of holding money, and rates of inflation can be understood as the rate at which money loses purchasing power over time. Central banks control the amount of money in circulation and parts of the interest rate environment directly, in order to affect inflation and the broader interest rate environment indirectly.

Central banks are given the power to manipulate to (hopefully) achieve macroeconomic stability. In broad strokes, the tangible features of macroeconomic stability include inflation being low and predictable, high employment rates (and associated low unemployment rates), and steadily growing output. The value of macroeconomic stability is easy to see if we look at instances of *instability*: in the wake of the 2008 financial crisis and during the COVID-19 crisis, uncertainty had direct and indirect negative effects. The instability of those periods eroded required conditions for beneficial (e.g. win-win) exchange like predictable prices, reliable access to credit, and confidence in counterparties' abilities to meet obligations.

These conditions relate to the functional properties of liquidity, speed, and transparency. Perry Mehrling, for example, emphasizes the role of money markets and financial institutions in generating liquidity and credit availability. He argues that central banks are the ultimate 'dealer' in financial markets rather than a mere supplier of money from which transactions and

investments occur.[30] The major views on monetary policy can be split into two umbrella categories: non-discretionary views like the gold standard or monetarist views, and discretionary views like Keynesian demand management policy. My project does not require adjudication among these views, just that we accept that money is a real phenomenon with economic uses (e.g., investment, income, spending) and consequences (e.g., bankruptcy, loss of funding power, inflationary impacts).

The Federal Reserve (the U.S. central bank) sets interest rates primarily by adjusting the federal funds rate, which is the rate at which banks lend their reserve to one another on an overnight basis.[31] This is deemed the 'risk-free' rate, as it is the rate at which major banks lend to one another, institutions that they estimate to have effectively zero risk over the timeline of the loan (which again, is literally over one business night into the next business day). By setting this short-term interest rate, "the Fed" (as it is affectionately known) affects borrowing costs

[30] Perry Mehrling, "Financialization and its Discontents," *Finance and Society* 3, no. 1 (2017): 1–6.

[31] Why would this be something that banks do, you may wonder. There are two things to explain here: why banks lend overnight, and why the federal reserve has a special relationship with banks and gives them privileges not afforded to everyday citizens. First, banks engage in overnight lending for a handful of reasons. 1) Capital requirements are regulations which require them to always keep a certain amount of cash on hand. Banks want to be right at this threshold as often as possible: if they are way higher, that means they could have lent more money. Of course, they cannot be lower, lest they face fines and other penalties. This means that if throughout the course of the day's events, a bank is a little under their capital requirements, they have a need to borrow cash. If a bank has found themselves over the requirements, they would like to loan that out for more than the 0% interest it will collect in-house. 2) Related to 1), but more generally, banks are in the business of managing liquidity. Short-term funding needs (the need for cash now, but not for a long time) is something that a bank might need for reasons unrelated to capital requirement regulations. This might relate to settlement of contracts, loans, etc., particularly in unexpected exogenous shocks that cause surges in customers withdrawals or repayments. 3) The Federal Bank, having the aim of providing stability to the financial system, gives banks an opportunity to borrow directly from them in times of need (to prevent them from going bust). This is called "going to the Fed's discount window." But doing so signals distress at a bank (psychologically, the story goes something like: Since the Fed is a lender of last resort, a bank must be in serious trouble for it to justify borrowing from the discount window. They must be facing a crisis. As I will get into detail more later, bank psychology is unique in its ability to self-manifest outcomes in a way that does not apply to normal businesses). 4) Interest rates, which can change over time, might change in such a way (or one might predict that they will change the following day) to behoove banks to lend overnight in the lead up to some longer or larger trade the following day. These reasons also relate more generally to the reasons for repo markets to exist amongst institutions other than commercial banks. Second, the federal government recognizes the importance of the role banks have in creating money via lending. The Fed's relationship with banks is premised on managing systemic risk and supporting economic stability. The nature of this relationship legitimizes regulation of banks. For more on this issue and its alternatives, see James and Hockett, *Money from Nothing*.

throughout the economy: when the Fed rate goes up, mortgage rates will go up, and when the Fed rate goes down, so too will mortgages.[32] Lower interest rates makes borrowing cheaper, encouraging businesses and consumers to take out more and bigger loans.[33] Conversely, higher interest rates discourage borrowing. The Fed has market power in the market for money, which it is pledged to use for the benefit of all, not just itself.[34]

The Fed can also alter the money supply through open market operations like buying and selling government securities (see below). These are U.S. treasuries (aka T-bills) which are bonds issued with the full backing of the federal government. The market in T-bills can serve to inject or remove liquidity from the banking system.

Fractional Reserve Banking

Most people understand what banks do from a consumer's perspective: You deposit your paychecks into a bank, earn a (very small) amount of interest on your savings account, and can use a debit card and to write checks tied to the money in that account. The bank earns its keep (the difference in the lowly rate they give to your savings and the higher rate they can earn lending your money to others) by keeping your money safe and keeping track of how much of it you have. This can go wrong in various ways, if the bank uses your money to make bad loans,

[32] Interestingly, there is some dislocation on the intensity of these movements. This is not straightforwardly the impact of any one thing. See Joe Weisenthal and Tracy Alloway. "Why Mortgage Rates Went Up After the Fed's Big Cut." Odd Lots. Bloomberg, October 21, 2024.

[33] Thinking through your experience as a consumer will help here: When interest rates were low post-'08 housing crash, it was easier to justify taking out a loan at, say, 2.5% APR for a large appliance than it is at the time of this writing, in 2025, when a consumer loan might be to the tune of 5% or higher. Similarly, if readers of this footnote have any experience buying a home, the expected mortgage payments depend both on the price of the home and the interest rate one can expect to have on their mortgage. As an unrealistic example: If your financial status requires that your mortgage payment be no higher than $1,000, and interest rates are at 2%, you can afford to take out a 30-year mortgage on a home purchased at $270,549. If rates double to 4%, you can only afford to take out a home worth $209,461 at time of purchase. If rates are 6%, you can only afford a house for $166,792 and so on.

[34] The system described here is a discretionary one with a fiat currency in currency in circulation. In non-discretionary (e.g., gold-standard) currency economies, the supply of money remains relatively constant, allowing market demand for cash to determine interest rates.

does not keep track of how much money you have, or otherwise mismanages it. The perception[35] of mismanagement can cause a bank run, where depositors all rush to get their money out of the bank at the same time.

Of course, the bank lends some of your money out, charging interest on loans issued. Banks operate under a fractional reserve system, meaning they are only required to keep a fraction of the money you give them in reserve (in case you want to withdraw it). Once a bank lends out money, it is not the case that this money evaporates. A bank may issue money to a business, who then deposits some of it into a different bank or sends money to other businesses for products and services– meaning money winds up in the bank accounts of those other businesses. When that other business deposits money into their account, their bank does the same thing over again. This process iterates, and in this way, banks *create* money (or, minimally, 'purchasing power') from the promise made by a loan recipient to repay them. In fact, the seemingly trivial service rendered (keeping money safe and sound and withdrawable) is a highly leveraged institution that allocates creditworthiness, mediates trust, and generates purchasing power.

Why Think Finance Should Exist at All?

The aforementioned farmer produces something of obvious social value (people need to eat), they hope to earn a living from it, and the process of doing this faces a certain set of

[35] Banks are uniquely sensitive to public perception because confidence in their solvency is itself a condition of their continued existence. Unlike banks, negative public attitudes towards a firm that produces consumer goods or services do not directly threaten that firm's ability to continue operating. Coordinated consumer boycotts, for example, affect firms indirectly, by lowering future demand or influencing investors' expectations. Shifts in public sentiment may move a firm's stock price in the short term, but such shifts in public sentiment reflect agents' beliefs about possible future behavior rather than constituting an immediate causal threat to the firm's survival. People may complain about Apple getting rid of a useful port on a particular device, but sales could still be up next quarter. In the case of banks however, changes in public confidence can themselves precipitate and perpetuate bank failure by causing a bank run.

uncontrollable risks that cannot be eliminated related to nature. Issues with weather or pests have both direct effects on whether a given crop can survive, and indirect effects on the prices that those crops will fetch come harvest. This process requires time, land, labor, and usually capital to bring to market.

What matters for our purposes is the diachronic structure of this activity. Farming is not merely labor initialized at a moment in time.[36] It is a months-long process which inevitably requires the management of uncertainty and risk. A farmer could bear this risk solitarily, hoping each year that their soil is fertile and that their crops will grow. But eventually, and more regularly in an era of climate change, a bad season will come. To continue being a farmer over time, they need ways to spread or transfer the irreducible risks inherent in their work.

Finance precisely provides these kinds of risk management tools which allow the farmer to trade some of their exposure for greater predictability. Each of these tools is a transaction with a counterparty: another agent is willing to accept some of the farmer's risk because doing so aligns with *their* self-interest (presumably including something like 'holding a diversified portfolio of assets').[37]

Important Financial Products

Securities are a kind of asset identified via the *Howey* test established by the Supreme Court in 1946. They are investment contracts where there is an investment of money in a

[36] A Marxist account, for example, captures the exploitation or compensation of labor at any given moment but not the fact that the activity unfolds under conditions of uncertainty that must be managed across long timeframes.
[37] A quick elaboration on this motivation: A good financier won't want to but all of their proverbial eggs in the same basket. Most investment managers pursue a strategy of diversification. This is good for their long-term success. For example, if they were to suppose that the S&P 500 is the best place to put their money, whether their portfolio is 'up' for a quarter will depend precisely and only on whether the S&P 500 index has gone up or down. What if (let's say, because the AI bubble has burst and tech stocks are down) the S&P 500 goes down? They would prefer to have an asset that is uncorrelated with the S&P 500 that softens the blow. Losses from one position are offset by stability or gains in another.

common enterprise, with the expectation of profit to be derived from the efforts of others.[38] Securities exist ostensibly to facilitate getting money to companies who don't have it now but expect to have even more of it in the future. They are generally more liquid assets, compared to owning a portion of a private company or a piece of fine art.[39] People put money into securities because they have money now, but would prefer more of it later (the sooner and the more the better, usually).

To make this analysis more concrete, let us use our example of a farmer. Suppose that the farmer wants to expand and start a company which turns some raw commodity of theirs (let us say, apples) into a set of consumer products (apple jam, apple juice, etc.). They do not have enough money in their savings account to pay for the startup costs, but they reasonably expect future revenue once the crops are grown and processed into apple-based products. Engaging with a financier will have the effect of giving the farmer an advance on anticipated future success.

There are numerous ways this could happen with a variety of different financial assets: A farmer could take out a second mortgage on a property to get capital, they could form a limited liability company and sell a stake in this partnership to a limited partner, they could issue bonds secured by future product revenues, or they could sell shares in their business once incorporated. Many financial products make better sense the larger a company is, and the greater its scale. Eventually, if the farmer's apple-products are well-received, they could 'go public'– issuing shares of their company to the public to fund expansion into new markets or to generate a mechanism which allows them to pass on the business to new management.

[38] There is a varied array of critiques of this definition, many of which are covered in Georgakopoulos (2017). His major critique is an ontological one that this definition is "a list of terms rather than a principled definition." p.19.
[39] Fine art may be a controversial asset, but the moral salience of this will not be discussed here.

Perhaps it was a coincidence that the farmer found the first financier they worked with, or perhaps their relationship came from personal networks. There are systematic ways for businesses to broadcast opportunities for investment. The securities discussed below– stocks, bonds, mortgages, futures, forwards, options, and swaps– are all variations on this core theme: they move money across time and distribute risk across counterparties. And the central reason any given counterparty is willing to take on another's risk is the prospect of making money by doing so.

What this does not tell us

I should note that while these tools are mechanisms for distributing risk, they do not tell us anything about how much risk we *should* be taking. They also appear to be mechanisms which, if not paid normative attention, could facilitate exposure to *limitless* risk. This is, in some sense, what went wrong with the credit-default swap mortgage back securities market leading up to '08. Without a normative sense of how much risk we should be taking on, the mechanisms which we can use to distribute it have no upper bound or limit outside the confines of the law– and there are currently no laws in place that conceive of risk in this way.

Selling of one's assets can occur for another reason too. The farmer can act in the capacity of an 'entrepreneur,' creating new products. Perhaps their experiences have inspired them with other ideas, meeting their customers' needs for, let us say, a receptacle to carry apple-based products. The farmer has something of value– ownership of their apple-products enterprise– that they want to convert to capital to both reward themself for their hard work, and to fund (perhaps only partially) their next big idea. So later, after the fact, there is a reward incentive: people who want to sell part of their business to get out of it want to be rewarded for

having run the business well. The converse of this is that by owning parts of business, people are incentivized not to make themselves worse-off (hopefully), because the value of the thing they own will go down.

A differentiating factor in how securities are regulated and how nonfinancial products are regulated is that nonfinancial products are largely regulated by contract law, which follows the principle of *caveat emptor* (buyer beware).[40] The principal logic here, as delineated by Georgakopoulos, is that within the realm of non-securities, the purpose of purchase for nonfinancial products is highly individualized and cannot be easily anticipated by the seller. One person may purchase bananas as green as possible, with the intent to mix them into smoothies, whereas others may want to buy ripe bananas because their intent is to make banana bread.[41] In contrast, securities law proceeds from a seller disclosure perspective, in which corporations issuing securities bear the burden of disclosing all possible risks.[42] This regulatory framework has implications for how financial markets distribute and manage risk over time. By requiring disclosure, securities law helps facilitate the transfer of risks from those unwilling or unable to bear them to those willing and able, creating value by facilitating informed capital allocation.[43] There are regulatory requirements which are designed with the intention of providing would-be investors with the information they need to make an informed decision on whether to invest in a given enterprise. We will delineate between how this dynamic operates in initial public offerings

[40] Of course, the principle of buyer beware is supplemented by various consumer protection regimes and business regulation which has the intent of making the marketplace for nonfinancial products less hazardous for buyers.
[41] Georgakopoulos uses tomatoes as an example: some want unripe tomatoes to can, other want ripe tomatoes to eat, and still others want overripe tomatoes for use as "projectiles for artistic or political criticism."
[42] The implications of this different regulatory structure will be further discussed in Chapter V.
[43] This has given rise to a problem Matt Levine calls 'Everything is Securities Fraud.' The story here goes that if something bad happens which results in a stock's price dropping, investors can sue that corporation for not disclosing (or not adequately disclosing) the risk of the bad thing happening. Matt Levine, "Everything, Everywhere Is Securities Fraud," *Bloomberg.* June 26, 2019.

(IPOs) of companies, where risk is concentrated at the point of issuance, versus the trading of publicly traded stocks in the secondary market, where risk is redistributed among a broad pool of investors over time for varied reasons. In both cases, securities serve as tools to manage known and unknown risks, and they connect present capital to future productive activity in the real economy.

The most common 'security' people think and talk about are publicly traded stocks of major companies like Apple (ticker: APPL), Microsoft (ticker: MSFT), and Google (ticker: GOOG).When you purchase a share of APPL, the traditional logic goes, you are investing in a common enterprise (the enterprise of whatever Apple is doing), you expect to profit rather than lose money, and the effort that will beget that future world where this share is more profitable will come from the efforts of the employees of Apple.[44]

One persistent misunderstanding of stock markets is that companies somehow receive money when *any* stock is sold. The only time that a publicly traded company receives capital from shares is at the point of share issuance: the time at which the initial or secondary stock shares are sold to investors. The reader might then wonder, why have secondary markets at all: why allow anyone to sell shares after an IPO? There are various time-horizon and liquidity reasons for this, but one simple way of seeing the value of secondary markets even if companies aren't getting any new capital when a previously-issued share of stock gets sold is to say that keeping track of their valuation is useful at key moments in the future, and stock price changes give those running the company important (albeit sometimes distracting or badly-incentivizing) signals on how various business decisions are perceived. These moments include future stock

[44] Here we usually ignore the 'efforts' of consumers in purchasing Apple's products or continuing to use their services.

issuances (the company wants to slice up more of the pie that they previously owned, or 'dilute' shares by issuing more of them [i.e., creating and selling additional shares to raise new capital, which is only attractive to the company if the existing secondary-market price is high enough to justify the dilution]), if they receive an offer from another company to acquire them. If they use their shares as collateral for borrowing or compensate employees using stock or stock options, it is not the case that they would be unable to do this without being publicly traded but having a stock price that is continuously updated makes these activities easier, faster, and more trustworthy.

People also tend to use this factor to attribute net worth to individuals: Jeff Bezos does not actually have trillions of dollars in a bank account right now. He owns a large portion of a company which is worth a lot of money, but if he tried to sell the entire amount of it that he owns, people would probably wonder why he is selling it and the value of the company will likely go down.

There are things other than publicly traded stock that are considered securities. Mortgages, for example, should be familiar. A mortgage is a loan secured by real property. The mortgage lender has the right to take the property if the borrower fails to repay. A single mortgage is not itself a security, however, because it, in isolation, is a private credit contract rather than a tradable investment instrument. Mortgages can be transformed into securities by way of '*securitization*': many individual loans can be pooled together and sold as 'mortgage-backed securities' or MBS. These kinds of securities will forever live in infamy, having been largely responsible for the 2007-08 financial crisis.

In some cases (and mortgages are a frequent target) people disagree about whether an asset 1) should be counted as a security at all[45] or 2) should be securitized.[46] Objections from 2) often appeal to moral disgust[47] or to a reactionary worry that securitization makes the underlying activity dangerously abstract or systemically risky.[48]

Mortgages are often the most valuable asset that people own, and they are also the most democratically distributed asset. Mortgages are important for the functioning of financial markets because of the risks that concentration poses. If a bank holds many mortgages (or MBS), many defaults (or, conversely, a wave of prepayments), can create outsized losses or disrupt predicted cash flow. This can threaten a bank's stability.

When a bank issues a mortgage, they are exposing themselves to the long-term risk that a borrower will be able to pay. Securitization spreads the risks of that future across many investors instead of leaving a single bank exposed to many years of uncertainty. One thing that makes mortgages unique is the way in which they are shaped by interest rates. Unlike other loans (imagine a small business loan that the farmer might take out to start their apple products business), mortgages have a pay-back risk: that is, the risk that homeowners will pay back their mortgages all of a sudden due to selling of the house or refinancing when interest rates are lower.

---

[45] This is an argument sometimes made about cryptocurrencies (because designating them as such would have the outcome of making them more heavily regulated by the SEC), and homes (because designating them as such has a corrosive effect on how humans conceive of housing).

[46] This is of course an argument made about mortgages post-'08.

[47] Viaticals– whereby someone with a life insurance policy can sell that policy to a financier and collect the money entailed by that policy upfront (e.g., before they die) is one financial product that receives particular disdain. This is because the financier receives higher profits the sooner the counterparty dies. Even without any wrongdoing in bringing that outcome about, the alignment of incentives can look exploitative since sellers are often ill people with large medical bills and limited bargaining power.

[48] You see this deployed by way of analogy to '08 MBS: "Look at this thing which is behaving mechanistically the same as mortgages transformed by way of aggregation into a mortgage-backed security. It's the next '08 crisis waiting to happen." Most often this occurs with auto-loans, likely because they bear conceptual symmetry to mortgage loans and are more readily understood by non-finance experts. See: Jeff Kirk. "Subprime Auto Loans: The Rising Menace of Wall Street's Latest Darling." Journal of Consumer & Commercial Law 18, no. 2 (2014): 72-75.

Prepayment risk means that the value of the mortgage depends not just on whether borrowers will pay (assessment of their individual creditworthiness), but on whether they will choose to keep paying into the future at all (an assessment of their likely future actions, and a prediction about future states of the world). A prepayment would change the expected cash flow of a bank. Banks anticipate receiving interest payments over 15-30 years, but prepayments reduce the actual interest income collected by the bank, meaning an asset has generated less profit than predicted. The sudden influx of capital in the event of prepayment, may also come at a time where putting that money into a new mortgage may beget a very different interest rate. If interest rates have gone down since the prepaid mortgage was issued, the bank will not be able to have the same level of predicted profit by reinvesting the cash in new mortgages.[49]

Bonds are securities (specifically instruments of debt) that entities use to fund projects. Bonds are structurally similar to loans: they provide an entity with money upfront, which will be repaid on a regular basis with a set amount of interest. There are two major kinds of bonds: corporate bonds and municipal bonds. Both have the following structure: the company or municipality wants to do a project, e.g., build a new factory or sewage treatment plant, and they need money for it now. Instead of asking a bank for a loan (since it would be a very large sum), the company/municipality issues debt: they, in essence, divide the amount they would need from a loan into slices, not all of which are necessarily equal. They then sell these slices to buyers with some attached 'yield' or interest rate. Over the lifetime of the bond, the company repays the slice of debt. If something happens to the company/municipality, their assets are sold off and the debtors are repaid according to some delineation of seniority (if everyone on the list gets paid out, stockholders are left with the remaining sum).

[49] This is not described with the intention of getting the reader to pity the bank, but rather to explain an unintuitive risk that banks face regarding mortgages.

Bonds differ from traditional loans. While stock represents partial ownership in a company, with rights to dividends and voting, a bond represents a slice of a debt obligation. Bonds divide up the lender's claim on repayment into tradable pieces. Bonds are more liquid than traditional loans because they can be bought and sold on an exchange. This tradability also means bonds typically have a more diversified set of holders, whereas a loan is usually issued by and owed by a single bank.

In ideal cases, this diffusion of bondholders entails a diffusion of risk that smooths out the volatility that would otherwise be concentrated in the banking system. A single bank holding a large loan is exposed to the borrower's business cycle: If the farmer's apple product revenues dip or a high-capital cost product underperforms, that bank alone must absorb the shock in the event that the farmer's business defaults. With bonds, the same fluctuation is absorbed by a wide set of investors with diverse portfolios, making the financial system more resilient. This is one reason bond markets are seen as stabilizing relative to bilateral loan markets: they spread credit risk, reduce the likelihood of concentrated failures, and allow borrowers to access capital without placing concentrating potential strain on any single lender.

'Derivatives' is a catchall term used for securities that *derive* from other securities, usually referred to as the underlying asset. Derivatives involve "the future exchange of cash flows whose value is derived from or based on an underlying value."[50] Some derivatives are traded on a centralized exchange, whereas others are called over-the-counter (OTC) derivatives. Derivatives that are traded on a centralized exchange include options and futures (explained below). OTC derivatives include bespoke contracts negotiated directly between two parties, allowing them to tailor the terms, risks, and time horizon to their specific needs. Because these

[50] CFA Institute, CFA Curriculum 2025, Level I, 2025.

instruments are customized and not traded on a public exchange, they can be more difficult to assess the value of and might be more vulnerable to counterparty risk (the risk that the person you are trading with won't or can't hold up their end of the trade). As a result, derivatives are the most impactful mechanisms for distributing risk, and potentially generating non-essential risk (e.g., there is no underlying risk in the real work based on real economic enterprise to distribute).

What real purpose can derivatives serve aside from speculation? In short, they can be used to manage known and unknown risks. Issuers of securities (e.g., corporations) use derivatives to try to manage the specific risks, usually those based on timing issues. E.g., if a corporation knows that it has debt coming up due and wants to 'lock in' a specific interest rate (realizing that interest rates might climb higher by the time they refinance, it will use derivatives. Imagine that a company has $100 million in bonds maturing in 20 months (the amounts here do not really matter). To refinance, it will issue 10-year bonds at the prevailing interest rate and use the proceeds to pay off the bonds at maturity. This has the mechanics of 'refinancing' a mortgage, which readers may be more familiar with. If interest rates go down, future borrowing costs will be lower. If interest rates go up, future borrowing costs will be higher. They do not mind paying a fraction of what they would stand to lose to prevent the latter from happening, and they are willing to forego some of their potential savings that the possible world with lower interest rates would entail. Let us suppose that current interest rates are around 4%. They can use an interest rate swap to achieve this objective: They enter into a swap agreement that stipulates that in 20 months, they will pay 4% (daily; annualized) plus some spread (a fraction of a percent, which is the fee the counterparty receives for their trouble) to a counterparty (likely an investment bank) to receive SOFR (the secured overnight financing rate). In exchange for the fixed rate of 4% plus a little bit, the company receives SOFR daily, which will float around

according to interest rate movements. If SOFR falls below 4% plus a little bit, the company will have lost out on possible savings. They are willing to open themselves up to this possibility because they can avoid ever paying more than 4% plus a little bit.

They have used a derivative to 'lock in' a specific interest rate. Similar stories can be told about a company that needs to purchase raw materials before it can sell finished products. A derivative can hedge the risk of input price fluctuation. Likewise, a retailer purchasing goods in one currency but selling those goods in another currency can use currency derivatives to manage the risk of exchange rate changes. In these cases, derivatives help lock in costs or revenues, reducing uncertainty and allowing counterparties to plan.[51]

These derivative uses are obviously value-additive activities in ultimate service of the real economy which are not themselves products or services.[52] The value added by derivatives in the context of investors' activities is less direct but can still be valuable. Investors might want derivatives that change the concentration[53] of or duration[54] of their exposure profile. An investor might have a specific view (or 'trade thesis') about the market that they lack the cash to actualize

[51] Not to mention the risk of changes in tariff policy they might need to mitigate all the while. This particular risk is top of mind as of writing in 2025.

[52] The reader might object that the service being rendered here is insurance, broadly construed. This might constitute a possible objection to my view that I could simply reconfigure this entire project as being about how finance is all essentially an insurance service. But I think this would do a disservice to our analytical perspective, as I'll address further in Chapter IV. There is also a linguistic issue with using the word 'insurance,' as most insurance providers can do what they do because of their use of financial products (discussed in main text below), leaving us open to an objection of reversing the explanatory order (treating 'insurance' as upstream rather than downstream from finance) and making a category mistake (conflating contractual risk-transfer with market-based risk distribution).

[53] A pension fund wants to have enough money in the long-term future to pay out pensioners. The tech sector makes up roughly 30% of the S&P 500. They worry that by investing broadly in the market, they are concentrating risk in the technology sector. So, they want to get exposure in something they think is uncorrelated with the S&P 500. By this, we mean that if the S&P index goes up, they want something that won't go up for the same reasons (and vice versa, if the S&P goes down). They can't purchase actual bushels of wheat or barrels of oil or any other physical commodity, but they might want to buy futures in commodities they determine to be uncorrelated with tech stocks to counteract the concentrated nature of the S&P 500.

[54] 'Duration risk' is a term used to describe risks due to interest rate changes. Again, let's use an example from a theoretical pension fund. Let's say a pension fund holds some 10-year treasury bonds. The price of bonds is inverse to their yield (or interest rate). So, if interest rates go up, bond prices go down and if interest rates go down, bond prices go up. A pension fund wants to mitigate the risks of interest rate changes over the next 10 years. So, they might use an interest rate swap.

using the actual securities.[55] Although this increased exposure can increase risk, it also helps facilitate accurate pricing. This is indirect because buying or selling options only indirectly influences the price of a security.

Why is any of this necessary? Couldn't a company - say, our farmer - simply take out an insurance policy with a 'normal' insurance company? Part of the answer here is that derivatives are a tool that insurance companies themselves often use in order to create insurance policies. Further, using derivatives enables a larger volume of commercial activity, and investors assign higher valuations to non-financial firms that use them—presumably because derivatives allow both insurers and operating companies to offset risks that would otherwise remain on their balance sheets.[56] There is a large volume of literature which examines the reasons behind corporations' use of derivatives, and most use derivatives to smooth out known and unknown risks related to taxes,[57] periods of possible fiscal stress,[58] and underinvestment.[59] Next, I delineate the four most common kinds of derivative products.

---

[55] Imagine here publicly traded stocks that have high prices. I might feel confident that Ulta (the cosmetics store) is undervalued and will go up by 10% in the next three months, but I don't have enough cash to buy very many shares of this stock. As of writing on September 8th, 2025, a single share of $ULTA costs around $520. If I have $1,000 to spend on this trade thesis, I could purchase a single share of $ULTA and wait for it to increase in price (since two shares costs more than I have to spend). Alternatively, I could buy a set of call options (each option is for 100 shares of $ULTA, which is standard for exchange-traded options, and imagine that each option costs $10 per share, which is unrealistic because options contracts are usually cheaper the sooner they expire and the closer they are to the expected price at expiry). This means I could buy up to 10 contracts with my $1,000 (I decide to buy 10 options contracts, which gives me exposure to 10 * 100 = 1,000 shares, or 1,000 times the number of shares I could get exposure to if I transacted in cash). If I purchased one share in cash and ULTA rises 10% over the next three months, I would receive 520 * 0.10 = $52 in profit. If I spent $1,000 on options contracts with a three-month expiry and a strike price of $540, I would make 1,000 * (572 − 540) − 10,000 = $22,000 in profit. This example illustrates the degree to which I can get more leverage using derivatives. This is not necessarily a good thing. In a world where I am totally wrong and $ULTA decreases by 10% over the duration of my trade, not only will my options expire worthless, but I will also have spent up to $1,000 on now-worthless contracts. In a world where I only bought one share of $ULTA, I still have $480 in cash and a single share which is now worth $468. Comparing these two worlds, a 10% decrease in the price of the underlying security has either led to a 100% loss of my $1000 or a mere ($1,000-($520-$468) = $948) 5.2% loss.

[56] Milken Institute, "Deriving the Economic Impact of Derivatives: Growth Through Risk Management." 2014.

[57] Donohoe, 2012; Graham and Rogers, 2002

[58] Smith and Stulz, 1985; Tufano, 1996; Bartram et. al., 2009

[59] Froot et. al., 1993

Forwards are OTC derivatives that take the following shape: a buyer agrees to purchase some underlying asset for future delivery from a seller for an agreed upon price. This structure is precisely what a farmer might use to stabilize the future: the farmer commits now to selling next year's harvest at a fixed price, ensuring predictable income, while the buyer accepts the risk that market conditions or crop yields may change by the time delivery occurs. Futures are types of forwards that are standardized such that they trade on exchanges. They are forwards with specified standard contract sizes, standardized delivery dates, and a clearly defined underlying asset. "Delivery" refers to the date when the underlying asset or its cash equivalent is to be exchanged according to the terms of the contract. Most of the time, futures are 'closed out' with cash being exchanged, rather than involving the cumbersome physical delivery of the underlying asset. Importantly, futures have daily settlement whereby the future is "marked to market" and margin[60] requirements are adjusted.[61]

Let us imagine the farmer again. In a forward contract, the farmer and a financier privately agree on the exact quantity, quality, and delivery date of next season's crop. This might be a resource intensive negotiation requiring that the farmer hire someone to represent their interests or simply spend the time discussing this trade with the financier. In a futures contract, by contrast, the farmer would use an exchange-traded corn future with standardized contract sizes and delivery months, and the farmer's account would be settled daily as prices changed. The forward reflects the farmer's specific situation, which if unique might mean that using futures is not the best product for their needs. The future reflects the standardized, liquid version

[60] Margin is the amount of cash each counterparty must keep with the exchange as a means of alleviating credit risk– the risk that you agreed to buy something later that at that future time you cannot afford. These margins are standardized but often differ based on exchange and underlying asset as they are influenced in part by regulations designed to prevent exchange failure.

[61] The settlement price used here is based on what the exchange calculates to be the average trading price over the course of the trading day. This differs from other settlement conventions which use the price at the end of the specified trading day (at "market close").

of the same idea, and might be 'cheaper' to use based on the resources required to negotiate and set up a specialized deal.

Options are exchange-traded *rights* to buy or sell an asset at a specific time and price. The buyer decides if the counterparties will transact, and the seller of the option is obligated to transact based on the buyer's wishes. For example, imagine an investor buys a call option on a stock with a strike price of $50 that expires in three months. If the stock increases in price to $70, the buyer can exercise the option and purchase the stock at $50. If they turn around and immediately sell the stock, they will profit from the difference in selling the stock for $70, having purchased it for $50 (minus what they spend on buying the option). If the stock stays below $50, the buyer can let the option expire. In that case, they have lost the amount they spent on the option.

This differentiates options from forwards and futures, the latter of which is guaranteed to transact.[62] Complex derivatives trades can be constructed by combining various options. Options can be either *Puts* or *Calls*; one can express confidence in a stock price (take a *long* position) or express pessimism (take a *short* position) with either sort of contract, depending on whether they are a buyer or seller. A Call contract is the right to purchase a security at some price *x (the strike price)*. The purchaser of a Call contract is said to be long in the security since they benefit when the security price exceeds the strike price. If the price of the security does in fact rise, they will be able to purchase the security at the strike price of their call, sell the security in the market at some higher price, and the profits from such a trade will be the difference between the market price and strike price, minus the premium they paid for the option contract itself.

The seller of a Call contract is said to be short in the security because they do not expect

[62] In the case of forwards and futures, they are simply transmuting a purchase *forward* into the *future*, whereas an option buyer has the *option* but not the obligation to transact in the future.

the price of security to exceed the strike price— they would prefer to keep the security if they thought its price was going to exceed the strike price. In selling a call, the trader will profit on the premium she sold the contract for but runs the risk of infinite losses if the stock price actually rises instead of falls (this is not the case in selling puts, because a stock price can only hit $0). They do not expect the call option to be exercised and so they are being paid for an eventuality they do not think will come to pass.

A Put contract is the right to sell a security at some price *x (the strike price)*. The purchaser of a Put contract is said to be short in the security since they benefit when the security price is lower than the strike price. In buying a Put, the trader is purchasing the right to sell a stock at a given strike price. If the stock falls below that price, the buyer of the put will execute the option, forcing the seller to buy the stock at the (now overvalued) strike price. The buyer can then sell the security at the (lower) market price and will make a profit between the strike price and market price, minus the premium paid for the option contract.

The seller of a Put contract is said to be long in the security because they do not expect the price of security to fall below the strike price. If they have sold puts and the stock price rises, they will have profited from the premium they sold the Put contracts at, and the contracts will remain unexecuted (as they have turned out to be unprofitable from the perspective of the trader that has purchased them). They do not expect the put option to be exercised and so they are being paid for an eventuality they do not think will come to pass.

The Call and Put contracts themselves are sold for some premium. Most simply, if a trader is long on a security (they think the price of the security will go up) they can buy Calls and/or sell Puts; if they are short on a security (they think its price will go down) they can buy puts and/or sell calls.

Swaps are a broad category of derivative contracts in which two parties agree to exchange streams of payments over time according to specified terms. The term 'swap' is useful because it names a common contractual structure that can support transactions motivated by disparate aims. For example, hedging risk, gaining exposure to certain price movements, or altering the timing of cash flows can all be accomplished using swaps.[63] A swap consists of two cash flows and usually has the structure of one fixed 'arm' and one floating 'arm.' These arms are streams of cash flows that are exchanged between two counterparties. One arm is usually fixed, meaning the payment amount is predetermined. The other arm is floating, meaning the payment varies according to some underlying benchmark (e.g., an interest rate, or commodity price).

In an interest rate swap, a counterparty who has a variable-rate loan might agree to pay a fixed rate to a bank while receiving a floating rate tied to SOFR. This allows the company to lock in a predictable series of interest payments (which is good for planning purposes) and hedge against the risk of rates rising. Conversely, the counterparty to them could be a trader who has the view that interest rates are likely to go down (so they predict that the floating rate will fall relative to the fixed rate they receive, generating a profit).

Swaps are often used to mitigate the risk of some variable increasing or decreasing, or to profit off the increase or decrease in some value. A commodity-based swap might allow an airline to fix their costs of jet fuel (in case the cost rises). A Chinese manufacturer could use a currency swap to hedge against the U.S. dollar weakening relative to the yuan, since they receive payments in dollars but must cover expenses in yuan. Most swaps are over-the-counter (meaning

[63] These objectives may be desired by individual counterparties for a variety of reasons and with various intentions.

not exchangetraded) with the notable exception of credit default swaps or CDS. CDS are more standardized and sometimes are cleared through exchanges.

These examples illustrate the variety of things derivatives can be used for, and the structural differences between products. To assess the functioning of financial markets, we also need to know the landscape in which these products are traded. I compare products across three dimensions, including how they relate to a market's thickness and liquidity, the differences between private and public markets, and comparing ways to short securities.

Comparing Products: Thickness versus Liquidity

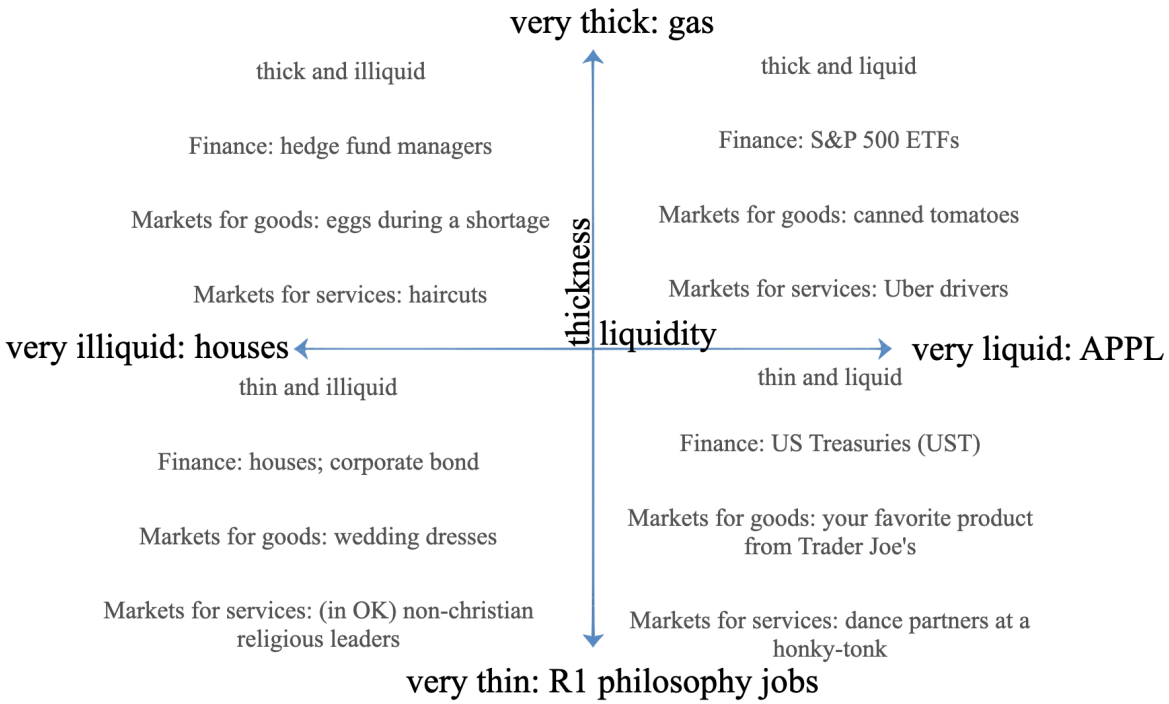

Figure 2: a visual representation of the differences between liquidity and thickness

In securities markets, bespoke derivatives are less liquid (or potentially completely illiquid, as in non-tradable) than stock in Apple, which trades upwards of fifty million shares

each day.[64] Assets can also be 'securitized' (like mortgages) to make them more tradable. Thickness is not often described as being a value in finance, but maybe it should be.[65] Thickness matters because it affects the availability of exchange participants at the moments when people want to trade, and thus opportunities for exchange. The more participants whose buying and selling intentions overlap in time, the easier it is for any individual to transact without moving prices in a volatility-inducing way.

Al Roth defines thickness as "having a lot of participants," but that is not sufficient for a market to be thick. Thickness is, in part, about the number of participants, but there is also the matter of their temporal distribution of intentions. It is not just about the number of participants, but also *when* they are buying and *when* they are selling. Roth uses examples which implicitly describe the temporal element of what makes thick markets, without explicitly recognizing it. A market with many participants who never show up at the same time is, functionally, thin.

Thickness is about the availability of possible transactions in time, while liquidity concerns the ease and cost of transacting with those participants once they are available. There can be thin but liquid markets when there are few participants, but their trading intentions are stable and well-coordinated over time.[66] Conversely, there can be thick but illiquid markets, where many participants exist but their activity is clustered on one side of the market (lots more buyers than sellers or vice versa), making the execution of trades difficult.

The temporal structure of participation helps make sense of these cases. A thin and liquid market is one in which there are few possible transactors (perhaps the threshold to participate in

[64] Apple Inc., "Stock Price," Investor Relations, accessed May 24, 2025, https://investor.apple.com/stock-price/default.aspx.
[65] I will connect this to my teleological account in Chapter IV.
[66] This can happen in niche bond markets. For example, in certain municipal bond markets, there aren't very many participants, but these participants trade regularly and continuously.

the market is high), but for those transactors, trading is easy because their activity is continuous, predictable, and temporally aligned. A thick but illiquid market has plenty of potential transactors and low barriers to participation, but mismatched timing, supply shocks, or habituated preferences make it difficult to buy or sell at the moment one possible transactor wishes to transact.

For example, R1 philosophy jobs and gasoline differ dramatically in their thickness. In the academic job market, there is asymmetry: there are few buyers (R1 institutions) and many sellers (Philosophy PhDs), and the timing of search and hiring is extremely rigid. Gasoline, by contrast, has plenty of buyers and sellers who regularly have both demand for gas and sellers able to produce it, but its liquidity is shaped by cartel-influenced supply dynamics.[67]

Houses and Apple stock are neither particularly thick or thin, but the friction of transacting makes them differ in terms of liquidity: A house may sit on the market for months before a seller and buyer overlap in time, whereas APPL shares can be bought or sold almost instantaneously during trading hours because there is nearly continuous temporal overlap in buy/sell intentions.[68]

Market thickness should be of concern because a scarcity of synchronized active participants increases the risk of volatility. Product illiquidity should be of concern because friction in executing trade also increases the risk of volatility.

[67] On my chart I have gas as being neither particularly liquid nor illiquid for this reason. Perhaps in some macroeconomic environments in certain periods of time, it is a very liquid and very thick market.
[68] The instantaneous nature of this market is largely due to high-frequency trading, but liquidity would otherwise still be very high due to the market-making function of investment banks.

Comparing Products: Primary versus Secondary Markets

When people use the term "the stock market," they are often referring to public financial markets. The defining feature of public markets is that they are not private: most people, if they have money in a bank account, can buy things traded on exchanges with minimal requirements.[69] In contrast, private markets include everything not transacted on an exchange. This includes OTC derivative products mentioned above, along with private equity (i.e., anything done by a private company), venture capital, real estate, and private lending. Private markets are more limited than public markets and limited to 'institutional investors' who meet certain wealth and income requirements. The gatekeeping in private markets was established due to the illiquidity common to these markets, along with a different and often lesser regulatory burden which implicitly places more emphasis on investor research and responsibility.

There are some interesting heuristics people tend to apply to public and private markets: It is often the case that longer-term funds like pension funds will invest a great deal in private markets, often by way of investing in venture capital funds, which fund private investments themselves. But it is not entirely clear that private markets are less risky than public markets. One serious difference between the two is that assets (like stocks) which trade on public markets fluctuate in price to a greater degree because their prices change with new information with little lag time. Private markets, in contrast, might go up and down without anyone paying much attention. This has to do with conventions on valuation. If you own 100% of a business, and I buy a 1% stake in your company for $1 billion, you now own 99% of a $100 billion dollar company. This might be intuitively puzzling, since I only wanted 1% of the company. Who is to

[69] These requirements can include that their money and/or bank are denominated/domiciled in the same country as the exchange. Depending on the nature of the asset in question, there may be certain rules around posting margin in order to transact in derivatives, which is not done with the intent to include market participants but rather to avoid default and credit risks.

say that, if you had wanted to, you could have sold the other 99% (or even just 10%) at the same rate? This presents a major issue for both private companies and for venture capital and private equity funds. Private companies do not want to sell more at a lower valuation to investors, even if they really need more capital to keep the business running. Private equity funds do not want to sell ownership stakes at a loss (because it is at that moment that they must 'realize' the losses on an accounting basis), so they will sometimes hold onto ownership for these extrinsic reasons. Private markets are less liquid than public markets, but we cannot deduce from either 1) evidence that valuations do not often go down or 2) the lack of volatility in the valuation of these companies that they are less 'risky' investments.

Comparing Products: Synthetic versus Physical Shorting

Short selling gets a bad reputation as being the epitome of speculation motivated by greed. There is an important function to short selling, without which markets would have prices less responsive to new information and thus at greater risk of volatility and inaccurate pricing.[70] When a market participant (let's say, a hedge fund) has a trade thesis which includes shorting a security (here I'll just be talking about stocks and bonds; shorting currencies, commodities, and other assets behave differently), there are two kinds of shorts they can put on: "synthetic" or "physical."

A synthetic short encompasses anything done using options, futures, and most swaps. In this case, the hedge fund does not directly borrow or sell the underlying security. The exposure is created through financial contracts that have the same payoff of a short position: the assets

[70] Perhaps an explanation of why inaccurate pricing is bad for society can be discussed later.

associated with the synthetic yield a higher profit if the security's price falls and decreases in value if the security's price rises.[71]

In a physical short, the hedge fund borrows a real share of the security, sells it at t_0 (for what they think is a high price), and buys it back at a later date (they are hoping it is a lower price) at t_{1+x}. For each day the trader is short, they pay rent to the people they borrowed it from (the investment bank).

Physical short selling relies on a supply of securities to be loaned. The investment bank might get these shares from their own balance sheet, but usually they borrow the shares from others. The supply side of physical short selling is largely composed of securities held in long-term accounts, like index funds and ETF providers, who are happy to lend out stocks for a fee ("rehypothecation") since they have to sit on them anyway. Investment banks here act as intermediaries by borrowing shares from long-term holders and lending them to hedge funds (and other firms) who want to short. The quality of this supply depends on both the quantity (shares available to be lent) and stability of the supply (how likely the lender is to need the share back)

If a supplier needs their stock back (which might happen for a variety of reasons including index rebalancing or voting), the investment bank tries really hard to find a replacement of the supply so they don't force the hedge fund's hand in closing out their position earlier than they might want to. So, from the investment bank's and the hedge fund's perspectives, a good supply dynamic is one in which there is a lot of stock to borrow and there is

[71] Here is a simplified example: The hedge fund wants to short XYZ, which is currently trading at $100. They buy a put option for $5 with a strike price of $100 which expires in three months, and they sell a call option for $3 with a strike price of $100 that expires in three months. The hedge fund has paid a net $2 for the privilege of putting on this trade. If the stock is $80 in three months' time, the hedge fund purchases a share of XYZ for $80 and uses their put option to sell XYZ for $100. Their counterparty on the call option is likely to not use it (because, if they wanted to buy XYZ, they could buy it for $80 instead of $100). The hedge fund has made a net $17 profit (their profit from the put opinion, minus the net costs of putting on the trade).

not a high likelihood that those renting out their stocks will need them back anytime soon. Before the rise of the index fund, both of those were bigger issues. This may be why the desks that manage this activity within investment banks have shrunk dramatically in the last two decades.

Securities lending (or "stock loan trading") is an especially opaque business because there is not anything like a public stock market. So, the rents being charged by other banks for hedge funds to borrow are mostly unknown to any individual bank. Hedge funds have some (or maybe most) of this information, but they usually do not have an incentive to share it. Securities lending has become increasingly automated, which is driven by algorithms that can reliably predict the supply dynamics of how many shares of whatever are going to be available for rent (and this would not be possible if the supply dynamics were more chaotic and unreliable). Given these two options, *why* would someone short 'physically' rather than synthetically? Most of the time it is because it is cheaper than the ‘rent’ they would pay using a synthetic swap.

Important Institutions

The institutions that are supposed to implement the financial market processes above include banks, brokers, investment firms, credit rating agencies, clearing houses, and regulatory bodies. The lines between some of these roles can blur, and there are some institutions that have both a ‘bank’ and a ‘broker’ within them (e.g., large investment banks). First, we can define the specific roles and then see how they might be combined within institutions (for both good and bad results). For the sake of this expository section, I will assume clean lines between these institutions. We focus here on aspects of each role that matter to the overall functioning of financial markets. The process of risk assessment and capital allocation has systemic effects that will require consideration in considering relevant moral features of these professional roles.

Traditionally, banks collected customer deposits, which they then used to fund other activities including loans and mortgages. They expand purchasing power: by issuing loans, they expand the money supply, effectively creating new purchasing power from nothing (as stated in a section above). Aaron James and Robert Hockett argue that this power makes banks and money a *res publica*. When banks extend credit, they are allocating capital and exerting control on which projects, firms, and individuals get access to newly created money, and which do not. Ideally, this delegated authority channels credit toward productive uses that support broad economic flourishing. It can, of course, also serve to concentrate risk, misallocate resources, or serve narrow private interests. James and Hockett use the facts of their socially delegated power to justify the regulation of banks, and they also advocate for public banking. For our purposes, we can think of banks as general contractors with a number of tools. Having been given those tools, banks have the power to shape the construction of future markets. They can be good contractors, who generate honest invoices, hire qualified carpenters, and plan carefully, or they can be bad contractors in the sense of making mistakes while building, or building with poor design, etc.

Banks are in the business of taking short-term, liquid liabilities in the form of customer deposits and transforming them into long-term illiquid assets (e.g., loans). This structure works well so long as everyone remains calm and unworried. But since depositors can withdraw quickly, often more quickly than assets can be sold without depressing prices, a shift in beliefs held by depositors can generate a downward spiral in solvency and resulting existential crisis for a bank.

Spreading rumors about a bank sufficient to cause a large enough (the definition of this depends on many things, including the bank's capital requirements) set of customers to withdraw

funds could cause the bank to collapse, or at least have significant problems to the extent that they might face investigation, be taken over by the FDIC, require an emergency infusion of equity, or be forced into the rapid sell-off of assets at fire-sale prices (which can subsequently trigger one of those other consequences). Unlike markets for products and services, banks (and other financial firms) can be directly and negatively affected by public perception, regardless of the merit of the concern. There is no bank-run equivalent for a publicly traded corporation– mania about their products tends to manifest in action which is good for a corporation.[72]

Since banks are firms which can be impacted by the mere *perception* of them by the public, reformation of how the public views banking and finance is a cogent potential project for stabilization. Theoretically, there could be cases where a fundamentally sound bank suffers a PR snafu or even an insidious smear campaign that triggers withdrawals. There are also cases where banks that are engaged in genuinely unethical or fraudulent activities see no stock-price consequences until the behavior is revealed somehow. In general, the public's perception of Wall Street in 2014 was far worse than in 2006, while the activities of financiers were almost certainly worse in 2006 during the lead up to the financial crisis.

At their core, investment banks are in the business of originating, structuring, and distributing financial claims. They have had varied roles in the last century. In the wake of the 2008 financial crisis, large investment banks are (mostly) no longer allowed to manage their own money for the purpose of making a profit. This has not stopped them from making money, but their orientation has been forced to become focused on fulfilling customer needs without trading in a 'proprietary' manner. This is a seismic shift from the kinds of dynamics reported in pop

[72] One might imagine pathological cases where customers become so enraged that they attack employees or destroy facilities, but these extreme routes are not part of the ordinary causal pathways through which product-market firms are harmed.

culture in books like *Liar's Poker*.[73] They now operate more like a consumer-facing bank, with day-to-day operations oriented more towards facilitating the needs of their clients.[74] The major clients an investment bank has are hedge funds, pension funds, mutual funds, private and public corporations, and other asset managers including private equity funds and insurance companies. Their revenue comes from fees for facilitating different kinds of transactions. The transactions occur with counterparty client companies, and include raising capital, underwriting securities, advising on mergers, and providing liquidity and access to markets.

Still, most investment banks have small pockets within areas of operation that are tasked with 'risk management.' This can serve to give them the justification to invest and trade in a quasi-proprietary manner (they are trying to make money from *this* trade to offset what they might lose from facilitating *that* trade, which they did because a client wanted to do it). Most investment banks have some capacity for proprietary trading and investing within their wealth management arm as well. When they serve as the counterparty to a trade, they earn a fee (their customers 'get a haircut') on the provision of this service. Here they clearly have something to be gained by volume: more deals, more trades, and more business activity like mergers means more business for them. To connect this to earlier-mentioned financial products: Investment banks underwrite initial public offerings and secondary offerings, which involved the issuance of stock. They do the same thing for bonds, structuring and pricing new issuances, locating buyers,

[73] Prior to the Dodd-Frank Act (in the U.S.) and Basel (in the UK and EU) regulatory changes post-2008 crisis, banks trading desks operated much the way hedge funds operate now: they had their own ideas about what prices would go up and down, and they used the bank's own money to instantiate those trade theses. The firm's balance sheet being used to make bets was obviously not always in the best interest of clients, and there were serious conflicts of interest between traders who may want to position clients as losing counterparties.

[74] Matt Levine writes: "investment banks used to be well capitalized, lightly regulated, adventurous, and willing and able to commit their own capital to make markets more liquid and efficient. Now they are more regulated and risk-averse, and the hedge funds have stepped into some of those roles." There's also a case to be made that private equity companies have taken over some other traditional lending activities that used to be largely done by banks. See "Pod Shops are the New Banks." Bloomberg, 2025.

and facilitating trading of those bonds in a secondary market. Investment banks are often the counterparties to derivative trades.

The stylized mechanism for the way this works is that for a given client with a lot of money, their wealth manager might present that client with various options that include that wealth managers subjective view on what equate to trade theses ("I think the market has underpriced the risk of a decrease in inflation in the next year"). In all of these cases, the investment bank's role is to serve a client's needs, whether that need be to raise capital, allocate risk, or do business with another company. Post-2008 financial crisis, investment banks are heavily constrained in doing things 'proprietarily' or for their own sake: they are supposed to, ideally, merely serve the needs of their clients. Their interest in expansion of financial products is mediated by regulatory requirements, and their self-interest in keeping their own balance sheet hedged to prevent negative outcomes for themselves (which would generate reputational risk and likely negatively impact their portfolio of clients).

Investment banks are oriented towards facilitating what their client wants to do, but in high-net-worth wealth management cases, what the client wants can be shaped by what their point person at the investment banks suggests they do. Banks are not supposed to trade proprietarily, but they are allowed to trade to accommodate "reasonably expected near-term demand."[75] There is some interpretive flexibility in this term, as it does not specify what precisely constitutes a client's demand. This ambiguity gives traders discretion to make predictions about what clients will want exposure to, given the general market environment.[76]

[75] See John Ramsay, "Remarks on the Volcker Rule's Market Making Exemption," speech at the New York City Bar Association, New York, NY, February 4, 2014. SEC. https://www.sec.gov/newsroom/speeches-statements/2014-spch020414jr

[76] Not to mention that at a major investment bank, most parties that would be doing the buying in case there is 'market at large' demand are themselves clients of the firm

Agents that work in investment banks that are discussed below include traders and investment bankers.

Hedge Funds receive a great deal of attention from pop culture and the media. There are several different types of hedge funds, but they all share a common structure: they deploy capital on proprietary trading ideas. The employees of the hedge funds (portfolio managers and analysts) come up with ideas (a 'trade thesis') about where prices will rise or fall, and they use the money clients invest in the hedge fund to instantiate those trading ideas. Their 'clients' are their investors, who typically have no involvement in the fund's investment decisions or day-to-day operations. Since the hedge fund is more expensive than putting money into an ETF or mutual fund, investors are not merely looking for net gains but gains higher than what they would make putting their money into a broad-based index. There are usually a variety of agreements investors must make with hedge funds to not take their money out for some period of time which gives hedge funds the ability to worry less about liquidity and panic-selling out of positions.[77] The basic interest of any hedge fund is to make money and outperform other (less expensive) mechanisms for making money. Since they charge a fee for the service of investing money, they need to make more money than their clients could have made by investing in a general index fund.

There are two central differences between hedge funds and pension funds: the time horizon and their client base. Hedge funds are seeking short-run (usually yearly) gains: in any given year, they want to maximize the dollars they will have the following year. Managers of hedge funds are often compensated with a bonus structure that is dictated by yearly gains. On the other hand, pension funds are seeking less risk because their focus is further into the future. The

[77] These 'lock-up periods' contrast with other popular investment funds and products like ETFs, where clients are often able to withdraw/liquidate at any time.

clients of a pension fund are usually employees of a specific employer or union, who are paid out when they become retirees. They would like to turn dollars this year into as many dollars as possible next year, with the proviso that they really do not want to lose any dollars next year or in 50 years. Pensions funds engage in the same kind of activities that other hedge funds do, but with a different set of preferences. Pension funds use the same set of financial products as tools but are much more risk averse. They are slower-moving and so have a stabilizing force on asset prices.[78] It should be unsurprising if hedge funds are more short-term minded than pension funds due to their different natures. However, the compensation structures for employees at both kinds of firms are quite similar. Employee reviews follow the review process common to many employers, and a yearly review will usually, in part, be about one's financial performance over the last year.[79]

Pension funds are perhaps the least controversial institution in finance: they exist to pay out retirees in the future, like teachers and first responders. A philosopher may naturally wonder, why not just have pension funds, since they seem immune to certain kinds of incentives associated with questionable behaviors? The short answer is that not all economic activity is aimed at meeting predictable, long-dated cash obligations. Pension funds are an engine for the distribution of future capital, but they are not interested in taking short term liquidity risk[80] (which is necessary to make markets), underwriting securities[81] (necessary for issuing new

[78] Pension funds are also subject to different regulations.
[79] In other words, I am suggesting that the claims made about how these operate differently are sometimes overstated. It can be the case that pension funds are aggressively investing and trading in the same manner as hedge funds, but with a small percentage of their overall funds.
[80] Short-term liquidity risk means there's a risk of an asset not being able to be sold quickly without a potentially substantial loss in value. Because the obligations of a pension fund are relatively stable (e.g. they know how much they will need to pay out in the future to pensioners), they cannot afford to rely on quickly selling assets at unpredictable prices or waiting for opportune times.
[81] Similarly, pension funds would like to avoid this due to the short-term risks involved. That they are liability-driven (e.g., they know how much they will need to pay out in the future to pensioners), they are not motivated by

stocks and bonds), owning stakes in uncertain enterprises (which is necessary for non-already-mature companies to exist), and they have a uniform set of risk preferences (which means they cannot be the counterparty for derivatives necessary for a system of risk management).

Asset Managers are large entities which usually invest in more conservative ways than hedge funds and more systematically compared to pension funds. Major asset managers include Blackrock, Vanguard, and Fidelity and their major products include ETFs and mutual funds which tend to be more 'passive' investing products. Here, passive investing contrasts with actively investing, in that most customers who invest in these products are saving for retirement and their investing is more long term. Asset managers often take no overarching proprietary positions. They may offer products which themselves have some trade thesis or sector specificity, but the asset manager is in the business of offering the investment product which itself tracks something. There are a variety of concerns regarding large asset managers related to the idea that universality of ownership via widescale participation in ETFs might have negative effects. These worries will be discussed in Chapter V.

Private Equity firms deploy capital from their clients into non-public markets. They were supposed to exist to take dysfunctional companies and give them the capital and expertise to improve and potentially restructure. That involves absorbing risks, fixing inefficiencies, and otherwise 'creating value.' They sold themselves as experts at improving companies. In recent years, an increasing number of traditional investment bank functions have also been shifted over to private equity firms.[82] This includes things like raising capital, advising companies, and underwriting.

---

the possible upside as much as they are worried about the potential downside of not being able to meet their predicted obligations.

[82] Matt Levine, "The Banks Are Re-Tranching," *Bloomberg Opinion*, June 9, 2025, https://www.bloomberg.com/opinion/newsletters/2025-06-09/the-banks-are-re-tranching

There are also common stereotypes about private equity companies, that they are leeches on the real economy who break up firms and 'sell them for parts'. Unfortunately, this stereotype has proven true in many cases. Private equity firms have short-term incentive structures, which means that in many cases it makes sense for them to borrow heavily against the company's assets to fund payouts or acquisitions (to 'lever up') companies, extract cash and liquid assets, and then sell the firm without putting in any of the work necessary to improve companies long term. There have been many cases where private equity entrants lead to harmful cost-cutting that leave companies worse off.

For the sake of this project, the biggest concern with private equity is that they are getting their cake and eating it too: Banks are stringently regulated because of the things they do and the negative impact that doing those things poorly can have on society. Private equity firms, in many cases, have begun performing similar functions (e.g., allocating large amounts of capital, leveraging companies, and managing complex debt instruments) without the oversight, transparency, and regulatory safeguards demanded of banks. This is a very reasonable source of concern.

Exchanges and Clearinghouses are responsible for facilitating and executing trades. Exchanges serve as the traditional 'arena' where trading occurs. The role of exchanges has decreased in recent years due to the rise of dark pools, which are essentially private exchanges. They have received criticisms for their opacity but can be valuable for investors who want to buy or sell large quantities of an asset without moving the price. Not causing price changes due to a large buy or sell order is a legitimate aim, and not only wanted for self-interested profit maximization reasons, but because for large trades both the investor and the (usually) investment bank trading desk assisting in the execution of the order can potentially be subject to charges of

market manipulation.[83] They are, however, freeriding in obvious ways on the price discovery provided by public exchange markets.

Clearinghouses are institutions which manage the nitty gritty details of getting securities from a seller to their buyer. Today, few paper share-certificates demonstrating ownership in a stock exist. Most stock ownership and changes of ownership are recorded electronically. Exchanges and clearinghouses are the basic infrastructure of financial markets; they are the pipes which allow the trades to flow and for assets to move between buyers and sellers.

Credit Rating Agencies have an important ontological role. Broadly, they are in charge of assessing riskiness. If a credit agency gives you a 'good' rating, this means they think you will be able to repay money loaned to you. If a credit agency gives you a bad rating, it means there are some doubts about your future solvency and ability to repay debts. The obvious way this can go wrong for the sake of risk distribution is if the credit rating agency rates the asset or company in question incorrectly. This was of course related to the mortgage-backed securities which, when compiled together into tranched products, were deemed of a higher credit rating than appropriate given their lowest-rated parts.

Central Banks and Regulators have been touched on lightly, but will be covered in more depth in Chapter V. For now, we can think of central banks and regulators as putting some conditions or restrictions on the market. They do this with the intent and directive of making society better off in various ways. Many people disagree about how best to do that.

[83] They can also protect market participants like pension funds trading on behalf of long-term beneficiaries from direct downward price movement and prevent front-running of high-frequency trading algorithms which might be able to detect that near term demand will be high due to signals of a large buy or sell order. In this sense the self-interest of the market participant is also in the self-interest of the pension fund's beneficiaries.

Important Agents

One thing that makes financial markets different from markets for products and services is that their participants are different. They differ not only in how much of society participates in each market, but also in the roles they play in society's economic life. You are more likely to need to buy apples (or at least food) at the grocery store, but no one needs to buy APPL, the publicly traded security. While almost everyone participates in product markets directly, far fewer participate actively in financial markets. Still, almost everyone is *affected* by financial markets– either directly in the form of things like retirement savings, banking, or indirectly in the form of being the customers of the future markets which finance creates.

Within the realm of markets for products and services, there is extensive business ethics literature on the firm.[84] Shareholder versus stakeholder views, theories of profit, and analyses of market failures are all helpful in answering normative questions about companies that produce goods or provide services. Financial markets also have firms, but this level of delineation is insufficient to capture the relevant details necessary for robust normative analysis. Financial markets involve actors and institutions that do very different things and play very different roles. These actors are without analogs in existing literature on ordinary markets and so they are not well-theorized.

These are the three main types of traditional financial agent. There are many permutations of these agent-types, and additional agents (like retail traders), all of which will prove to be important for (and potentially disruptive to) the operation of financial markets.

Investment Managers allocate capital from some group of investors (which may or may not include the manager) to some set of investments, with a specific desired output along a

[84] See the work of Joseph Heath, Thomas Donaldson, and Edward Freeman.

certain timescale. Robert Shiller describes them as sharing a common purpose to "determine the composition of portfolio of investments on behalf of their clients and buy and hold those portfolios for those clients, the ultimate investors." Their risk appetite is shaped by the nature of their fund. What rule-based approaches will not tell us about this kind of agent is what their maximum risk level should be. It also will not tell us what kinds of innovative technologies they should invest in.[85]

Traders too encompass a wide variety of specific responsibilities. Two contrasting traders to keep in mind are those oriented towards market-making (e.g., supplying continuous bid–ask quotes to provide liquidity) and those oriented towards proprietary trading (e.g., taking short-term positions based on statistical arbitrage signals). To be clear, methods for market-making are often proprietary, but the purpose of a market maker is still wholly different. Shiller describes the reward mechanism that traders manage as a kind of parallel evolution to the reward system found in the human brain, whereby "because of traders' activities, the reward is virtually instantaneous, and it serves to encourage similar activities again and again."[86] To understand market makers, imagine a large store that not only sells virtually everything but is also willing to buy almost anything from you.[87] We are used to this, at least from the consumer perspective: Walmart is incentivized to make prices as low as possible to attract consumers. They also want to have enough of the trendy water bottles *de jour*, Tickle Me Elmos, etc., to satisfy consumer demand (lest enterprising buyers snatch up all the limited quantities of whatever trendy goods can demand a higher price on the resale market). Market makers are incentivized to do this too, regardless of whether they think the thing their customer wants to buy or sell will go up or down

[85] A utilitarian theory here is virtually impossible to implement due to the unforeseen nature of new technologies.
[86] Shiller, *Finance and the Good Society*, 58.
[87] Let's pretend for the moment that marketing, self-identification with consumer brands, and convenience of locations don't exist for this example.

in value. In this regard, they are on the frontlines for providing liquidity in public markets: if, in a time of uncertainty or economic downtown, no one wants to buy what anyone is selling, markets would be worse off.[88] Proprietary traders are said to also provide liquidity (not by posting continuous two-sided quotes like market makers, but because their attempts to execute trading ideas still adds depth to the market and creates additional opportunities for others to transact, or transact as lower prices), but they are more important in the price discovery role: using whatever means they have at their disposal, they come up with trading theses ("this is overpriced" or "that is underpriced") which they enact in the market. By many small, invisible hands, prices move toward being more accurate.

Investment Bankers help companies buy and sell parts of themselves and of other companies. This function tends to take one of the following mechanisms: an initial public offering or IPO, where a private company owned by a few becomes a public company owned by many; a merger or acquisition, where two companies (or parts of two companies) become one; or financing, using loans, additional share issuance, or debt instruments to help companies raise capital. In order for bankers to do these things, they must be well versed in evaluating companies. This includes identifying the riskiness of a company, which contributes to the overall assessment of how much their capital will 'cost,' or what price a portion ownership ought to fetch.

For all of these agents, the general directive is to maximize returns (capital, money, profit) with legality as a constraint. A more careful account of the role of financial markets will not, by itself, resolve the full range of social problems. It will not determine which forms of

[88] This is one objection raised by the emergence of algorithmic trading: in 'unprecedented' times, models won't want to take on risk by buying up assets, which might supercharge downturns, as evidenced by Black Friday (and suspected of the COVID-19 and Liberation Day selloffs). See: Odd Lots 'Vineer Bhansali on Losing Fed Independence as the Biggest Tail Risk Right Now.' September 13, 20205.

innovation ought to be pursued within particular sectors given certain risks and costs. However, as will be argued in Chapter V, incorporating a component of collective deliberation can help guide financial investment in a manner that is democratically determined, while remaining subject to the familiar advantages and limitations of democratic decision-making.

Summary and Preview of Coming Attractions

This chapter will be used to support the idea that financial markets should be treated as conceptually distinct from markets for products and services. This applies to the normative theorizing done by philosophers, and ought to have downstream effects for how we see the role of regulation and both individual and institutional behavior. As I will describe in the next chapter, when financial markets receive attention from philosophers, economists, sociologists, and historians, they mostly think about financial markets using similar mental scaffolding used for thinking about markets for products and services. Part of this dynamic can be explained by the pervasive myth of the barter system which supposed that finance and financialization is layered on top of more 'fundamental' market activities. It is also true that finance used to be much simpler, and we could get further by merely examining lending markets.

Given that the complexity of finance has steadily increased over the past century, analytically distinguishing it from markets for goods and services provides a useful framework for interpreting more recent phenomena than the classical political economists could predict would emerge such as meme stocks, cryptocurrencies, and the expansion of index funds. From a normative perspective, a generalized conception of markets, when applied to finance, will be insufficient so long as it treats finance like a "black box."

Getting this right matters because even if we have markets for products and services that are operating perfectly, and moral exemplars working within the profession of finance, systemic

failure remains possible if risk is miscalculated and innovation might not occur if investment decisions are made without a clear understanding of how markets can change over time. Most market interaction in the products and services realm is synchronic, and so too is analysis of those markets, focusing largely on distributional issues and other questions related to the proverbial economic 'pie' as-is. Finance, on the other hand, is largely diachronic, and is oriented towards economic pie-changing: how big it will be, what products and services will emerge and what will become less popular, and so on.

In Chapter IV, I outline what I see to be the distinct telos of financial markets, from which we can better understand i) precisely what ways finance avoids taking into account social good, and ii) how we might improve it for the better.

Before that, Chapter III reviews relevant literature in moral philosophy and political economy to demonstrate that philosophical frameworks developed to analyze markets for products and services will be incomplete when applied to financial markets. Whereas theories of product and service markets typically center welfare or preference satisfaction as justifying their existence, we should see finance as being justified for its potential role in distributing risk for the sake of social goods like stability. I cover historically dominant approaches shaped by welfare economics, Hayekian and Keynesian views and their descendants, and behavioral psychology. I argue that each overlooks this role. I diagnose this omission as a possible effect of the "myth of barter" which lends itself to the historically inaccurate view that exchange economies preceded financial abstraction and organized risk distribution. I also examine a set of common critiques of financial markets and show how they motivate the need for a more robust theoretical framework.

Chapter III: Literature Review

In this chapter I identify some major threads in the literature that connect to my project. Following this literature review is a taxonomy of seven common criticisms of *both* markets for products and services and financial markets. In addition to providing a traditional literature review, I want to establish that the literature on markets and finance largely fails to address the components of futurity and risk that my project draws out. The philosophical and economic literature on markets focuses primarily on the conditions and results of exchange. This, I argue, leaves the systemic role of risk in society undertheorized.

One explanatory reason for this undertheorization of finance, I suggest, is in how economic life is commonly imagined to have developed. Financial abstraction is often treated as a relatively recent technological development layered on top of (and thus separable from) an earlier, more basic system of exchange for products and services. On this account, early economic life consisted primarily of direct barter, with money and then finance as we know it emerging only later as tools to facilitate and streamline that trade. But both historical and anthropological evidence suggest this narrative is inaccurate. Human societies have long had either informal or institutional mechanisms for the financial functions of extending credit, managing uncertainty, and distributing risk across time. What later becomes formalized and thought of as *finance* is therefore not a separable, add-on layer to, or derivative of, economic life. Rather, it is an institutionalization and elaboration of functions already present in earlier systems of exchange. This matters for the account I develop in Chapter IV, because if economic life has always involved coordinating action under conditions of uncertainty, then risk management is not a peripheral feature of markets but rather a mechanism for handling a central elemental component of economic life.

I sort the literature related to my project into three intellectual lineages or threads. First is the thread of teleological moral philosophy that begins with Aristotle. Second is a classical economics thread which ties in Friedrick Hayek's information theory, the Chicago School, welfare economics, and rational choice theory. The third thread connects Keynes's psychological account of markets with work in behavioral economics, critiques of concepts of 'economic rationality,' and work in market design. These three threads in the literature are valuable resources, and I will identify common criticisms of financial markets that emerge from these three strands. What I hope to demonstrate is missing, however, are concerns related to futurity and risk that this literature hardly touches. While the term "risk" does appear in this economic and philosophical literature, it does not really get to the kind of society-wide risk I am thinking about. If mismanaged via finance, risky situations can be a social disaster; if finance can distribute it appropriately, society would benefit tremendously.

The goal of this chapter is to reorient philosophical discussion of finance around its capacity to distribute risk across individuals and institutions, thereby enabling economic coordination and growth. After discussing the emerging threads in literature from connected disciplines, I survey seven prominent criticisms of financial markets that emerge both in academic literature and in public discourse. These criticisms are important and deserve serious consideration. However, many of them apply equally to markets for products and services and therefore do not capture what is distinctive about finance. By contrasting how these criticisms apply across these different market domains, we can see what is distinctive about finance and why theoretical frameworks developed for ordinary markets may be insufficient for understanding it and providing normative guidance. This paves the way for my own account in Chapter IV, and consideration of alternative frameworks that *are* attuned to the unique features

of finance. Understood this way, the subjects of financial criticism (e.g. financial crises) become a symptom of our lack of collective deliberation rather than a testament to the intrinsic evils of finance or, alternately, a pill we must swallow to keep the benefits of a market system.

The Origins of Finance: Myth v Reality

There is a familiar historical narrative about the origins and evolution of economic exchange. The story goes that once upon a time, early man first engaged in barter, exchanging goods directly with one another when opportunities arise– the traditional "double coincidence of wants".[89] Unfortunately, the process of trying to find someone with ten chickens who wants to trade them for your cow (say) became inconvenient, cumbersome, and inefficient. Then, the story goes, mankind invented currency, which allowed exchange to be facilitated more conveniently. Financial instruments and abstract financial products are then imagined as much later technological developments. This narrative implies a linear progression from simple exchange to financialization, but anthropological work has largely discredited this account.[90] The historical record indicates that, in fact, systems of credit and debt appear earlier than markets organized around spot barter and alongside or in some cases prior to the development of coinage. It does not appear to be the case that many (if any) small communities relied on a double coincidence of wants to facilitate exchange.[91]

Instead, goods moved around and circulated in communities via kinship and relationality and structured by trust, obligation, and future expectations of reciprocation. Consider the following situation: If one villager (a rancher) gives another villager (a farmer) a cow in January,

[89]Jevons, *Money and the Mechanism of Exchange,* 8.
[90] See anthropological and sociological work conducted by Marcel Mauss, Alfred Mitchell-Innes, Karl Polyani, and David Graeber.
[91] Graeber, *Debt: The First 5,000 Years,* 36.

it would be odd to expect the farmer to be in a position to 'transact' with reciprocal goods of similar value given the seasonality of harvest time for most crops. What seems reasonable is that the cow given in January generates a continuing obligation. The recipient (the farmer) might repay the debt gradually in 'installments' over time (let us say, in sets of peas come harvest time). There might be no formal repayment schedule, no interest rate, and no written contract. But this does not mean that repayment expectations did not exist, so long as there exist embedded social norms and community enforcement.

If we suppose that the farmer who received the cow soon appears with a brand-new hoe for the upcoming pea harvest or some other conspicuous purchase while still owing goods in return, we might expect some degree of community enforcement against that violation of understood norms. Perhaps this behavior will be scrutinized, and a third community member (let us say a baker, acting on behalf of the rancher) asks the farmer for a reasonable justification for their use of resources to acquire new products instead of repaying outstanding obligations. In the event of delayed repayment, the farmer 'debtor' has some set of justificatory obligations owed to the 'creditor' and perhaps the rest of the community to avoid negative impact to their reputation and thus capability for future cooperation. Even in the absence of a formal contract or currency, the exchange is mediated by expectations about creditworthiness, repayment and risk. If we suppose the farmer never repays the rancher in peas, the rancher seems less likely to engage in exchange with the farmer again in the future.

So, what might, on the surface, appear as 'barter' seems to have many embedded social relations that plausibly resemble analogs in modern finance. In this mundane scenario (albeit hypothetical), credit was extended to the farmer, risk was assessed by the rancher, and the farmer's reputation (and risk of not qualifying for future exchange) functioned as collateral.

Social norms provided mechanisms of enforcement, and community members served informally as ratings agencies. Seen through this lens, even 'simple' exchange involved ongoing financial relationships between community members.

This is aside from the historical record which demonstrates as early as the Code of Hammurabi that proto-commodities futures contracts were being used to mitigate risk from farmers. This fact refutes the timeline implied by the myth of barter that abstract financialization only occurred after the development of markets for products and services. It seems that from the beginning, society used finance to meet their needs for managing (e.g. distributing away from isolated holders of) risk.

The influential myth of barter influences much of the economic work on modern markets. This is a descriptive hypothesis of how we have managed to avoid talking much about futurity and risk. Discussions of efficiency in textbooks, for example, often work with exchange models without any production sector at all. The historical focus on the direct exchange of goods and services, rather than on the role of credit, debt, trust, and financial abstraction, has contributed to the modern neglect of theorizing risk, the management of which is essential to the stability of society. This is acute in modern, globalized, society, but was needed in ancient societies too: psychological justification to do things at present given the reality of future uncertainty. What was previously unofficial and unenforced other than by the laws of physics and how difficult it was to move around is now officially sanctioned and enforced by financial and banking institutions: You were unlikely to 'default' on your loan of a cow's worth of goods, because there was no interest on the loan, no official repayment schedule to make or miss payments on, and a low likelihood that you would leave your home area. Now there are official payment networks that convince the steakhouse where you eat or grocery store you shop at that if they

perform a specified ritual with a rectangular piece of plastic, they will be paid with due haste and will serve as backstops in case you aren't good for the money.[92] Still, my account does not rest on the accuracy of this hypothesized explanation for why we've neglected to philosophize about risk. It is sufficient for our purposes to reject the notion that financial abstraction is a recent human development that can be separable from economic development and to make clear that humanity has always needed some mechanism (institutional or individual) for distributing risk.

Thread One: Moral Philosophy

Philosophers have long contemplated the moral dimension of markets: which transactions are just? should the state support or interfere with market operations? in what ways do markets shape human psychology?[93] This discussion traces back at least to Aristotle, who in Book I of *Politics* distinguishes between *oikonomikos* and *chrematistiks*. *Oikonomikos* has to do with managing one's household (and the *polis*, by extension). This includes securing the materials goods necessary for sustaining it.[94] *Chrematistiks* is what Aristotle describes when one accumulates wealth for its own sake. Some forms of wealth acquisition are natural extensions of household management ('natural *chrematistiks*,' which is part of *oikonomikos*), but he argues that other forms can become detached from the needs wealth is supposed to serve ('unnatural *chrematistiks*,' which is not part of *oikonomikos*. For Aristotle, this distinction is instructive in understanding when economic activity is getting off track. This underlying intuition that economic activity gone 'too far' can be damaging, and the broader questions of whether the

[92] Obviously, a cash-based transaction would be different, and involves less risk so is less interesting of an example.
[93] Smith 1759, Rawls 1971, Polyani 1944
[94] Meikle, Scott. *Aristotle's economic thought*. Oxford University Press, 1995.

pursuit of wealth serves or undermines the broader goals of flourishing remain central to the philosophy literature on markets.[95]

In discussing the transition from pre-market economies to those that existed in the wake of the industrial revolution, Karl Polyani frames Aristotle's focus on *oikonomikos* (in his work, translated at "œconomia") as a key point of development from hunter-gathering societies: working towards satisfying the needs of a household, as he construes it, only occurs in civilizations with advanced levels of agriculture, and this motivation (as he sees it) need not have anything to do with the concepts of 'gain' or a sense that such pursuits are mediated through market institutions.[96] This sets up a contrast in Polyani's work, from the pre-market society with embedded social relations, to the modern market society with disembedded market mechanisms that serve to organize social life, including the 'fictitious' commodification of land, labor, and money. Polyani uses the term 'fictitious' to describe things that are treated as commodities even though they were not originally produced for sale in a market. Markets subsume society, forcing society to conform to markets running according to market logic.[97] In this view, there is a necessary tension between society and markets, and society has the responsibility of keeping markets in check, lest they are allowed to undermine and erode society.

There is a clear influence of Polyani in the work of contemporary theorists like Perry Mehrling, who applies a related view to finance, in particular. He argues that finance is not a distinct layer of human activity separate from society. Society failed to keep markets– financial

[95] This is discussed further in Chapter IV, but there is a difficulty in figuring out where precisely to delineate between 'natural' and 'unnatural' *chrematistiks*.
[96] Polyani p.55
[97] Polyani p.60

markets in particular– in check, so markets succeed in imposing their logic. Rather, financial relations *are* the constitutive relationships of the system, which he calls Financial Society.[98]

Both Polyani's and Mehrling's critiques are apt when discussing the economic mechanisms that govern the distribution of valuable things in society. Indeed, these distribution features are what almost everyone in the literature is focused on. There is a resulting split between the moral philosophers who critique how the spoils of the system are distributed, and those that focus on how those spoils are generated. This gives rise to an unfortunate tendency for people to talk past one another. Generation and distribution are both crucial. Unless we limit the production of humans (population growth), the economy will need to provide opportunities for economic growth in order to provide for future generations.[99]

One type of neo-Aristotelian virtue ethical view is focused on individuals, identifying general virtues that are particularly appropriate to economic circumstances. This is the approach taken by John Boatright, who discusses ethical issues involved in professional responsibility, fiduciary duties, conflicts of interest, and so on. In *Seeking Virtue in Finance*, JC de Swaan focuses on an exemplarist version of this approach, providing exemplars in finance for others to emulate. He has, for example, a vignette about Erin Goddard, a Canadian accountant who moved to Rwanda to start an organization dedicated to expanding the local population of qualified accountants, having seen that a limiting factor in non-profit development in Rwanda was local accounting and auditing expertise.[100] In this analysis, agents are brave, temperate, generous, etc. in a financial context, and the purpose of financial markets is usually taken to be the distribution of capital. Risk management is reduced to a necessary procedure involved in the allocation of

[98] Mehrling, "Financialization and its discontents." Finance and Society 2017.
[99] And even under a scenario of population stagnation or population decrease, mechanisms would need to be in place to manage that stagnation or shrinkage in the economy.
[100] de Swaan, *Seeking Virtue in Finance,* 182.

capital, but its uniqueness and the relation of finance to markets for products and services is not discussed.

A different kind of virtue ethics view tries to articulate the purposes and goals of an entity or practice in order to critique and improve it on those grounds. The value of such an approach has been more readily accepted by philosophers, at least historically, with less uptake in other disciplines. Some economists, however, have begun to appreciate using this virtue framework– notably Bruni and Sugden. Their view, explicitly inspired by neo-Aristotelian virtue ethics, articulates the *telos* of a market as mutual benefit.[101] This account gives us a plausible normative framework for evaluating markets for goods and services, where the benefits of exchange can often be assessed in relatively direct terms. This is more difficult to apply to financial transactions however, because they are fundamentally motivated by uncertainty about the future state of the world, and their consequences may only become visible long after the point of exchange. If we evaluate financial markets solely in terms of the mutual benefit of participants at the moment of exchange, we will be overlooking the role finance plays in distributing uncertainty and shaping future economic states (e.g., the goods available on the future exchange market).

Another approach rooted in neo-Aristotelian views is advanced by Martha Nussbaum and Amartya Sen. Although they do not analyze financial markets specifically, their *capabilities approach* does address markets more generally. The capabilities approach focuses on the substantive opportunities people have to achieve what is necessary for human flourishing. On Nussbaum's account, human well-being consists in a set of central capabilities such as the ability to live a healthy life, participate in political and social relationships, and exercise practical

[101] Bruni and Sugden, "Reclaiming virtue ethics for economics," *Journal of economic Perspectives* 27, no. 4 (2013): 141-164 141-64.

reason. She is clearly developing her account to try and remain sensitive to cultural differences in what constitutes the good life, so the basic capabilities she highlights are those she sees as transcending any reasonable cultural difference. This approach is really interesting because it places demands on institutions to be set up in a certain way (e.g., the way that would enable capabilities in the population). This allows a capabilities theorist to normatively critique institutions. Thus, this approach gives us a way of understanding ways that market-based economies can generate inequality of capabilities made available to those on different ends of the supply chain. The fruits of globalization, for example, have been inequitably harvested by countries that are disproportionately western and already well-developed.

Sen suggests that "there is no escape from the necessity of critical scrutiny" of markets and what they generate.[102] In practice, Sen argues, this means discarding views that contain a single-minded focus on liberalization or getting prices right. He agrees with classical economists that markets are well-suited for the allocation of private goods (that is, goods owned exclusively by an individual, like apples and shirts), because individual ownership and exchange drive efficient outcomes.[103] However, he points out that markets are not geared towards the provision of public goods (such as living in a disease-free or reduced environment), and he uses this point to argue for "the provisioning of public goods, going beyond what the private markets would foster."[104] Sen, starting from this acknowledgement of market impurity, takes a kind of *wait-and-see-and-then-reform* approach. We ought to consider markets as they are, what their effects are (for good and ill, both related to straightforward 'welfare' considerations along with other political and ethical ones), and adjust accordingly.[105] Here, I interject that the provision of risk

[102] Sen, *Development as Freedom,* 126.
[103] Here, he's thinking of efficiency as any pareto optimal outcome
[104] Ibid., 128.
[105] Ibid.,122-130

management is something that the national federal governments of many countries do. But there is little attention paid to such provisions, and they are not often seen as kinds of public goods.[106]

In a similar vein to Aristotle, Adam Smith's moral philosophy preceded his work on political economy. *The Theory of Moral Sentiments* precedes and grounds the arguments of *Wealth of Nations*. As Glory M. Liu argues in *Adam Smith's America*, later uses of Smith in defense of a particular strand of laissez-faire capitalism can be distinguished from a more historically accurate interpretation of his political economy grounded in his moral psychology and accounts of sympathy, sociability, and moral judgement. Despite centuries of debate around the 'Das Adam Smith problem,' Liu strongly argues that *Wealth of Nations* cannot be understood independently of the framework Smith develops in *Theory of Moral Sentiments.*

Interpretive disputes notwithstanding, it is sufficient for present purposes to observe that Smith addresses both the moral character of individuals and the role of markets within society as mutually constitutive concerns. He recognized the transformative power of markets for products and services and was also attentive to the resulting social consequences that arose out of them. Historically, it makes sense that Smith is mostly remembered for his focus on economic activity related to agriculture and production, as well as views on emergent coordination and efficiency (the 'invisible hand') and the need to create savings ('wealth') in order to increase production in the future. While Smith does address finance issues when he discusses saving, taxation, trade policy, etc., these discussions are difficult to extend to features of modern financialized economies and have not been as influential as his discussions of markets for products and services.

[106] These provisions include the implicit backing of most mortgages by the federal government, the existence of 'too big to fail' institutions. Of course, there has never been any referendum or voting on whether taxpayers would like their dollars to be spent this way when many see as a valid criticism such risk-mitigation measures. I analyze this issue in more depth in the following chapter.

For Smith, specific features of a poorly functioning market system generate specific harms or failures. Self-serving producers[107] and monopolies undermine benefits to consumers, and so must be checked. The interests of producers should only be accommodated insofar as they yield the promotion of the consumer's interests.[108] Smith recognizes means for protecting consumers, at least partially. The emotion of sympathy should be fostered, since we are disposed to quasi-worshipping wealth which is "the great and most universal cause of the corruption of our moral sentiments."[109] This is important because an all-out pursuit of wealth is bad.[110] He is concerned that although humans have a longstanding history or charity and philanthropy, it is rare that poor people have their "fair share of the wealth produced by society."[111]

Debra Satz has a compelling analysis of Smith. She argues that he rejects *laissez-faire* economics for three reasons: parties in exchanges have idiosyncratic vulnerabilities; markets can shape the parties in the exchange themselves; and there are social goods the market is unlikely to supply. She critiques views of both general egalitarians like Thomas Nagel (who think the any redistribution should occur on a tax and transfer basis, opposed to limitations on the scope of market mechanisms) and specific egalitarians (who think there's a subset of scarce goods that should be distributed in kind to all equally) as missing something the classical economists knew– that market interactions can influence the preferences of participants in the market and exacerbate inequalities.[112] If markets themselves can shape the wants and desires of agents, then mere redistribution of goods based on existing preferences or allocating scarce goods equally

[107] 'Self-serving producers' is redundant phrasing since all humans tend toward self-interest, on Smith's view.
[108] *Wealth of Nations* Book IV, Ch VII. And on exploitation, Smith writes, "Exploitation "What are the common wages of labor, depends everywhere upon the contract usually made between those two parties, whose interests are by no means the same. The workmen desire to get as much, the masters to give as little as possible."
[109]Smith, *Theory of Moral Sentiments* Pt1, III, Ch IV
[110] Trachtenberg (2020) articulates this as the difference between necessities and convenience
[111]Smith, *Theory of Moral Sentiments* Part VI, Section 1 chapter 1
[112] See Murphy and Nagel, *The Myth of Ownership: Taxes and Justice* and Walzer, *Spheres of Justice*

won't address all structural inequalities. Without attending to how preferences emerge from the social and economic context in which people exist, redistribution may not correct the inequalities we want to fix.

Satz herself is focused on the ways markets can reinforce or erode a liberal social order, and she analyzes the dynamics between equality and market interactions. Her motivations are to analyze markets for products and services that might yield what she calls 'noxious' markets.[113] Noxious markets are those which meet at least one of the following four criteria: if they tend to include participants with 1) vulnerability or 2) weak agency, or the market tends to result in an extremely harmful outcome to 3) the individual or 4) society.[114]

It should be noted that egalitarians are largely concerned about distribution (as opposed to production) issues, although not exclusively focused on material goods alone. There is a way to construe my project as one that could arguably expanding the egalitarian projects. Distribution is not the only concern. Living in a society with a poor mechanism for managing growth and risk is certain to infringe on stability, whether we classify stability as a civic virtue or simply a desirable state of affairs requisite for fulfilling other important human goals. Uncertainty is a harbinger of both economic downturns and social strife.

John Rawls is a popular philosopher who has analyzed major human institutions through the lens of justice, and that finance is an obvious area of human life in which injustice may occur, so I should say a few words on how Rawlsians might read my project. Rawls himself does not extensively discuss finance, with the exception of commentary on concerns about

[113] Some examples of the kinds of markets who meet her criteria for being noxious include markets for prostitution, organs, and drugs.
[114] Interestingly, a market being noxious, according to Satz, does not entail that the market should be banned outright. This is because total prohibition of a certain market could worsen the noxious aspects of the market (e.g., making prostitution illegal may end up selecting for even more vulnerable participants than otherwise).

intergenerational wealth transfer (e.g., a distributive concern). Contemporary finance, though, performs functions that directly shape citizens' outcomes in virtue of their risk distribution, credit access (or not), and intertemporality. So, it would be incorrect to treat finance merely as a domain of private exchange operating within rules governed by justice. Finance seems to clearly be a set of partially constitutive pieces of the institutional framework or 'basic structure' within which distributive outcomes are generated. This project won't disagree with Rawls, but like many of the moral philosophers, a view from Rawls of finance will be insufficient in virtue of its focus on distributional issues.

Most of the philosophers surveyed here see financial markets and the issues with them as part of a distribution problem: they reward some (often the lucky or the wealthy) and damage others (often the resource poor) in a way that does not seem to track justice. Of course, there are these sorts of distributional issues in finance, but there are others as well. Financial markets can create and exacerbate crises, which have obvious detrimental effects on the well-being of society as a whole. In general, some degree of financialization is necessary to stave off or mitigate crises, because there are always risks associated with economic activity. Financial markets give us a way to respond to and mitigate those problems. So, if finance is an institution that is badly instantiated, we need to find a way to make the institution better at fulfilling its function, not to do away with it entirely.

Polyani, for example, takes crises as an inevitable output of financial markets. Discussing the depression of 1873-76 and the Great Depression, he regards both as inevitable, emergent, and not something we can do much about. He acknowledges that we "suffer grievously" from "dislocations involved in any sudden adjustment to the needs of foreign trade or foreign exchanges," (in the context of speaking about imperialism and protectionism as a trade policy),

but he doesn't speak to either the fact that there are necessary economic activities (e.g. food production) that necessarily come with risks, or the manner in which any institution (but namely finance) might have a positive influence on alleviating the grievous suffering experienced by those subject to dislocations.

In an earlier section on productive organizations, he mentions communities could not withstand the abrupt changes in the price level necessitated by the existence of stock exchanges unless this shock was dampened by central banking policy. His focus in this section is on bankers and their role in orchestrating price levels of currency as a refutation of the idea that markets are self-regulating. However, he misses the point that it is not the market's lack of self-regulation that necessitates stabilizing (e.g., risk mitigating) activities. 'Stability' is an important goal, but it is one that will need to be deliberated about no matter what role markets play in our society, or the degree to which they are embedded into the social fabric of our society.

Many of these thinkers engage a similar narrative framing, using major financial crises like the 1929 and 2007-'08 Wall Street crashes and subsequent depressions as rationales for several, often interconnected, points: Polyani in particular, uses this example to justify the need of an independent central bank, arguing that the "ruinous effects" of sudden price changes should be "cushioned by the means of an independent central banking policy."[115] Most frequently, explanations of crises expound upon the ways in which financial markets contribute to additional risk and uncertainty. This seems to make sense because there are clearly things that have gone wrong, and people are made worse off as the result of financial crises, depressions, and recessions. However, I would suggest we view those periods as examples of finance providing us with a failed solution, as opposed to causing the issue of increased risk. The risk is there, not

[115] Polyani, *The Great Transformation*, 208.

depending on the degree to which finance occurs in society. Some risks are realized as losses, and unrealized risks are returns or a continued state without change. Finance in society is the result of individual or collective pursuits of higher growth. In cases where things go wrong, it is easy to collapse this into the simple causal explanation of, 'the existence of our financial system has yielded these bad things'– but the problem is finding a counterfactual world where the bad thing doesn't come about without also removing the positive risk-mitigating and growth-enabling features of finance. Speculation is an easy target in finance, being an 'unproductive' behavior that is only sought after for the sake of wealth creation. The difficulties with preventing 'mere speculation' in markets are twofold: First, an action taken cannot easily be discerned to have been taken for the sake of speculation or something else. Second, there are cases where pure speculation for the sake of monetary reward is undertaken by one party to a transaction, the counterparty to which has some productive social end.[116]

That depressions and stock market crashes are used as examples to argue for diametrically opposing viewpoints makes it difficult to present these crises as unassailable reasons to implement any particular solution. In an application of Smith's framework, it was argued using historical data that financial markets were *over*regulated, claiming that if only markets were really left alone, they would have reduced the negative externalities of the Great Depression.[117] If this strikes one as ridiculous, than so too should the assertion that financial markets are the sole culprit of systemic crises, and that by doing away with the former we can eliminate any occurrence of the latter.

[116] Such as those cases of commodities futures trading which operate as *de facto* insurance for farmers.
[117] Friedman, Milton, and Anna J. Schwartz. *A Monetary History of the United States, 1867-1960*. Princeton: Princeton University Press, 1963.

These are all philosophers who start from some conception of morality, and apply it to issues in markets, some specifically in the context of finance and other in the more general context of markets for goods and services.[118] These thinkers can see clearly the impact that markets have in society. All of them accept that markets are both beneficial and harmful in various ways. None of them, however, seem particularly concerned about the negative consequences of sudden shocks (or risks thereof) in our economic system. Of course, the development of financial markets was not entirely relevant in the time periods these thinkers are writing. But there would have been very relevant systemic risks posed by industries that would have been of pertinent concern in their time. Weather-related famine, natural disasters, and other exogenous shocks which certainly have the potential to beget the same harms each philosopher is concerned with.

Thread Two: Hayekian Economics

Modern economics, dominated in the United States by *laissez faire* views, exemplified by 'Chicago School' thinking has several major tenets relevant to my project. I will not generally address the validity of this approach or its presuppositions, as regardless of its true nature, it has been, and is, the dominant economic ideology of American economic policy making.

As modern economics evolved as a distinct discipline from political economy, there emerged (at least) two strongly opposing lines of thought. For the sake of simplicity, I will call these as Hayekian views and Keynesian views.[119] Very roughly, those I consider to be Hayekian

[118] Interestingly, Gauthier does something like the opposite: he starts with a conception of economically rational behavior which grounds his moral theory. Nozick too, in his political philosophy seems to start from certain rational game theoretic conceptions about rational man's action, which may or may not entirely line up as an accurate portrayal of how humans behave.

[119] To be clear, the field of economics contains multitudes, and most economists would find this bipolar view overly reductive. More accurately, there are Austrian views, which can be separated into at least two camps (Mises and Hayek). From there, the Austrian views more generally inspired the Chicago School view in the mid 20th- century.

thinkers that markets function as self-organizing systems that efficiently coordinate disparate knowledge, and that attempts at central planning of regulatory intervention create unintended consequences that muddle the well-functioning of markets because singular authorities can never access all relevant knowledge that gets embedded in markets because of prices which emerge from individual market agents' actions. Those that I consider to be Keynesian generally agree that markets can be unstable, and that government intervention is a useful way to stabilize the economy for the sake of preventing recession and reducing periods of prolonged or high unemployment.[120] The uptake of issues from economics into philosophy, then, has often traced the ideological boundary in economics. We can see, for example, the economic delineation in libertarian rejections of egalitarian concerns about the scope of markets (e.g., Satz) on the basis of individual freedom and allocative efficiency.[121] This intellectual history helps explain the tendency of the philosophers I've highlighted to be focused on discussions of distributions in markets for goods and services. That part of this history is required for the sake of covering the bases of what has been talked about but does not get into the features of financial markets I want to draw out. The part of economic history involving monetary policy and macroeconomic accounts is of interest to my account.

The classic Hayekian view is skeptical towards regulation, as they take markets to work naturally, invisible-hand style. Markets are supposed to be good resource allocators. Individuals

---

See Liu (2022) for an account of economic history that avoids the crime of such heavy-handed reductiveness as I have committed here.

[120] An under-appreciated point here is that a Keynesian needn't disagree with the Hayekian information theory's essence: that prices are useful because they aggregate information. Keynes, as I'll describe below, however sees this as a mechanism susceptible to human psychology, and thus is not something sacred to be preserved from state influence or intervention.

[121] See Brennan and Jaworski, *Markets without Limits: Moral Virtues and Commercial Interests.* 2022 for a direct counterargument to Debra Satz. They argue that if it would be morally permissible to give something away for no payment, it is morally permissible to sell that thing, and the morally problematic features of exchange Satz highlight rest in the background conditions of the market or the thing rather than in the marketization of it.

act rationally to maximize their welfare, buying things they want, selling things they make, and only paying prices they are willing to in the event of exchange. Firms act this way too, in order to maximize their profits. Prices dynamically balance supply and demand, ensuring that goods and services are allocated to entities with the greatest effective demand. But markets are only supposed to work well, however, so long as there is not stifling intervention from the state.[122]

The Arrow-Debreu general equilibrium model puts into discrete mathematics the kind of picture painted by Adam Smith's invisible hand: given a world with perfect competition and completeness of markets, their model delineated prices at which each good is bought and sold in equilibrium (i.e., the goods that are produced are exactly all the goods that are purchased). *Perfect competition* entails the following: no interdependencies amongst agents' utility functions (e.g., my utility doesn't go down just because yours goes up); markets for goods and services are complete, such that all utility functions can be addressed in the market; no barriers to entry into or exit from a given market, so that no one is 'trapped' in a market; and a sufficient number of traders such that no single agent may influence prices (i.e., no one has market power). Under conditions of perfect competition, market transactions will be *Pareto improvements*: the results of a transaction will be strictly preferred by at least one person, and no one will strictly prefer the pre-transaction situation. As a result, the equilibrium allocation that results from a market under conditions of perfect competition will be *Pareto optimal*: no further Pareto improvements will be possible. This result is known as the First Fundamental Theorem of Welfare Economics. *Completeness of markets* entails that transaction costs are negligible, there is perfect information, and that everything that can be bought or sold has a price.[123]

[122] Such interventions include but are not limited to price controls, restrictions on entry and competition, subsidies, and regulatory mandates that increase costs.
[123] As philosophers, we might get ahead of ourselves here. For the sake of the terminological definition as used in economics, the things which can be bought or sold are limited to what are traditionally classified as assets.

Hausman and McPherson sketch out a definitive summary of the assumptions of the standard welfare economics view. The prototypical welfare economist will agree with the following: Economics should appraise outcomes (as opposed to processes) using a single method of appraisal (as opposed to using multiple perspectives depending on the issue at hand). What should matter in such appraisal are the consequences for individuals (as opposed to the consequences for groups or for some other thing like the environment). What we ought to care about are not values like freedom, rights or justice, but rather individual welfare, which are taken to be constitutive of the satisfactions of one's preferences (as opposed to mental states like happiness or some set of objective goods). Even neoclassical economists will admit that the Arrow-Debreau equilibrium model omits important factors in the real world of finance, such as possibility of default, the limits of collateral requirements and market friction (namely, transaction costs).[124] Of course, these three major factors all pertain to risks. The risk that counterparties will not fulfill their agreed upon obligations is the risk of default. A counterparty's collateral requirements are an institution's assessment of their riskiness. And market friction often (but not always) arises due to the untrusting marketplace which contains a few loci of trust (in the form of institutions like clearing houses).

Hausman and McPherson's research is unified in its attempt to explain the ways in which mere preference satisfaction can be inadequate for maximization of welfare. They disagree with the common notion in economics that it can exist as a value-neutral discipline. In rejecting the positivist claim that economics can be wholly objective, they describe how we insert normative value judgements into calculations or conceptions of utility, welfare, and efficiency. For them, the proverbial smoke is values, and where there are values, there are ethical

[124] Geanakoplos, 2018 "Promises, Promises" and 2010 "The Leverage Cycle"

questions (the fire). They take aim at the first welfare theorem of markets, arguing that although markets are, indeed, institutions which facilitate the satisfaction of preferences, those preferences are not reliably welfare-improving.

Sen also has work that centrally concerns the methodological issues to this section. In his *Royer Lectures on Economic Behavior and Moral Sentiments* delivered to Berkeley, he presents a graphic description illustrating his key disagreement with neoclassical economic models, arguing that "to try to use the demands of rationality in going to battle on behalf of the standard behavioral assumption of economic theory is like leading a cavalry charge on a lame donkey."[125] He finds it disagreeable that economics neglects the influence of ethical considerations in the characterization of human behavior, as it does not require any concessions to the contrary (e.g. that humans, acting sometimes in non-self-interested ways, always act selflessly).[126]

Thread Three: Keynes and Psychology

The Keynesian view contrasts the Chicago School most notably in its understanding of the role of the state in regulation and monetary policy. Keynes departs significantly from the Hayekians in both his descriptive account of markets and his normative assessment of how public institutions should influence them. It is not easy to identify a unified 'Chicago School' position on monetary policy, as there are a number of variations (ranging from commodity-backed money, like gold, to Friedman's monetarism). Generally speaking, however, economists influenced by Austrian traditions tend to prefer monetary regimes with less discretionary

[125] Sen, *On ethics and economics.* 1987, Royer lecture on Economic Behavior and Moral Sentiments, 16.
[126] Sen, *On ethics and economics,*19.

intervention by the state and reject centralized monetary policy that sets interest rates, does quantitative easing, etc. as some officials judge appropriate.[127]

Keynes recognizes that individuals have the freedom to choose whether to spend or save their money. He spends considerable attention on the relationship between inducing investment and people's propensity to save money. In an economy where investment is required for innovation, a population's propensity to save too little or too much can be a concern for future economic growth.[128] The macroeconomic concern is that increased saving reduces current consumption (people are leaving money in their bank account instead of giving it to businesses), which can lower aggregate demand. This theoretically in turn weakens firms' expectations about future profits (e.g., expectations will be revised down). So, the trouble that additional savings by consumers spells for investment is more psychological (it's about expectations) than about availability of capital (which is, in theory, being stored in a bank and which can be deployed by banks to finance loans). In an effort to explain macroeconomic instability and provide policy recommendations for appropriate investment, Keynes challenges several assumptions about what determines interest rates and prices.

Classical theory held that interest rates were determined by the supply of savings with relation to the demand for investment at a given time. Keynes' liquidity preference theory instead argued that interest rates are primarily determined by the demand for liquidity (the desire to hold money or money-like assets) rather than by the aggregate stock of savings available for investment. There is a root issue here about what the mechanism is whereby individuals convert

[127] Quantitative easing (e.g., large-scale central bank purchases of longer-term government bonds and mortgage-backed securities using bank reserves, intended to lower long-term interest rates and increase liquidity), the other tool in the U.S. Central Bankers toolbelt, is far more controversial for Chicago School (and Austrian) economists.
[128] This was borne out in the 21st century in China, where the population demonstrated a strong capacity for saving overspending. See: Chamon, Marcos D., and Eswar S. Prasad. "Why are saving rates of urban households in China rising?" *American Economic Journal: Macroeconomics* 2, no. 1 (2010): 93-130.

savings into investment. If it were the case that savings naturally turned into productive investment, encouraging additional savings would necessarily be bolstering investment. Keynes disagrees with this mechanistic explanation, and intuitively it seems plausible that many individuals save simply to save, without any natural 'trigger point' at which they convert some portion of that savings to investment. Keynes instead argued that one's motivation to invest depends on expectations about future profitability.[129] This perspective entails that should a government wish to induce more investment, they should adjust interest rates and increase demand via fiscal expenditure.[130]

Throughout his work, Keynes operates with a distinct psychological conception of human behavior. At moments, his account echoes that of Hume, recognizing that human beings possess potentially destabilizing dispositions that still may be socially beneficial when redirected toward productive economic pursuits.[131] So if we imagine individuals who have a desire to obtain status or a sense of superiority in society, they will require some outlet for this disposition. Absent a market economy in which they can 'dominate' by, say, working hard to make superior trading

[129] Again, this is less about the expectations about future profitability that the consumer/saver has, and more about the expectations that those with the authority to invest the funds have. So, this ignores whatever individual level investment that a single consumer could do (e.g. perhaps they're saving in order to self-fund a business of their own).

[130] The basic mechanism works as follows: When interest rates are low, individuals cannot earn much return by holding money in highly liquid forms such as savings accounts. When interest rates are higher, individuals can earn a reasonable return simply by leaving funds in savings, making them less inclined to seek out riskier investments. Fiscal expenditure aims to stimulate aggregate demand by increasing overall spending in the economy. This can happen via government purchases (of bonds, etc.) or direct transfers of money that raise household consumption (e.g., COVID checks). As demand rises, firms observe stronger sales and improved expectations about future profitability. These improved expectations create incentives for businesses and potential investors to expand production or undertake new projects, thereby inducing greater investment.

[131] On page 324, Keynes writes: "dangerous human proclivities can be canalized into comparatively harmless channels by the existence of opportunities for money-making and private wealth, which, if they cannot be satisfied in this way, may find their outlet in cruelty, the reckless pursuit of personal power and authority, and other forms of aggrandizement."

and investment decisions, this disposition might find its outlet in violence or other destructive behavior.[132]

Another interesting psychological element of Keynes' approach is his exploration of the motivational forces or 'inducement' necessary for someone to invest.[133] In his chapter on long-term expectations, he describes the psychological basis for price changes in financial markets. He famously compares the stock market to a beauty queen contest, in which "we devote our intelligences to anticipating what average opinion expects the average opinion to be."[134] Price formation is, on this view, often a mere reflection of recursive expectations about others' beliefs rather than a straightforward assessment of underlying value. Keynes expresses concerns about this sort of speculation but argues that it can enhance liquidity and price discovery.[135] For Keynes, the danger arises when speculation dominates motivations for productive enterprise. Speculators are harmless insofar as they are "bubbles on a steady stream of enterprise," but can destabilize markets when "enterprise becomes the bubble on a whirlpool of speculation."[136] Although he is less worried than some, Keynes is worried about the role of speculation, especially in more developed financial markets.[137]

Keynes' concept of 'animal spirits' is central to his psychological explanation of markets, which he describes as a spontaneous urge toward action rather than the outcome of calculated

---

[132] An ethicist might object to the idea of finding outlets for negative dispositions and might argue instead that one ought to assist the moral development or evolution of this individual to not have such a disposition in the first place.
[133] "There has been a chronic tendency throughout human history for the propensity to save to be stronger than the inducement to invest." [this is a problem] "Today the explanation of the weakness of this inducement may chiefly lie in the extent of existing accumulations; whereas, formerly, risks and hazards of all kinds may have played a larger part.", 301
[134] Ibid.,134-135.
[135] Both of which remain the strongest arguments for a variety of otherwise non-value additive market action like that taken by high-frequency trading algorithms
[136] This metaphor bears striking similarity to that of expressing concern via the phrase "the tail is wagging the dog." Keynes mentions a radical solution, which even he seems uncertain of, which would be to make financial investments permanent and indissoluble, as this would eradicate the profitability of short-term thinking.
[137] "As the organization of investment markets improves, the risk of the predominance of speculation does, however, increase," 136.

reasoning.[138] Investment decisions are not reducible to rational calculations, on his view, because there is uncertainty about future outcomes.[139] But without this non-rational (or perhaps 'irrational') element of investment decision-making, investors might never undertake projects whose risks include the possibility of loss, despite potential social benefits and financial gain.

In his view this irrationality is required however, for someone to put aside the genuine risks "of ultimate loss which often overtakes pioneers" in order to induce one to invest.[140] In several sections, he talks about risk in the context of the borrower-lender relationship. He provides a taxonomy of three kinds of risk: First is borrower's risk, arising from uncertainty regarding whether anticipated returns will materialize.[141] Second is lender's risk, which reflects doubts about repayment and represents an additional cost introduced by the separation between borrower and lender.[142] Finally, Keynes identifies a monetary risk associated with adverse changes in the value of money itself, which may render monetary loans less secure than ownership of real assets.[143] Keynes' discussion of this third category is vague, but appears to concern inflationary or monetary instability that in some way alters the real value of contractual obligations.

An interesting insight is that riskiness is baked in at several points in the financing of an investment[144]: For example, a risky venture will require the borrower to demand a greater

[138] On page 139 he speaks of two kinds of stability, that due to speculation and that due to our animal spirits or "a spontaneous urge to action rather than inaction, and not as the outcome of a weighted average of quantitative benefits multiplied by quantitative probabilities."
[139]The distinction between 'risk' and 'uncertainty' is discussed further below.
[140] He compares this to the way "a healthy man puts aside the expectation of death."
[141] On page 124 he describes the borrower's risk which "arises out of doubts in his own mind as to the probability of his actually earning the prospective yield for which he hopes."
[142] Here I think Keynes means that the risk that would exist would essentially collapse or fold into the Lender's risk
[143] He describes a risk of "a possible adverse change in the value of the monetary standard which renders a money-loan to this extent less secure than a real asset; though all or most of this should be already reflected and therefore absorbed, in the price of durable real assets."
[144] Keynes, on pages 124-125, he calls this "in part a duplication of a proportion of the entrepreneur's risk, which is added twice to the pure rate of interest to give the minimum prospective yield which will include the investment."

expected margin between anticipated yield and borrowing costs, and the lender will charge a higher premium above the pure rate of interest.[145] Boom periods occur when both parties underestimate these risks in tandem. Despite this valuable insight, Keynes avoids extrapolation from the single-investment lens into a discussion of systemic risks, or the role of investment in either mitigating or intensifying broader risks borne by society.

If Kenneth Arrow can be seen as formalizing aspects of Adam Smith's invisible hand, Richard Thaler can be seen as formalizing, by way of empirical evidence, aspects of Keynes' psychological model of human behavior. Behavioral economics focuses on observed (rather than rationally theorized) human behavior in economic decision marking. Thaler's account of economics combines the cognitive psychology literature, most notably popularized by Kahneman and Tversky, with experimental evidence. When real humans buy, sell, borrow, and save, they are less economically rational than traditional economics would suggest. Thaler contrasts 'Econs'– the characterization of the perfectly rational actor– from the real human beings who substantially depart from theoretically predicted rational behavior. Because policy design relies on assumptions about human behavior, erroneous assumptions have serious consequences for important social economic outcomes like retirement planning, credit markets, loans, and monetary policy. That human 'misbehavior' is predicated by these behavioral models is little consolation if the financial structures are predicated on what rational Econs will do and there are grave consequences if wrong.

Thaler does talk about risk at an individual level, as a 'loophole' that efficient market hypothesis supporters use to wave away his concerns about the intrinsic flaws in the efficient

[145] "for is a venture is a risky one, the borrower will require a wider margin between his expectation of yield and the rate of interest at which he will think it worth his while to borrow; whist the very same reason will lead the lender to require a wider margin between what he charges and the pure rate of interest in order to induce him to lend (Except where the borrower is so strong and wealthy that he is in a position to offer an exceptional margin of security)."

market hypothesis.[146] However, he doesn't speak to the risks that are necessary or beneficial for society to take on at a collective level. In his estimation, the financial crisis of '08 was caused by our flawed models of human rationality. But it is unclear, if we were to someday rid ourselves of the flawed picture of agents as rational Econs, what kinds of risk we should take on or how we might manage them.

Behavioral economics also reveals variation within populations. Experimental evidence suggests that individuals with extensive economic training are more likely to behave in ways consistent with rational choice predictions.[147] Whether this reflects selection effects or educational conditioning is still up for debate. If economic reasoning can be cultivated through education, however, then limitations in risk assessment need not be treated as permanent features of human psychology. Later in this project I argue that determining the socially appropriate level of risk (particularly those undertaken in the pursuit of economic growth) should be subject to deliberative democratic processes. A straightforward objection to this would be that such a system would invite uninformed participation. Yet the possibility of public error does not foreclose the possibility of public education to better individual understanding of risk, as supported by this literature. Institutional design can include mechanisms that cultivate more sophisticated public reasoning about risk.

There is a contrast between financial markets and markets for products and services that emerges in the market-design work of Al Roth. Roth emphasizes the importance of market "thickness," timing, and congestion in determining whether markets function effectively. Many problems he addresses arise precisely because traditional markets fail to manage risk and

[146]Thaler, *Misbehaving*, 225-229.
[147]Camerer, Colin, and Richard H. Thaler. "Anomalies: Ultimatums, dictators and manners." *Journal of Economic perspectives* 9, no. 2 (1995): 209-219.

coordination problems adequately.[148] Although not framed in these terms, his work could be described as making non-finance markets more like finance by importing financial market techniques into non-financial contexts and non-exchange traded forms of 'exchange.' Roth, in other words, can be understood as being motivated by risk reduction. By applying liquidity-enhancing measures, matching algorithms, and game theory, he is using institutional design to mitigate uncertainty faced by participants in a market. This is not an approach to apply universally, though. While efficiency might theoretically improve if, say, restaurant diners submitted ranked meal preferences in advance, such a system would likely undermine other social features of dining whose value exceeds marginal gains in coordination efficiency. Still, Roth demonstrates how market design can be used for risk mitigation and bears an interesting contrast to Satz and Sandel in that he seems to demonstrate ways in which arenas of human life can be *improved* by marketization.

Seven Common Criticisms of Both Kinds of Markets

After a tour through the literature, it is clear that there are several popular criticisms launched against financial markets. These critiques are similar to analogous critiques made about markets for products and services. Most of the criticisms detailed below have something to them – this section is not to make light of the concerns they raise nor to disregard them as wrong or unimportant. Instead, I mean to make clear that this set of concerns is not exhaustive and does not address some of the unique features of financial markets. Most of these views treat finance

[148] For example, school choice systems previously suffered from congestion caused by sequential decision-making. Allowing families to rank preferences simultaneously enables clearinghouse mechanisms to produce stable matches without strategic delay. Kidney exchange programs allow incompatible donor-recipient pairs to be matched across a larger network. One patient's donor can then be given to another recipient (who they likely don't know personally) and another donor (who the patient likely doesn't know) is willing to provide their compatible kidney to the patient. This increases efficiency of matches and thus the number of successful transplants.

'as is', as a 'black box', one whose inner workings to which we do not have access. They might seek to change the inputs (attempting to engender virtue in future finance professionals, for example), and they often seek to change the context of the outputs (e.g., by adding a layer of taxation or universal basic income). There are views that attempt to alleviate wealth inequality not by changing how finance works, but by taking for granted what it does, and fixing it after the fact. Of course, finance is not a 'black box'. We actually can look 'into' it to see what is going on, and we need not only control the inputs and the context of the outputs to fix what it might get wrong. But that is how these criticisms proceed.

Figure 3: The Seven Common Criticisms of Financial Markets

#	**Criticism**	**Popular Critics**
1	Financial markets infect areas which should not be financialized	Michael Sandel, Karl Polanyi, David Graeber
2	Finance is treated as a science when it is not	Donald MacKenzie, Hyman Minsky, Robert Shiller
3	Financial actors are treated as rational agents when they are not	Daniel Kahneman, Amos Tversky, Richard Thaler
4	Finance incentivizes bad behavior	Michael Lewis,
5	Finance breeds a culture of bad behavior	John Kay

6	Financial markets yield inequality	Joseph Stiglitz
7	Financial markets are vicious	Rana Foroohar

What follows is a taxonomy of frequently cited issues that apply to financial markets. These theorists often treat them as disputes about state intervention versus individual responsibility, which I think is an unnecessarily limited lens through which to address the underlying issues. Analyzing these criticisms in the context of markets for goods and services highlights key differences, reinforcing the claim that financial markets warrant distinct theoretical treatment.

Financial Markets Infect Areas which Should not Be Financialized

These criticisms are connected to a few of the others below but have a distinct feature in that they model financialization as a sort of infectious disease. In other words, even if financial markets themselves (or finance as an academic discipline) were to be in good standing, the attitudes which finance inspires us to take on orient one towards expanding the reach of financialization into arenas of human life that corrupt those areas in some way. As a criticism of markets more generally, this view has received much attention in philosophy and in other disciplines, with popular works by Debra Satz, Michael Sandel, and Robert Kuttner. The markets for products and services version of this 'marketization' criticism says markets for some products and services erode some facet of the human experience by rendering it paradoxical or incoherent. Examples of the former include markets for organs, sex, or weapons of mass destruction. The latter would look like having a market for friendship or love. When we use the term 'marketization', we usually mean standard market procedures applying to new products or

services. The term 'financialization' usually means applying financial thinking, including an abstraction of some risk or uncertainty which we price in (often in conjunction with those things described as marketization). Viaticals, for example, are agreements in which an individual sells their life insurance policy to a third party in exchange for an upfront lump sum payment. The individual who sells their policy then has the benefit of immediate liquidity, and the third party receives the future financial benefits upon their death. Viaticals are commonly used for individuals who need immediate liquidity to pay for medical expenses, and there is a clear sense in which this practice can been seen as predatory or coercive due to other compounding injustices (i.e., since the healthcare system is unjust, an individual has restricted choices including entering into this sort of agreement). Viaticals do seem to have a potentially exploitative element due to the perverse incentives it creates. Even apart from outcomes, viaticals can be seen as unethical in how they degenerate an aspect of someone's humanity.[149]

Both marketization and financialization can be harmful, but it is important to distinguish between these concepts. The existence of markets is often tolerated, at least to some degree, by most theorists. On the other hand, financialization is more frequently viewed as wrong in every case. In other words, one more frequently encounters the argument that nothing should be financialized than one encounters that nothing should be for sale. [150]

The criticism that financial markets *infect* other areas is commonly found in the humanities and social 'marketization' sciences literature. When these scholars criticize financial

[149] There are other kinds of insurance swaps either between individuals and third parties or between insurance providers that have a similar structure but are perhaps less coercive or icky in the same kind of way because they are several degrees removed from the patient being insured. Or perhaps they receive less criticism simply because people don't know about them.

[150] Nancy Frazier probably expresses the strongest anti-financialization/marketization view in her book *Can Societies Be Commodities All the Way Down?* Other theories more generally support the view that marketization (financialization is talked about as specific to general marketization less frequently) generally erode as positive social norms or features. The latter view is a large tent encompassing most of the theorists mentioned in the previous section.

markets (and even more specifically, stock markets), they are usually doing so based on moral principles and theories which they bring in from elsewhere and apply to financial markets. They say little about how those financial markets ought to operate or the extent to which they think they should exist in the first place. No one much discusses which aspects of human society *should* belong within the purview of financial markets. Likewise, there is not much consideration about how to build strong enough walls between these two kinds of things that might help deter or prevent similar corruption in the future. It is one thing to argue that financial products like viaticals and prediction-market contracts are vicious or should not be allowed. It is another to clarify the appropriate role for financial markets.

'Sports betting', or prediction markets, both on traditional athletic achievements and other political, business, or otherwise pop-culture-related outcomes, exemplify the risk-calculation (or speculative) feature of financial thinking. This is discussed in more detail in Chapter V. Platforms like Kalshi and Polymarket boast that they let users wager on anything, turning one's ability to assess any kind of risk or uncertainty (or to have a propensity for not caring about those things) into a tradable asset.[151] The quantification of the minutia of human activity and a risk-calculated bed on such action strikes some as bizarre, and others as inversions or distortions of the human experience, if not in itself, for the addictive dopamine cycles such behaviors tend to give rise to.[152]

<u>Finance is treated as a science when it is not</u>

[151] Prop bets related to specifics around the Super Bowl are some of the most niche bets individuals can make. Books take bets on the length of the national anthem or what color cleats someone might wear. When Peyton Manning played his last Super Bowl the most popular prop bet was the over/under on how many times he'd scream "Omaha." as an audible. Controversially, these platforms also gave users the opportunity to bet on the outcome of the 2024 presidential election. See Levine, Matt. "Prediction Markets are a Thing Now." *Bloomberg.* November 7, 2024.

[152] Smith, Annie, Colleen Poon, Garrett Jones, and Zainab Thawer. *Understanding the odds: Gambling among BC youth aged 12-18*. McCreary Centre Society, 2021.

Critics of this sort argue that although finance is an academic discipline, and financial regulation and policy are treated in a similar way that science is, there are sufficient dissimilarities between finance and the sciences that finance should not be warranted such a treatment. Although finance includes numbers and rules that purport to explain things like what a security should be priced at, it often fails to accurately predict price moments and lacks universal laws. The subjectivity of valuation is also connected to the interrelated criticism below that the agents within finance are treated as rational agents when they are not.

This critique is a serious one, rooted in criticism of the economics of markets for goods and services. Markets for products and services do not predictably respond to exogenous shocks despite the existence of economic tools like demand curves which purport to explain consumer behavior. in the ways economists might predict. Real evidence suggests that consumer demand does not behave according to this model in, say, inflationary periods.[153]

Another example of the markets for products and services version of this problem can be seen in relation to the behavior of individuals when items go on sale, or when they at least look as though they are on sale. Many major shopping holidays, like Black Friday (the existence of which might be an example of the criticism), take advantage of this by artificially inflating the prices of items before marking them on a significant 'discount'.

In finance, this sort of criticism can take aim at any number of analytic accounts. Consider, for example, traditional portfolio theory of how publicly traded securities ought to be priced given certain 'fundamentals' of the company including earnings, profits, etc. In modern times, prices do not seem to correspond with such models, with some stocks trading at prices that represent large multiples of traditional 'ratios' of classic financial metrics (like the price of the

[153] Katona, George. "Psychology and Consumer Economics." *Journal of Consumer Research* 1, no. 1 (1974): 1–8.

security relative to that company's earnings).[154] There is also well-evidenced herding behavior.[155]

Financial actors are treated as rational agents when they are not

This criticism has become increasingly popular as the once-demonized subfield of behavioral economics has become more widely accepted amongst scholars in the economics, finance, and psychology disciplines.

Given that it is true that agents in systems do not, in fact, always behave rationally, at least in the traditional sense, what are we to do? Critics vary in the degree to which they think we must rethink the system in which sometimes-not-rational agents are operating in. To some, the solution is just to call for less hubris when presenting market claims. People tend to be rational, so markets tend to be efficient but it is all approximate and a matter of degree. Others fear that this mistaken assumption about human cognition and behavior is unassailable, and that the core of the way we have structured financial or economic markets is rotten as a result.

The markets for products and services version of this problem is backed by many examples in Richard Thaler's book *Misbehaving*. In graduate school, he experienced puzzling behavior that did not align with economic theory: A friend had been given two tickets to a basketball game, which was an hour away. But on the day of the game, there was inclement weather that made driving conditions materially worse. They elected not to attend the game due to the weather, but the friend remarked that had they actually paid for the tickets, they would have insisted on braving the bad road conditions to attend the game. Intuitively, this might make

[154] The explanation for why this may be is complicated. Many blame this dislocation on any number of things, including but not limited to the proliferation of passive investing, the increased presence of retail traders, high frequency trading and algorithm trading.
[155] The measurement of such behavior in modern times is likely also picking up algorithmic trading set up to buy or sell securities if a set of previously coded conditions are met.

sense, and it is tempting to think that in a gifted-tickets scenario, there is less 'at stake' than the one in which one had purchased one's own tickets: you are not 'wasting' any money. But according to economic rationality, this is a fallacy termed the 'endowment effect.'[156] In the moment when you're deciding whether or not to go to the game after all, either decision you make (to go or not to go to the game) isn't going to change the amount of money in your bank account. Going to the game in the self-purchased ticket scenario doesn't save you any money because you've already spent it.[157] It's also been demonstrated that people are happier when given, for example, a vacation worth $500 rather than $500 in cash as a bonus from their employer.[158] This is not what standard economic rationality would predict, since, under the assumption of stable preferences and fungibility, a rational agent should be indifferent between equivalent monetary value (they could simply use the $500 cash to pay for the vacation). To generalize– there clearly exist certain habits and tendencies that govern our behavior that are not consistent with economic rationality.

Similar behavior crops up in finance: those who remain invested in a mutual fund, as opposed to an ETF, whose fees are higher and where the former provides inferior annualized returns. Likewise, trading in meme stocks for non-financially motivated reasons such as nostalgia, desire to be a part of a social group, etc. Professional traders and money managers are not impervious to this criticism either, as evidenced by extensive research proving that herding exists in markets to an extent that is unlikely to be caused merely by algorithmic trading and/or

[156] Thaler, 1980
[157] Here, we see illustrated both the endowment effect (where one values something more because they own something or have paid for something) and sunk-cost fallacy behavior (where one commits to a behavior to justify previous expenditure, even if doing so no longer makes sense).
[158] Hidden Brain with Shankar Vedantam, "Misbehaving with Richard Thaler." October 24, 2017.

retail trader activity.[159] The basic criticism is a more significant risk to our understanding of financial markets to the extent that professionals engage irrationally too.

Individual Bad Behavior

Critics of this kind focus on the unethical actions of individuals, often coalescing around instances of fraud or mismanagement that have been uncovered by journalists and the news. Sometimes, these criticisms are used to extrapolate that more regulation of these professionals or firms is needed while others dismiss these cases as involving 'bad eggs' in an otherwise okay basket. Others use these (frequent) examples of bad behavior to dismiss capitalistic finance overall, taking the 'bad apples' to be an indicator the whole bunch is rotten. Most frequently, the treatments of unethical individuals do not make specific reference to the unique functions of financial markets. Finance fraudsters do receive extensive (disproportionately large, perhaps) because our public perception treats them as particularly egregious compared to frauds involving goods and services.[160] To be clear, any instances of fraud (by an individual or at an institutional or systemic level) ought to get attention from the public and academics. An issue arises, however, when the coverage of bad behavior disincentives theorists from diagnosing problems from first principles due to preoccupation with today's financial crime.

The process for catching and prosecuting fraud by businesses that sell products and services is more complicated than the process which yields enforcement action by the SEC. Whereas two small businesses might engage in a transaction in which one party is perpetrating

[159]Guo, Xu, Chen Gu, Chengping Zhang, and Shenru Li. "Institutional herding and investor sentiment." *Journal of Financial Markets* 68 (2024): 100891.

[160] It is perhaps true that the stakes are higher, and the ability for a fraudster to defend themselves narrower: Relative to the SEC, the FTC has fewer resources and tools to prosecute fraud. One major example of this discrepancy is the difference in recording requirements: Professionals in finance are required to have their communications recorded, which allows for more discovery in investigations. But the FTC places no such demands on businesses which (allegedly) offer goods and services to everyday people. And everyday people are less likely to bring up incidents of fraud than investors are. See: FTC Staff Report August 2004 on Consumer Fraud in the United States.

some crime or fraud, the standard of evidence for prosecution of such misdeeds means that they are not so simple to catch and prosecute. On the other hand, the SEC has a robust whistleblower incentive program (which has itself been the subject of criticism) where 'tippers' can in some cases be awarded millions of dollars for coming forward. In addition, much of the financial crime which is caught and subsequently reported has been detectable from public records in the form of public regulatory filings, mandated recorded conversations, and exchange trading data. From a normative perspective, it is unclear whether individual bad behavior is a bigger problem in markets for products and services or financial markets in terms of what that can tell us about how the respective institutions/markets are organized. Using Bloomberg columnist Matt Levine's imagery of a dial, we can think of our regulatory environment (specifically as it relates to consumer protection) as falling between two opposing ends: a maximal consumer protectionary environment with little to no fraud and lots of control and a minimal consumer protectionary environment with some to a lot of fraud and less control. There are tradeoffs to be had at any point along this dial, and it is the regulator's (and arguably the philosopher's) duty to figure out the extent of the government's responsibility in protecting consumers.

Culture of Bad Behavior

Going a step further than addressing individual instances of bad behavior are those criticisms which take issue with finance culture more generally. These usually incorporate references to the vice of greed, dishonesty, and excess. Often, these criticisms present finance as unpalatable without substantial reform, which can set them apart from the previous kind of criticism which might grant that finance itself is well-functioning, vicious individuals notwithstanding. This criticism rightly focuses not merely on instances of individual 'bad actor' schemes, but on system-wide failures, ethical or otherwise. In these criticisms there is an implicit

or explicit explanation of such misdeeds as being perpetuated or at least supported by a bad culture which promotes vices over virtues, or is too focused on the benefit of the individual rather than that of the system or society overall, etc. In markets for products and services, we might think of our stereotypes of used car salesman as emerging from a used car sales culture which places individual profit over customer satisfaction. More than other criticisms, the markets for products and services and finance examples are quite intertwined: In terms of markets for products and services we could think of, on a corporate perspective, something like planned obsolescence: companies are not actually looking to maximize customer welfare or provide the best service, but are instead incentivized to juice up quarterly earnings. Why do that? Because large corporations are increasingly financialized: they answer to their shareholders more than stakeholders like customers.

Examples of bad corporate culture in finance abound in pop culture, and even fictionalized accounts stem from underlying truths about traditional 'finance culture.' Cultural touchstones include *Wolf of Wall Street* (based on the true story of Jordan Belfort's pump and dump schemes), *The Big Short*, *Wall Street* ("greed is good"), and so on. There are no shortage of statements testifying to the abominable culture of finance pre-2008 financial crisis: a culture of misogyny, and other vices. In some ways we might think of finance corporate culture as Virtue-Ethics-but-Evil: it seems as though young interns and financiers, faced with 'exemplars' of vicious behavior, attuned themselves towards behaving like them. This can in general be said of corporate culture generally: it is quite simple in nature and tends to take the form of young interns and associates modeling their behavior off of successful, more senior employees. Although it is simple, corporate culture is very difficult to reorient or change exogenously.

These criticisms ring true and fixing an industry's culture is never simple. Efforts to change recruiting practices are a promising option. These criticisms are helpful and needed, but they do not address norms around how financial markets ought to work, just how the individuals in them ought to behave.

Financial Markets Yield Inequality

These criticisms are outwardly oriented and consequentialist in nature. Generally speaking, they have the form of looking at the reality that can come about, pointing to distasteful features of that reality, and concluding that there are some number of things wrong with finance (because otherwise, we would have a less unequal, or unfair, etc. result of such a system). The mortgage crisis of 2008 proved disastrous to the global economic system, and so something about this system is flawed. Some argue that from the results of this crisis (and others), we may conclude that there is too much speculation and risk in the market.[161] I estimate that this is likely true but lacking in these analyses is an appreciation of what we *get* as a result of taking on these risks. Another example of this partial analysis occurs when politicians, economists, and philosophers alike propose fiscal austerity on the premise that deficits ought to be close to zero and so are too high. What this approach often forgets is that, like any outstanding loan, every dollar of debt includes an equal (at time of issuance) and opposite asset.[162] Were I a homeowner, it would be ridiculous to think of myself as 'only' being thousands of dollars in debt. I would be, in fact, in debt to the tune of my total mortgage amount. But I also have a house! To focus on the deficit of a country without considering the infrastructure (amongst other things) it funds, is like focusing on your indebtedness to a bank that issued you a mortgage without remembering you also have a

[161] Stiglitz, *America, Free Markets, and the Sinking of the World.*
[162] This is explained in detail in James and Hockett, 2020.

house.[163] Still, to speak only of the risks that financial markets take is to forget the real growth they enable. I do not mean to disagree with the above literature or suggest that it is wrong or inaccurate in any way. Rather, I mean to suggest that it is considering a puzzle with a few pieces missing.

These criticisms largely focus on distribution, although a few look at counterfactual claims about what could have been funded or innovated had the system been set up differently.[164] The growing divide in wealth held by the *haves* and the *have-nots* is the most popular (and most difficult to explain-away) problem pointed to. This broad critique applies equally as a problem for markets for products and services and finance. A market for products and services-specific subset of this problem is the interrelation between the rate of inflation and the cost of living compared to the increase in wages which have not gone up as much. A finance-specific subset of this problem includes the way in which having access to wealth facilitates one's ability to access even more wealth (via advisors and premium trading opportunities not available to all). A certain subset of wealthy families have so much wealth, that their family business becomes the management of that wealth (in what is called a "family fund" and is essentially a hedge fund run with the purpose of preserving and growing that single family's wealth). Having capital also enables one to become an 'angel investor.' There is no guarantee that someone with wealth will not squander it away in poor financial investments, but those with wealth have access to opportunities that promise the same or increasing rates of return as one's wealth increases.

---

[163] To be clear, this does not absolve anyone from ever overextending themselves on their mortgage. Just because a mortgage corresponds to a house, does not mean that, in this alternative world where I can be a homeowner, I could own *any* home. In any world, there will exist an upper limit to the dollar amount it would be reasonable for a bank to loan me in the form of a mortgage. Similarly, considering the growth that is enabled by risk-taking, is not to absolve us of our responsibility to deliberate about how much risk-growth we should take on in the first place.

[164] Foroohar comes close to making this claim explicitly, but it can be inferred from her complaints about how little innovation finance has ended up doing that with a different set of priorities, more money would have been spent on research and development at some of our largest publicly traded firms

A subset of these criticisms see political will as a prerequisite to change anything in our current system, which needs additional regulatory reform. There are a variety of reforms suggested that in combination are assumed to address the problems put forth. The SEC, holding the requisite power to enact such reform, is presented as under-resources and/or under-motivated to pursue such deep reforms. Accordingly, it is asserted that the public's attention on the matter (or lack thereof) is, in part, enabling this oversight (or lack thereof). Often, these criticisms are rooted in a consequentialist logic, and suppose, on a consequentialist basis, that our system for distributing wealth is bad and that we ought to distribute it more equally. The proposed manner that we ought to do these things include raising taxes and/or giving more people access to financial markets. I will refer to these as "tax-based redistribution" and "financial market democratization" respectfully. Although I have more hope for the former than the latter, neither of these solutions seek to restructure or change anything about how the financial market operates.[165]

Financial Markets are Vicious

This kind of criticism goes further than criticism 5. Articulating that there are certain toxic features of a culture is distinct from disagreeing with an entire practice on the grounds that the industry or institution is corrupt at its core. In other words, proponents of criticism 5 could hope for a different culture to exist in finance (one that prioritizes community impact or environmental sustainability, for example), so that such an industry would be suitable for the distribution of capital to enterprises in need of financing. Separate from these folks are those who argue that capitalism, and thus the structure of our financial system, is, overall, vicious. That is, a system which allocates capital and distributes it according to the principles of capitalist economics is

[165] In other work, I argue that financial market democratization actually does change financial markets, for the worse. But that is not what proponents of this argue.

still bad, even if all individual conduct is above board and in good faith. Counterfactually, there is no possible way for a culture to be sufficiently oriented towards any particular virtues that could make up for the structural inequities that may result as the system allocates reward and risk. This kind of criticism can focus on what it does to either individuals or society as a whole. In the individual case, the idea is that being into (and good at) finance leads you to value aspects of finance that should only be instrumentally valuable as having intrinsic value.[166] These are a bit stronger claims than those in the criticism about bad behavior. The concern is deeper in the sense that critics are concerned that engaging in finance is a formative experience that changes one's personal character.

On the institutional level, these criticisms echo writers like Rana Foroohar who argue that the financial sector has become oversized and extractive, prioritizing short-term shareholder value and financial engineering over long-term productive investment in the real economy. Polyani would also fit into this category in his description of an intrinsic tension between markets and society: in his view we are always in a state of reacting to the overreach of marketization and financialization.[167]

Analysis of Critiques

All of the foregoing criticisms are warranted, and there are almost certainly others of the same ilk. These are the criticisms that tend to dominate the attention of academics, however, and they all overlook the qualities of financial markets that make them unique when compared to markets for products and services. In many instances where philosophers criticize financial

[166] Again, this is a distinct type of criticism that the kind that references specific known vices like greed, etc., which I categorize as these

[167] We could classify stakeholder theory as a response to this. See Freeman, R. Edward. "A stakeholder theory of the modern corporation." *The corporation and its stakeholders* (1998): 125-138.

markets in one of the foregoing ways, the technical knowledge required to offer productive alternatives or reforms is often lacking, and so many of the solutions proposed are ill-informed, doing little to address potential changes to financial mechanisms that would prove useful.

As a practice, philosophers who criticize financial markets tend to deploy an ethical theory of their choosing to the problematic behavior of individuals or institutions without considering unique features of the market and its environment which gave rise to the ethical impasse. Professionals may then easily discount their critics as ill-informed or foolhardy.

It is often impossible for critics to offer insight into how to improve or reform a system which has permitted or encouraged ethical failure when the intricacies of how financial markets work are not readily made available for people to learn. There are unique aspects of finance markets– their capacity to be used as engines for growth and mitigation of risk– that differentiate it from markets for goods and services in important ways. My deepest concern here is that even if we address the criticisms above, there will still be large problems that went overlooked. If we implement both tax-based redistribution and financial market democratization to assuage the above criticisms, we will still be left with a system which lacks focus on its true purpose.[168]

To move forward with a productive theory of finance from which we may draw productive reforms and criticisms, we ought to both take note of the unique attributes or norms associated with finance and posit a philosophical theory we might use to guide us. The next chapter addresses normative responses and argues against deontological or utilitarian notions of financial markets as unhelpful for addressing some of the most fundamental concerns. I will offer a teleological conception of markets which posits the true purpose of this social institution is to enable future stability by way of distributing capital to things which might yield some

[168] This is not to say that I agree with these proposed solutions, even in the limited scope of the criticisms they are offered as solutions to.

improvement of social products and services, and as a means of (re)distributing risk, with the ultimate aim of being able to tolerate more risk than we would otherwise. This is a project for which we ought to have collective deliberation.

Other Philosophical Literature Which Talks about 'Risk'

A core distinction made by economist Frank Knight was between risk and uncertainty, where risk describes situations in which outcomes can be assigned measurable probabilities and therefore priced or insured against. Uncertainty concerns circumstances in which future states cannot be reliably quantified. Risk is typically understood as either a problem of harm avoidance or as a problem of rational decision-making under incomplete information. One finds discussions of both within epistemology, decision theory, and climate ethics. Once again, I wish not to critique these works on their own grounds but rather show the bas-relief image of what is not discussed by this literature that my project covers. What remains underexplored is the first a social need for a mechanism to redistribute and organize risks involved in productive enterprise, and finance's ability to suit that need.

Climate ethics and some parts of applied technology ethics talk specifically about risks to society and how those risks should be managed. These approaches are largely non-market based, and suggest various policies of taxing or outlawing certain human behaviors that contribute to climate change, and policies inducing or financially supporting human behaviors that alleviate it.[169] Another strand of this literature concerns how decision making should be done given

[169]Two examples of a justice-based approach to climate questions that focus on individual responsibility include Bell, Derek R., 2008, "Carbon Justice? The Case Against a Universal Right to Equal Carbon Emissions", in Seeking Environmental Justice, Sarah Wilks (ed.), Amsterdam and Hiller, Avram. "Climate change and individual responsibility." The Monist 94, no. 3 (2011): 349-368.

uncertainty about what consequences human behavior will bring about.[170] Climate ethics in particular has begun to address intergenerational concerns that pertain to distributive justice.[171] Climate finance, in particular, attempts to deal with risk and uncertainty issues. In some sense, climate change is itself evidence that we have not been thinking about and distributing risk in the appropriate way.

Within epistemology, risk is most commonly understood as the risk of false belief. Debates about epistemic luck, safety conditions of knowledge, and justified belief discuss when agents are insufficiently protected against error. Duncan Pritchard challenges probabilistic understandings of risk using a 'modal' account. Modal risk depends not only on statistical likelihood but also on how easily an undesirable outcome could have occurred across nearby possible worlds. His account captures intuitions about near misses that probabilistic reasoning does not. For example, a driver who narrowly avoids a fatal accident (let's say, by turning left instead of right) may reasonably feel like their life was at risk even if, retrospectively, the harmful outcome did not occur (because they were only one small decision away from it having occurred). Modal reasoning therefore gives us a contrast to purely probabilistic accounts of risk assessment. Still, the epistemology literature on risk is focused on individual cognition and reliability and does not extend to systems-level analysis.[172]

The moral risk literature assumes that risk is intrinsically negative. Much of the literature emerges from paradigmatic cases involving unjustified harm to others such as cases of reckless

[170] For example, there is uncertainty about what the long-term environmental consequences will be from various increases in global temperature.
[171] See Moellendorf, Darrel. "Justice and the assignment of the intergenerational costs of climate change." *Journal of Social Philosophy* 40, no. 2 (2009). And Fritsch, Matthias. "Climate Ethics and Intergenerational Reciprocity in Indigenous Philosophies." *Intercultural Philosophy and Environmental Justice between Generations: Indigenous, African, Asian, and Western Perspectives*. p. 33–58. Cambridge University Press, 2024.
[172] Nor does it content with what I take to be the genuinely epistemologically interesting phenomena in which collective expectations themselves reshape future outcomes (which is an interesting feature of financial markets).

endangerment or interpersonal violence. This framing makes risk seems like that which ought to be wholly avoided, since there is no endorsable 'gain' to be had for the associated risk of, say, playing a game of Russian roulette where the 'gain' is a perverse enjoyment from the inflicting of fear on part of the would-be shooter, at the expense of the would-be shooting victim.[173] Colloquial language rarely caches out possible positive outcomes in the language of risk because every day moral discourse typically concerns unwanted harms. But both losses and gains are imbued with uncertainty and risk. If x-risk is only associated with y-harms, or y-perverse-gains, we are systematically obscuring the kind of risk versus reward framework dominant in financial analysis. A more robust way of speaking about risk is to apply that kind of language to all kinds of uncertain outcomes, encompassing both upside and downside risk so that we are better able to take into account all-things-considered judgements about actions which may cause chains of events that could potentially make us better-off.

Knight's distinction between risk and uncertainty concerns measurability: how we can account for our inability to discern precisely what effects will be caused by some action or outcome. Pritchard's epistemological modal account gives us a way to think about proximity to close worlds, which is of value. Neither gives us normative tools though, which we will need if we are to better understand how much risk societies ought to take, and how that risk should be distributed.

[173] Maheshwari, Kritika. "A unificationist approach to wrongful pure risking." *Inquiry* (2024): 1-28.

Chapter IV: Positive Account of the Telos of Financial Markets

As we saw in the last chapter, academics have plenty to say about financial markets. Economists are concerned with understanding pricing, market efficiency, etc. by means of formal models. Ethicists concern themselves with processes and outcomes, just or unjust. The wealth and income inequality rife in our current system, for example, leads to the moral evaluation of financial markets. And there are others, previously unmentioned. Sociologists concerned with financial markets, for example, focus on the social and institutional forces that shape market behavior, power relations, and the cultural meanings attached to economic actions. Regulators too look at issues like stability and risk. But all of these various critical approaches are at risk of being summarily dismissed by market participants. One problem is that these views often talk past each other. The bigger problem, however, is that financial practitioners tend to disregard 'big picture' issues like income and wealth inequality or market stability because they are just 'doing their jobs'. There are many reasons (epistemic, ideological, practical, etc.) for seemingly 'reasonable' individuals in financial markets do not feel responsible for many of the issues academic critics bring up. Some of this can be chalked up to standard issues in communication. But a root issue, I argue, is that no one is thinking about financial markets in the right way. Most people (e.g., participants in financial markets as well as those who study them) have views about what financial behavior and institution are *for*. As Chapter II makes clear, financial markets are shot through with *this-is-for-that* reasoning– assets, actions, and institutions are designed to achieve certain ends. But people have different views about how to characterize those ends and whether the various options on the table are normatively appropriate or constitutive of all there is to care about. These differences make it easy for people to talk past one another in crucial conversations.

This chapter focuses on defining an *appropriate* teleological account of financial markets, so that we develop a purpose-driven criticism of current financial markets. This project is both apt[174] and useful.[175] Here I argue for the correctness of the teleological approach. First, I think that social institutions can be best understood as having a point or purpose, i.e., a *telos*. Further, I hold that in the case of finance, the *telos* is to determine future markets for products and services by promoting stability and distributing risks, enabling those markets to provide for the benefit of all.[176] That's the point of this chapter.

None of this is to say that current financial markets actually fulfil their *telos*. In fact, they pretty clearly do not. One risk of not properly understanding this *telos* is evidenced by the current state of financial markets. That we are not currently[177] in the midst of a debilitating recession is more happenstance rather than the outcome of rational decision making and regulation. A financial system can be set up to make us more vulnerable when risks happen to manifest, or better at tolerating risks. There is considerable room for improvement.

What we want philosophy of finance to do

Good philosophy of finance should be able to distinguish 'good finance' from 'bad finance,' in terms of functional (as opposed to dysfunctional) systems, not just in terms of outcomes or rule compliance. I want to differentiate instances where a financial market is working well from those in which there is merely no fraud or bad behavior. Dysfunction itself

[174] If teleological approaches - virtue theory, eudaimonism, essentialism, etc. - are the right way to do ethics then we need to think about the telos of financial markets if discovering a good-making feature of financial markets helps us get better lives (*a la'* the school example from Chapter I) then that discovery provides support for a teleological approach to ethics.
[175] We can grant for the pragmatists the weaker notion that it would behoove us to act as if institutions have a specific purpose, so as to best calibrate their functioning pursuant to certain ends.
[176] This latter point usually involves economic growth.
[177] As of this writing in March of 2026.

has moral upshot in some systems and institutions. When, for example, an educational system is dysfunctional, it is generating (or at risk of generating) moral harms (e.g., ignorance) even if everyone involved behaves themselves by appropriate standards. This makes an educational system different from, say, a block party.[178] In similar fashion, a restaurant can be a poor one even if the owners are upstanding people, the waitstaff is friendly, and the kitchen staff is hygienic and motivated to please. Why? The location might be terrible, the recipes might suck, or the cooks, waiters, or owners might just be incompetent. Likewise, in finance. A bank or brokerage firm can be staffed by saints, but they at least need to be competent for the institution to work well. Rule-based ethical frameworks will not show us why certain financial practices are still wrong in some way. If there is not some institutional-level normative guidance, individuals in finance can avoid individual-wrongdoing without supporting the right-doing of the institution. A properly oriented teleological approach should be able to explain why dysfunctional finance matters morally even when no individual rule violations occur.

The first bridge to cross is whether we ought to think about financial markets as having a *telos* at all. Finance shapes what markets for products and services will be like in the future. It therefore has important social power. Shaping the future poorly would be a moral failure with significant consequences. Financial markets, then, need a social justification for their existence that goes beyond "this is what has emerged so we can't alter it now." There is no reason not to tinker with institutions without some sense that we have got it right. But to think about markets as having a *telos* means that we can do a meta systems-level analysis too, which can yield system-level solutions.

[178] There may be a matter of degree here. I suppose a block party could put a neighborhood kid on an overly hedonistic path even if everyone at the event acted within the bounds of basic decency.

Non-Teleological Approaches

I start by trying to demonstrate the inferiority of non-teleological approaches to finance. I broadly group alternatives to the teleological view as being either consequentialist or deontological. I argue that these approaches miss something important about what social institutions should be doing for society, and what we should do to reform or orient institutions towards what we, as a society, think they should be doing for us in a systematic way. In what follows, I illustrate what a few types of consequentialist- and libertarian-style approaches would beget. I think the approaches of each type that I describe are the strongest of each type.[179]In each case, I look at both approaches that would focus on the individual (e.g., act utilitarianism and Kantian deontology) and those what would address the institution as the locus of the theory (e.g. rule utilitarianism and libertarianism). I do this to show that none of these approaches will give us *enough* tools to say whether financial markets are functioning well or badly, or, more importantly, how to address issues that arise from the material development of human society. All of this, of course, is meant to motivate the need for a different kind of ethical reasoning (a teleological view!)

Utilitarian Approaches

Many common analyses of financial markets implicitly use a consequentialist lens. This, anecdotally, seems to be the most popular approach that both scholars and everyday people use to gauge the financial system. The thought process here, for example, might look at all the

[179] But I could be wrong in two ways. First, I could get wrong what a libertarian argument (say) about financial markets would say. Second, I could fail to address an approach that will prove to be better than the neo-Aristotelian one I describe below. I don't think this will happen, but then again, I am, after writing this section, perhaps one of the least likely to see where I have gone wrong (if that is the case). I have tried to identify not just a set of straw man approaches, but what I think are the best approaches that consequentialists and deontologist could come up with respectively.

wonderful things modern society has, therefore concluding financial markets are good. Or, in the negative (but using the same methodological lens), people look at all the inequality that exists and conclude that financial markets are bad. Few thinkers address financial markets in a teleological way, where the purpose might acknowledge a productivity/equality tradeoff. The welfarist and capabilities approaches to political theory questions get close, but economically they only deeply touch issues related to individual consumer finance. This kind of critic often highlights individual-level problems with individual-level solutions: this financier was bad, do not be like that; this practice is harmful, let us regulate it. What I take to be a generally utilitarian approach to financial markets is one which uses their outcome or consequences to evaluate them, and so to shape their structure by way of regulation, design of incentives, and individual action-guidance.

Problem 1: What counts as a good effect?

A utilitarian conception of markets– for products and services *and* for finance– would place at the forefront welfare arguments, issues of distribution, and risk mitigation. This focus would allow consequentialists to argue that markets (financial and otherwise) are good because they make us 'better off' in some way, e.g., more desirable products and services. The negative consequences (e.g. however the market works, it is unevenly distributing the gains from it) are externalities that are (or are not) outweighed by the more robust prosperity they provide. A utilitarian theory, of course, will suggest we do that which will bring about the best consequences.

This is a reasonable way to understand what might be going wrong with our current financial system but does not help as much to guide market reforms or solutions, just amelioration of the market's effects. So, consequentialism does give us a plausible diagnostic

framework to use at a point of time: is the sum total of the effects of our current financial system on the whole, one in which the positive effects outweigh the negative ones. My concern here is twofold: First, that consequentialist theories nudge us towards downstream intervention and the bracketing of upstream structural causes (by treating complex systems and problems as 'black boxes' too opaque to see into or change). My proposal is meant to 'open the black box' to the extent that it tries to keep the inner workings of financial markets focused on their role in supporting human flourishing. Secondly, even if consequentialist theories can avoid treating financial systems as 'black boxes', when they are applied to markets, there is a serious calculation burden involved in the prospect of understanding which actions bring about the best effects in a complex system like financial markets. Even if bringing about the best consequences overall is a criterion for morality and not a decision rule, we need some serious insights into how market forces work in order to form judgments about when things are going well and what to actually do when we do not think so.

In the case of unequal wealth distribution, a utilitarian framework might rightly conclude that the consequences of our current financial system are unjust, and that a concentration of vast wealth in the hands of a few is not the best distribution of resources to maximize the well-being of our society. But this diagnosis often treats the underlying structure of financial markets as fixed, and is a 'black box' account which presumes the machinations of the financial system which gave rise to problems aren't *fixable* in any relevant sense.[180] One utilitarian solution, then, is to shift resources to foot the bill for a fix instead of trying to understand how the bill got so high in the first place, and to focus on the last 'domino' in the causal chain of events leading to

[180] There are black box accounts for other social ills too, such as affordable housing, healthcare costs, and college costs. These theories have the common feature of taking for granted how relevant social institutions and practices operate within society.

the poor consequences. One reason this public policy approach is so pervasive is that it is simpler than addressing the complicated realm in which the underlying problem occurs.

We see this dynamic in other arenas of life too, like education and housing. In higher education, rising tuition is often addressed by subsidizing higher rates of scholarship and loans, instead of addressing or inquiring as to why the price of tuition is rising so precipitously in the first place.[181] With regard to housing affordability (or lack thereof), many promote solutions like rent subsidies (e.g., money given to tenants by the state to help them pay for rent). This 'solution' enables residents to pay existing market prices, rather than solving any underlying issues which gave rise to unaffordable housing in the first place.[182] In both cases, the proposed solution is operating externally to the mechanism producing the problem, taking anything prior to the emergence of the problem as unchanging.[183]

---

[181] This does ignore the way in which tuition pricing has started to function in some ways like a dynamic pricing model, where the sticker price of tuition is only paid by a fraction of students, many others of which receive a large discount rate depending on their family's financial circumstances.

[182] In the case of housing, most places where affordability is an issue there is also a supply issue (e.g. the building of new housing hasn't kept up with the increase in population). So, in most cases the unaffordability of housing is directly related to a lack of supply. Building more housing is costly and is usually more profitable to be built on a single-family and/or luxury basis than the provision of affordable rental units directly. In this manner, even when new housing is built, there must be an indirect (and medium-term) shuffling of households out of housing which is subsequently reduced in price as more expensive supply is snatched up by those of means which suffice them to move into newer rentals or homeownership. I have suspicions about this trickle-down state of nature-style story about housing but shall not get into criticisms of it here. The point to make is simply that giving money to individuals to pay rents that are too high for them to afford does nothing to address the underlying environment which has led to their rent being unaffordable, and any solution which does address the root cause is more complicated either in terms of it coming to fruition or the timeliness of its solving any underlying supply-related issues.

[183] A consequentialist way of seeing things steers us toward 'black box' theories of problems, which take for granted some system, institution, or dynamic which has certain effects. Still, it is plausible that in many cases, we can indeed deduce reliable relationships between certain causes and certain effects (e.g., that giving a mortgage to a prospective homeowner with a credit score under a certain threshold is overwhelmingly likely to default on the mortgage). The consequentialist might argue that although financial markets are a domain in which many calculations of cause and effect must be calculated, with the power of, say, large-language models, we can actually reliably calculate, for example, what traits are held by prospective homeowners will make them reliably pay back their mortgage. Indeed, this sort of thing already happens. And this sort of thing can be done reliably for *bad* ends, as noted by those working on the negative possibilities made possible by the use of AI in finance.

Problem 2: Temporal Issues in Calculating Effects

Even if we grant that it was possible to determine which available actions in the present time period would produce the best consequences, this would only work within that time period. Financial systems are diachronic in that they shape the markets for goods and services that will exist in the next time period, and there are innovations specific to finance, where new financial products are created and then spread and are adopted. The relevant set of choices which correspond to causes and effects is not fixed over time.

So, the agents (or regulators or other institutional managers) are not choosing from a stable set of options. In the multiple-choice exam of decision making, even if one could confidently determine that out of choice set A, B, C, and D, B is the option that yields the best outcome, that becomes stale reasoning once new options (let's say, E, F and G) are introduced to the choice set. In that time period, the existence of E, F, and G 1) necessitate a recalculation of consequences and 2) might erode or change the consequences from choice B (what was, in the previous time period, deduced to be the best option for the sake of consequences). In other words, the domain of possible actions itself is changing in unpredictable ways between time periods. Thus, the consequentialist's aspiration to select the action which yields the best consequences is not one we can make sufficiently deterministic institutional rules about.

A teleological account also cannot predict the future with regard to new technological developments. But it need not rely on the continuous identification of optimal action or rule. On my view, what matters is that the set of available options is itself structured by reference to the *telos* of financial markets. Wanting an agent to always select the best option is unstable as the option set evolves but constraining how the option set evolves in the first place (by attending to the purpose of an institution) is more plausible. The idea here is to effectively constrain the

option set (including new developments) in such a way that all choices have a meaningful connection to the proper functioning of financial markets. Unintended negative consequences can still possibly emerge, but there is a smaller universe of permissible new developments with connection to the purpose than there are under a consequentialist theory. A teleological framework will not entirely eliminate the risk of bad outcomes, but it seems equipped to narrow the possible set of innovations to those that are intelligibly connected to the market's purpose.

Problem 3: Systemic Risks are Left Unaccounted For

Consequentialist approaches struggle to give us a satisfying way of thinking about systemic risk. Systemic risk cases are those where the relevant normative concern is not reducible to the outcomes of discrete observable actions (like, say, making a loan to an individual or not), but instead concern the possibilities related to large-scale and low-probability harms that may never materialize. These are kinds of risks we have strong reasons to prevent. A framework that evaluates actions primarily in terms of their realized consequences or past patterns of cause and effect is one that is poorly equipped to register the normative importance of preventing crises that, if adequately protected against, will not happen.

From the perspective of individual transactions, there are plenty of financial activities that seem permissible or to support human happiness that are only harmful when we think through aggregate and systemic possible harms. If we aggregate some behavior, those activities collectively can generate correlated exposures that only become visible at the level of the system. Consequentialist analysis risks missing the forest for the trees because it can vindicate locally rational or welfare-improving behavior while failing to capture how such behavior contributes to overall states which may be undesirable and put society at risk of harm.

One might hope that rule consequentialism could address this by endorsing rules designed to mitigate systemic risk, but the resources here are also limited. Rules that are implementable are rules that tend to do little to address the structural and emergent features that generate systemic vulnerability. Giving individuals rules to follow will not address these issues that do not emerge at the individual level. Telling financial market participants to act fairly or be honest will not, for instance, prevent the buildup of correlated risk in highly leveraged markets or speculative bubbles.

On the institutional level, the rules that will bring about the best consequences will be those that select for the least amount of risk (e.g., the most conservative option). This approach, however, would result in a regime with extreme risk aversion across the system. This would undermine the opportunities that taking risk allows the possibility for. Any genuine risk would seem not worth taking, and it is difficult to see how the innovation and capital allocation that financial markets enable could survive.

The upshot here is that consequentialist views are categorically likely to generate guidance that is either too narrow, too limiting for future persons, or too indeterminate to address the kinds of large-scale risks that are of major concern for financial markets. They will successfully rule out clearly harmful individual and institutional behaviors but will not capture all relevant issues of moral concern.

Deontological Approaches

Basing financial markets on rights and duties would emphasize respect for private property and adherence to contractual obligations. These issues are part of contemporary views of finance, but not necessarily the priority such considerations have over outcomes. A simple Kantian view (based on something akin to a categorical imperative) applies a deontological– and so non-

consequentialist– framework to individuals in finance.[184] A libertarian viewpoint would emphasize the threat government or regulatory intervention poses to personal rights and property. These philosophies see the *moral* value of the financial system in the voluntary (and so rights-respecting) transactions they support. A morally permissible financial market that respects the rights of others allows individuals to engage in markets simply to satisfy their preferences. People can act prudentially, with no mind toward any ultimate moral goal of such activity.

A deontological conception of markets will miss some of the key intuitions about what financial markets should do. This kind of view entails that a transaction that follows the rules is not only necessary for morality, but is also sufficient. We can only endorse or reject individual contracts or engagements with the market based on their particulars. This leaves macro-effects and institutional issues to the wind. This view suggests that there is nothing wrong with (say) the meme-stock phenomenon, so long as everyone is obeying the relevant regulatory frameworks. But there is something intuitively disturbing about our financial system making space for nonsensical trading behaviors.[185] But a teleological view (even the teleological views I think are wrong) can make sense of why meme stocks are 'bad.'

Deontological views do not help us think through issues of systematic risk. A libertarian view, for example, is basically blind to outcomes and cannot endorse governmental limits on transactions between consenting parties This commitment serves to amplify the negative effects

[184] This is what I take John R. Boatright's project to be in *Ethics in Finance.*

[185] This too, ironically, seems to be the operating intuitions of those who want to reject a certain flavor of ESG (a term used to describe concern for environmental, sustainability, and governance factors in investing decisions) and corporate responsibility trends. There are two importantly different versions of ESG. Originally, it was conceived of as a set of factors which would have effects on traditional profit maximization. If an oil company has no plans made for the onset of climate change in the medium term, such an oversight is a serious risk factor to that company's success. A different approach says we should be focused on things other than profit maximization (e.g. the environment). The latter camp is fine with sacrificing some profit for the sake of some other value; the former makes no claim about value. It simply notes that the environment is a factor in expected profits, if that's what you want.

of inevitable downturns in the business cycle[186] in ways that will be harmful to society as a whole (not to mention the worsening of individual-level impacts). There is no reason to carefully investigate the general functioning of financial markets if all that matters is that each particular transaction follows the moral rules. In fact, these kinds of views plausibly incentivize the obfuscation of risk, by means of lessening regulatory compliance burdens without substituting any other safeguards.[187]

I am open to accepting any view from either of these theoretical traditions which solves the respective issues I have highlighted here. As they stand now, the best of each kind of view highlighted here will still necessarily miss key features we should be looking for, such as a way of conceptualizing risk.

Neo-Aristotelian Views

Current work that looks at financial markets through a teleological (Aristotelian or Virtue-Theoretic) lens tends to focus on individuals and their roles ('micro' issues) rather than the point of financial markets themselves (a 'macro' account). These applications are often susceptible to one or more of the criticisms outlined above. These views usually presuppose that to redeem finance as an industry (the public perception of finance) that it suffices to reform individuals who work in finance.[188] The view endorsed by de Swaan in *Seeking Virtue in Finance*, for example, is that there are exemplars in finance we ought to aspire to be like (even though he seems to have the view that finance in pure self-interest isn't wrong *per se*).

[186] To be clear, we should not take for granted that this is what we'd want to do.
[187] The key here is that although we might legitimately protest the cumulative burden of regulatory compliance on corporations and businesses, we ought to conclude not that we should simply do away with such regulation but that we ought to find a less resource intensive compliance regime.
[188] I think it is important to distinguish between the ways in which we might reform finance itself versus the reformation or transformation of the *perception* that the public has of finance. Not everyone tracks this distinction.

There are a few authors who argue (some indirectly) that the *telos* of financial markets is to efficiently allocate capital to potentially socially productive uses.[189] This is, in my view, an insufficient definition because it obscures specifically the functional benefit of distributing risks and inclines us only to think of the entrepreneurial opportunities. Still, this seems to be a rather agreeable definition: *let us just use finance for good, socially helpful stuff, and not for any bad or unhelpful stuff!* Roncella and Ferrero's apply Aristotle's distinctions among *oikonomikos*, natural *chrematistiks,* and unnatural *chrematistiks* as a way of delineating these various market activities. Actions which fall under the umbrella of *oikonomikos* are those directly related to household management, such as the actual farming of the crops, or sewing of wheat. Actions that fall under *chrematistiks* are those which pertain to moneymaking, but there is a significant difference between what is considered natural and unnatural. Natural *chrematistiks* is moneymaking for the sake of the household (e.g., the farmer who wishes to sell some of their crop in order to pay for other household expenses) and is part of *oikonomikos*.[190] Unnatural *chrematistiks* are those moneymaking actions which are done for no other purpose than to accumulate wealth.

These sorts of thinkers are specifically oriented toward assessing innovations in finance, and they take specific aim at high-frequency trading as a *purely-for-profit* activity. High-frequency trading does seem like a waste of resources, in that it enables prices to change faster than humans can perceive that change. On the other hand, the 'efficiency' of price discovery, or the degree to which prices reflect what the aggregation of market participants thinks the value is, is an important value. So, in theory, the faster prices change, the better (i.e. more accurate) price discovery must be. Still, there is an upper limit to this helpfulness: if prices change faster than

[189] Fama, Eugene F. "Efficient capital markets: A review of theory and empirical work." *The Journal of Finance* 25, no. 2 (1970): 383-417.

[190] Bueno, José Luis Cendejas. "Economics, *chrematistiks*, oikos and polis in Aristotle and St. Thomas Aquinas." *Journal of Philosophical Economics* 10, no. Articles (2017)

human perception, it is not changing as a result of what humans think the thing is worth.[191] Roncella and Ferrero's framework is helpful here: it is clear that the speed of price changes, past a certain threshold (human perception), is unconnected to any service to the real economy. It is, explicitly, unnatural *chrematistiks*.

Roncella and Ferrero are not alone. There is a long tradition of criticism of speculation on the ground that speculators and short sellers are parasitic on the real economy, merely betting on prices and outcomes than contributing to productive economic activity.[192] But it is perhaps only in cases like high-frequency trading that we can so easily pull apart the 'natural' *chrematistiks* behaviors that facilitate social productive ends from 'unnatural' *chrematistiks* actions. If we consider an example using the good-natured risk adverse farmer from Chapter I, this delineation becomes more difficult. We can think through a standard two-sided trade. In such a trade, one counterparty may be participating for reasons directly tied to the economy (this will be our farmer). The other counterparty is perhaps only participating purely for profit aims (this, we can imagine, is our financier). The farmer wants to offload some of their risk, and they are happy to forego the theoretical maximum profit they might make in a good crop year. The financier is willing to take on that farmer's risk, in the hope that it is a good crop year, and they get to keep that profit. The financier is acting, we can assume, for the sake of pure profit. The farmer's ability to manage the risk of their needed economic activity (that which would certainly qualify as 'natural' or part of the basic household management component of Aristotle's view) is only made possible by the counterparty's willingness to assume that risk. In modern times, it is

[191] It is changing based on what algorithms can predict other algorithms will do, those latter algorithms having been based on predictions of human behavior

[192] Banner, Stuart. *Speculation: a history of the fine line between gambling and investing*. Oxford University Press, 2017.

unlikely that this counterparty is motivated by anything other than the prospect of gain. The same transaction, in this case, instantiates 'natural' *chrematistiks* for one counterparty and 'unnatural' *chrematistiks* for the other.

Most importantly, the farmer is unlikely to find someone willing to take on their risk for anything other than the prospect of possible monetary gain. Because risk-bearing is typically undertaken for the sake of potential profit, our condemnation of wealth-seeking activity for its own sake would also undermine the mechanism by which an enterprise clearly to the benefit of society (growing food) can manage irreducible risks to its activity. In rejecting unnatural *chrematistiks*, we are also rejecting anyone on the other side of trades those seeking wealth for its own sake would make.

The ultimate issue with this teleological approach is that it does not say what speculation we should keep, and what we should prevent. The idea that finance should only do the good things and avoid doing the bad things implies both that we can sort good from bad, and that the good could go on without the bad. In reality, these are both flawed assumptions– the former for reasons of unintended consequences, and the latter for more basic mechanistic reasons.

Other teleological views

In this part of the dissertation, I work through several possible conceptions of a financial market *telos*, including the most common (implicit and explicit) versions. I demonstrate that other teleological views are some combination of flawed, misguided, or wrong. I sort these alternative views into a few categories, representing them with titles that express implicit or explicit differences in what they consider to be the purpose of financial markets. I assume that the proponents of each of the following views think we should think explicitly about the purpose

of financial markets, and that such a purpose is separable from (although perhaps ultimately the same as) the purpose of markets for goods and services.

Telos: wealth creation

A more plausible version of the consequentialist, preference-satisfaction view treats the *telos* of financial markets as wealth creation (they would presume, via the satisfaction of preferences). On this account, the good at stake is instrumental. Financial markets are valuable insofar as they generate greater prosperity.[193] If Pareto-improving exchanges are allowed to proceed without interference, the view goes, individual preferences will be satisfied, and the resultant trading activity will produce wealth. Proponents typically add that this mechanism serves broader social ends with spillover effects. The lure of wealth serves as an incentive for people to innovate and become entrepreneurs, making new products that will fulfill existing consumer preferences or generate new ones.

This account is compatible with the conception of the market for products and services advanced by economists Bruni and Sugden in their paper "Reclaiming Virtue Ethics for Economists." On their view, a market's virtues will include traits that make it better for individuals to enact mutually beneficial transactions. Financial markets will have a special, specific role to play in creating mutually beneficial exchanges. In brief, Wall Street is supposed to underwrite the possibility of beneficial exchange by directing resources in a way that facilitates more and better mutual exchanges. More fully, financial markets are supposed to allocate capital to firms that innovate in ways that make for better transactions, that is, to direct resources to firms that build more desirable (more functional, cheaper, etc.) products and

[193] Some might care about the distribution of this greater prosperity, others might not and will regard any increase in, say, GDP, as a measure of increased prosperity such that it satisfies the condition of making us collectively better off.

services. In this sense, financial markets are not valuable primarily because they generate wealth, but because they direct resources toward firms that improve the terms of exchange by producing goods and services that are better somehow.

On this picture, economic growth is not itself the *telos* of financial markets. It is a byproduct of their success in facilitating mutually beneficial exchange. Growth emerges because financial markets enable the distribution and management of risk in ways that make innovation possible.[194] The connection between growth and mutual benefit is indirect. 'Growth' tracks the expansion of opportunities for mutually advantageous market exchange, rather than existing as an intrinsic aim or good of the system.

Advocates of financial markets *also* tout them as a means of wealth creation. This is connected, of course, with the idea of mutual benefit– successful investing is often lucrative for many. Paradigmatically, companies get upfront capital for exciting new ideas that might succeed or fail, and once those ideas come to fruition (or not) the investor is rewarded with a healthy return (or loss) on their investment. Everyday people are encouraged to buy into the market– either as traders themselves or investees in various funds managed for them– as a way of growing their wealth.[195] Arbitrageurs are incentivized to locate discrepancies and profit from closing such gaps. Overall, this *investment-builds-wealth* principle has held true for over a century.[196] From 1965-2019, for example, the annualized rate of return on the S&P 500 index has held at ten percent on average.[197] During the twentieth century, U.S. equities had a real return of

[194] This specific point is my explanation of why this occurs but is not made explicit by Bruni and Sugden.
[195] Stock-market invested funds from pensions and 401(k)s are touted as a way to save for retirement and are incentivized by offering tax benefits. Still, only 32% of Americans were enrolled in 401(k) plans according to the US Census Bureau.
[196] Royal, "Average Return for the S&P 500."
[197] Buffet, "Berkshire's Performance vs. the S&P 500." (2014)

4.3 percent between 1921 and 1996.[198] An individual who invests some amount of their excess money (funds they have that they don't wish to spend yet) *could*, under the right conditions, create a great deal of wealth. But individuals can and do lose money, while others gain: Some risks are realized as losses, and unrealized risks are returns. The (mere) possibility of wealth-generation is certainly distinct from the reality of individual wealth building.

Descriptively, this wealth-creation *telos* is the most widely held teleological view about financial markets (either implicitly or explicitly).[199] Folks who work professionally in finance tend to see creation of wealth for themselves as their ultimate goal in market participation. It is worthwhile, however, to distinguish between the purposes of the individuals that work in performing stock market actions and the point of the financial market as a whole. Those employed in the lucrative financial sector might not usually seek out those roles from a sense of duty to appropriately allocate capital. Income-seeking is, perhaps implicitly, the *telos* of traders. But it does not follow that the market itself is about income-seeking or wealth creation. Appropriate market mechanisms would allow individuals to act in self-interested ways without undermining the purpose of the market itself. For example, a country's military services are *for* (let us suppose) the protection of their country by use of force. Someone can join the armed forces (voluntarily or otherwise) without being motivated by that objective. The institution is set up in a way that allows it to employ individuals with different aims in the fulfillment of its overarching purpose. Someone might volunteer for the army to receive otherwise expensive medical training, and others might be entirely indifferent to the goals of the military, joining because it is compulsory to do so. Indeed, the whole point of the 'invisible hand' view is that the

[198] Philippe, and Goetzmann, "Global stock markets in the twentieth century," *The journal of finance* 54, no. 3 (1999): 953-980.
[199] Sometimes this is cashed out as capital allocation and its necessary entailments.

market can harness self-interest for the public good.[200] Even those who are motivated by the goals of the armed forces are often also motivated by more individual concerns. None of that changes the point of the armed forces. Indeed, the whole point of the 'invisible hand' view is that the market can harness self-interest for *another* purpose– the public good.[201]

We cannot just dismiss human efforts at wealth creation as immoral, *per se*. Wealth is an important means for enhancing human well-being. It can help secure not only our own futures but also those of our neighbors and our descendants. As Aristotle observes in his discussion of *chrematistiks*, there is an important distinction between forms of money-making that are oriented toward meeting human needs and household management and those that are oriented to money-making itself. Natural *chrematistiks*, where one is engaged in acquiring resources in order to support a flourishing life is permissible and necessary. The problem arises with unnatural *chrematistiks*, where the accumulation of money becomes separated from ends related to human survival and well-being. This suggests that 'money-making' alone is too indeterminate (because it could be for either kind of end) to serve as the organizing purpose of financial markets.

My central concern is that one legitimate aim is being mistaken for the *telos* of the system as a whole. Regardless of whether we take wealth to be merely instrumental or also intrinsically valuable, the mistake here is in treating what is properly a byproduct of well-functioning markets as their defining purpose. Even on a conservative view where financial market activity is understood as generating wealth in the course of serving socially productive ends, we would still,

---

[200] Of course, there are critics of the invisible hand, like Stiglitz who argue that "the reason that the invisible hand often seems invisible is that it is often not there." (*The Roaring Nineties*, 2006)

[201] Perfectly competitive markets, for example, are often considered an ideal because of their efficiency properties. Producers on the supply side of an economy, however, do not want to be in a perfectly competitive environment. In such a circumstance, their (economic) profits would be driven to zero. Producers instead seek to differentiate themselves from competitors to gain market (and pricing) power, all to increase profits. A market approximates the perfect-competition ideal only to the extent that it frustrates these producer ambitions. The market itself, however, might need to be structured and regulated to make sure it does its job.

on my account, be confusing the reward for successfully allocating capital (on the standard view) or distributing risk (on mine) with the purpose of those activities themselves.

Wealth creation for ourselves is, therefore, a plausibly legitimate moral aim, but it is not the only one. Still there are those who see the lure of wealth as corrosive to any society-focused *telos*. Financial markets, they fear, are being corrupted by a preoccupation with wealth creation. If the financial system is only about wealth creation, there is clearly something missing. Wealth creation as a sole institutional orientation is morally insufficient precisely because financial markets influence and can interfere with so many more human ends.

Telos: Not Applicable

I do consider the possibility that financial markets have no *telos*. This view suggests that markets are not themselves the subjects of moral evaluation, but rather frameworks or background conditions within which we can ask moral questions. Some economists and philosophers, for example, hold that financial markets are emergent systems – the natural results of self-organizing human interaction – they just are what they are. This is a claim not only about how markets come about, but also about the nature and limitations of ethics. On this view, morality is about making the appropriate decision in the situations we find ourselves in and does not include evaluation of the constraints on our decisions. Many would agree that we cannot evaluate the moral status of a tornado (or the weather more generally) because none of it is under our control. Weather is merely the backdrop to the moral stage on which we act. Thinkers in this camp see financial markets as a kind of natural outcome of social organization, one that is not directly accessible to philosophical theorizing. If participation in financial exchange is taken to be an inevitable feature of human social life, and we are not attuned to ways in which we might

meaningfully modify the institution, then criticism will tend to target the conduct of agents within markets rather than the structure of the market.

I disagree with this view on two fronts. First, I doubt that financial markets are emergent in the relevant sense. Second, I do not think that something having the property of being emergent or natural prevents it from being something with a (morally relevant) purpose. On the first point, it is important to note that financial markets are paradigm examples of social institutions that social-contract theorists think we coordinate to create (in a pre-moral state of nature), so they are at least as morally salient as governing institutions. Complex social organizations do not *merely* happen the way, say, a thunderstorm simply emerges from some set of conditions that occur in the world. We generate institutions iteratively based on needs. More importantly, I think it is important to consult the sociological and historical record. As examined in Chapter III, the legend that money was invented because it was more efficient than barter is a myth that plays the same role in economic history that 'state of nature' myths were employed by political philosophers of yore. Although recorded history suggests that humans have been involved in broadly 'economic' activities for their entire history, there is no evidence that centralized exchanges to trade fractional ownership in the form of 'stock' as we know them evolved from some non-financial predecessor. What does appear in ancient records are rudimentary kinds of derivative options used as insurance. In this sense, 'financial products' were around for the expressed purpose of risk-mitigation. It is demonstrably true that futures contracts existed before decentralized exchange and that abstract financialization actually existed prior to currency. The difference between the historical record and our economic myth is that this abstract financialization was mediated between individuals in a community, not in any kind of centralized marketplace or exchange.

To make my second point– that being emergent does not disqualify something from having purpose– I use the example of language. Human language, which by all accounts just emerges when members of our species hang around together enough, also has a purpose. We might say the purpose of language includes several facets, but communication between individuals is surely included. There is disagreement on how to improve, but various language reforms (formal and informal) have been suggested to better communicate with one another.[202] In the last decade, gender activists have advocated for the use of 'they' as a singular pronoun for those who do not wish to go by 'she' or 'he.' Pronouns were under scrutiny even earlier, with suffragist movements in the late 19th century arguing against the pervasive and singular use of 'he' in legal writing.[203] The movement by the trans and queer communities to use different pronouns which better approximate their gender is not a change that occurs immediately, and there has been significant resistance to changing the pronouns one uses for others. Still, there are many folks like me who are now eager to use an individual's preferred pronouns after speaking for many years with linguistic ignorance and apathy. Another reform of this kind has occurred with the term used to describe indigenous people of North America. In the early era of colonization, the mistaken perception of Europeans that they had landed in the Indies prompted them to call the people they found here 'Indians.' Despite their intention to set sail for the Indies, they had landed on the opposite side of the world.[204] Upon learning of their mistake, they were collectively unmoved to change their positionally incorrect term. More than four centuries later, indigenous activists in the United States led a vastly successful effort to forgo use of 'Indian' in favor of accurate terms like 'Native American,' 'Indigenous American,' 'American Indian,' or

[202] An excellent work focused on some of these suggestions and implementation throughout recent western history is detailed in McConnell-Ginet's *Words Matter*, which I thank Wayne Riggs for bringing to my attention.
[203] Baron, 2020.
[204] Columbus, 1924.

simply 'Indigenous.[205] Although we still have progress to make in both cases, the fact that positive change has occurred is evidence that the emergent nature of something does not preclude it from critique and reform.

Another version of the no-*telos* view is the one championed by David Gauthier in his book *Morals by Agreement*. For Gauthier, morality is a matter of rationally limiting your own pursuit of your goals. In his view, under the conditions of a perfectly competitive market there is no need for such goal-throttling. Where markets work appropriately, morality is just not an issue.[206] I disagree with Gauthier on two fronts. First, I would point out that the conditions for a perfectly competitive market are rarely found in the actual world, making for few occasions in which Gauthier's conditional will hold. In a perfectly competitive market, people's preferences aren't entangled: each individual cares only about their own desires *and* those desires don't include what happens to others, just one's own consumption of goods and services.[207] In Kantian terms, a perfectly competitive market describes individuals as treating others merely as means to their own ends, and not as ends in themselves. But this is in contrast with behavioral psychology which seems to suggest otherwise.[208] Our preferences are often entangled with one another; e.g., it is common to want wins for your offspring and losses for your enemies.

The more complex – and probably more important – disagreement concerns Gauthier's notion that morality is a matter of limiting one's pursuit of goals. Gauthier's view is that humans aim to maximize their own ends, which generates tension in the world as the maximization of

---

[205] Bird, 1990.
[206] Gauthier, 1987.
[207] If preferences are entangled, market transactions can't be guaranteed to be Pareto-improvements, undermining the application of the First Welfare Theorem.
[208] See Sen, Amartya K. "Rational fools: A critique of the behavioral foundations of economic theory." *Philosophy & public affairs* (1977): 317-344. and Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler. "Fairness and the assumptions of economics." *Journal of business* (1986): S285-S300. on how individuals' preferences are often socially valanced and sensitive to others.

our own ends usually interferes or conflicts with the ability for someone else or others to maximize their ends. For Gauthier, the perfectly competitive market is a panacea for the self-interest problem, because he sees it as eliminating that tension between ends. In a perfectly competitive market, the market mechanisms themselves do much of the work, relieving us of our need for morality. The main mechanism here that facilitates morality's taking of leave is the notion that rational self-interest will induce us to cooperate through the promise of mutual advantage. As Daniel Hausman makes clear, this requires that we have transparency and perfect knowledge of all economically relevant information.[209]

Telos: Preference Satisfaction

The idea that the purpose of financial markets is to satisfy preferences may have the same regulatory suggestions as a more deontological libertarian view described above, but the foundations of such a view are slightly different. The deontological, 'hands off', view argues, in essence, that to disallow financial markets to be an arena for preference satisfaction would be to violate something like 'individual rights.' Unlike the preference satisfaction theorists who see preference satisfaction as the mechanism by which some good (e.g. welfare improvement) occurs, there are those who would see preference satisfaction as the intrinsic aim, regardless of any good or bad effects that financial markets have on society.[210]

The teleological preference satisfaction idea would start from the position of what financial markets are for, rather than what values they support. In this form of the argument, to not disallow financial markets to be an arena for preference satisfaction would simply be to create something else– an institution that is not disrespecting anything (necessarily), but to

[209] Hausman, 1989.
[210] This view is simply that the institution's purpose is to allow for the non-violation of rights, regardless of whether that allowance actually makes us (individually or collectively) better off.

reform or change one institution into an altogether different one.[211]

My Account

In this section, I will defend my account of the telos of financial markets: to determine future markets for products and services by promoting stability and distributing risks, enabling those markets to provide for the benefit of all. I 'sneak up' on the issue, however, first getting clear on what a teleological concept does, and what metaphysical commitments it involves. I then proceed in describing two things we would want financial institutions to do: manage the economy's change in size (given facts about population and environmental changes) and manage the level of risk a society takes on (given a spectrum that ranges from the least amount of risk to infinite possible risk). I then describe why the state– even though there is clearly a role for it in finance– is ill-suited to meet these objectives.

Aristotle's Teleology

I start with a few words about what Aristotle is doing methodologically, and thus what I am doing when I use the language of 'teleology.' There is a misconception that what Aristotle does with teleological explanation is metaphysically 'spooky.'[212] Monte Johnson argues that it is a mistake to interpret Aristotle's teleology as being metaphysically stringent about some sort of cosmic or mystical purpose. He argues Aristotle's teleological arguments are best understood as forms of explanation used to make sense of organized systems. When Aristotle explains

[211] I'm not aware of anyone that makes this sort of argument explicitly, but some views which lack an overt appeal to individual rights (and their potential violation by disallowing financial markets as they exist now) look like they could be endorsing this view. This is more frequently found in arguments made by economists and financial professionals, who may very well fail to spell out that their views are founded in a notion of individual rights.
[212] See Chapter 5 Section 55 in Irwin, Terence. *Aristotle's first principles*. Oxford University Press, 1990., and note 18 on page 522 for a detailed description of these debates in the literature.

something by reference to its *telos*, he is, according to Johnson, identifying the activity or outcome that makes the organization of its parts intelligible. Teleological explanations answer a particular "why" question: why does a system have the structure that it does? In biological organisms, organs are explained by the functions they perform within the organism, such as the heart pumping blood or the eye enabling sights. Man-made artifacts like tools have ends derived from human intentions they have been imbued with. Political and social structures can be understood teleologically by looking at how their organization reflects the functions they serve. When Aristotle talks about the *polis*, he treats the structure of the *polis* in relation to the ends it enables for human life.

Johnson's account demonstrates that Aristotle's teleology is 1) not limited to natural organisms and is thus a general methodology for understanding organized systems whose parts coordinate toward some defining activity or outcome and 2) demonstrates that there need not be any strong metaphysical commitments required to use this form of explanatory analysis. The key implication for my work is that teleological language can be used coherently without committing to strong metaphysical claims about institutions or what their attributes are 'in nature.' We can identify a system's end, which allows us to explain its structure and evaluate how well it performs its characteristic function. It also allows us to critique developments, compare alternative (hypothetically) proposed purposes, and recursively orient the institution towards fulfillment of a socially productive end or set of ends. The normative method is, as follows: first, identify the function that makes sense of an institution's organization; second, figure out whether that institution fulfills that function effectively or not; third, figure out what arrangement would allow that institution to best fulfill its function. Of course, one criticism of this view is that we might plausibly find more than one function by engaging in this method. It is unclear what

Johnson would say in response, but one possible move would be to simply state that institutions can have multiple sub-purposes which jointly fulfil some overarching purpose.[213]

Other Normative Institutional Analysis

Seumas Miller posits an explicitly teleological normative account of social institutions.[214] He argues that institutions are normatively structured practices oriented toward morally significant (collective) ends. The legitimacy of social institutions depends on their capacity to realize those goods. This is clearly a teleological view, since realizing certain collective goods is the explicit point of certain institutions. The relevant goods are ones that can only be produced or sustained through coordinated social action. Examples include security, education, and so on. These goods share several features: they are generated through the joint activities of institutional actors, they are available to the broader community, and members of the community have a moral right to their provision.[215] So institutions are morally justified insofar as they help produce things that help people flourish. He combines 'methodological individualism' with a theory of joint action, which he uses to justify how groups and institutions can bear responsibility for outcomes even though action can be traced back to individuals having acted within structured

[213] In fact, my view on the purpose of financial markets could coherently be split into several sub-purposes, and this does not appear to be a conceptual issue. An issue would arise in cases where an institution has multiple plausible functions which in specific cases of human action might contradict the other. Johnson might also reject the idea that one could find more than a singular plausible purpose.

[214] In *Social Action: A Teleological Account*, he gives an analytical account of the theory and connects it to theories of join action, convention, and social norms. In *The Moral Foundations of Social Institutions: A Philosophical Study* he connects his previously published work into a narrative that first sketches out the range of theoretical issues that arise in the context of talking about a philosophy of social institutions, and then proceeds in the second half of the book to apply his teleological theory to specific institutions.

[215] This moral right is based on the requirement that individuals have the provisions for their basic needs to be met. Of course, provision of one's basic needs is necessary but not sufficient for an individual to flourish. Miller argues that "it is only when a certain threshold of aggregate need (or, for example, aggregated rights violations) exists that the establishment of an institution takes place" (pg.66, The Moral Foundations of Social Institutions). In other words, although institutions are required for flourishing, they are instantiated by the lower rung of aggregate needs (e.g. it is not sufficient for flourishing that a given institution, or perhaps set of institutions, exists).

roles governed by shared practices. I do not rely on Miller's account in full, nor do I attempt to defend it here. The relevance of his work is primarily to illustrate that teleological approaches to institutions already exist in contemporary political and social philosophy. My use of teleological language should not therefore be taken as introducing a novel or metaphysically unusual framework.

There are also several philosophers who pursue this kind of project without framing it in terms of Aristotle's teleology. For those who remain uneasy with Aristotelian terminology, the project can be understood more generally as part of a broader tradition of normative institutional analysis. Predecessors in this camp analyze institutions in terms of the goods or functions they serve within social life. For example, Michael Waltzer's *Spheres of Justice* is a normative account of institutions where different social 'spheres' have their respective internal logics which ground how they ought to function. The institutions are structured around the particular goods associated with each sphere, and injustice arises when the norms governing one sphere improperly dominate other sphere(s). An Aristotelian thinker would hold that the basic human *telos* (eudaimonia, flourishing, etc.) mediates the appropriate relations among these spheres. Lisa Herzog's *Citizen Knowledge* is also a work that engages in normative institutional theory. She looks at the epistemic conditions under which markets can function legitimately.[216] She argues markets depend on shared knowledge and structures such that citizens are enabled to understand and participate in economic life. Herzog also thinks there is a 'comeback' in institutional

[216] Her major concerns are that markets as they exist now, do not appear to have the information aggregation functions purported to exist. In particular, she argues for several 'knowledge issues' that undermine this function of markets including features of living in a 'post-truth' world (e.g. access to and proliferation of fake-news and misleading information), a decreasing deference and/or trust in expert testimony, and asymmetric information (which is problematic not just in itself, but in that many people are below a needed 'threshold' of information to have a relevant understanding of a particular issue.

political theory.[217] The methodological structure of this project, then, is not itself unique, but fits squarely into an established kind of approach to institutional analysis.

Institutional Desiderata

Contemporary economies continually change in response to environmental conditions. Unemployment rates fluctuate, demand for products rise and fall, and exogenous shocks like weather events, political interventions, or geopolitical conflict can suddenly disrupt how goods move through the world. These changes can produce large and rapid shifts in prices and production. For example, if conflict breaks out in a major oil-exporting region, the price of gasoline may spike almost immediately. Price adjustments are a normal feature of market systems. If we accept Hayekian information theory, price movements convey important information about scarcity and changing conditions and generate signals that are eventually supposed to induce appropriate changes in supply and demand that will stabilize prices.[218] Still, sudden changes in prices can have cascading effects on economic stability, producing inflationary pressures, financial stress on firms, and causing broader economic contractions. Ideally, there would be a way for the price changes to occur for the benefits that information theory says we should yield, while protecting society in which the price movement would otherwise effect from those secondary and tertiary effects of volatile prices.

---

[217] The language of 'ethnography' is used in this literature as a way of talking about what I describe to be 'institutional analysis.' See Herzog and Zacka (2019) who argue for an "ethnographic sensibility" in political theory. Longo and Zacka (2019) also argue for ethnographic research as a means to generate "problematizing redescriptions" that show us a different angle than we would see through traditional modes of analysis.

[218] This mechanism is typically understood in the following way: new information alters agents' expectations about future conditions, which indicates to them the quantities of goods they should either supply or demand. These adjustments pick out price levels that coordinate voluntary behaviors across market participants. In the ideal cases, this process will 'clear the market' - prices will adjust such that the quantity supplied equals the quantity demanded, which leaves no surplus product (something someone wanted to sell but that no one bought) or unmet demand (someone who wanted to buy something but couldn't).

Of course, I am appealing to the globalized market economy that most people are participants in, but these challenges are not unique to market economies. Any large-scale economic system will have to deal with uncertainty generated by environmental variability, demographic change in their population, and unpredictable shocks. Even a centrally planned socialist economy would face similar risks. A society organized without capitalist financial markets would still benefit from mechanisms that help manage uncertainty. Institutions capable of hedging against weather-related crop failures or stabilizing food supplies in the face of environmental volatility would still be important. In this sense, the problem of managing economic uncertainty is not specifically a capitalist or free market problem. It is actually a structural feature of complex economies of any type.

Every society needs an institution to help manage fluctuation in the size and activity of an economy. Economic expansion driven by entrepreneurship and innovation can generate large gains in welfare, and this is a central argument for a free market system. But rapid growth can also produce asset bubbles, resource misallocation, and systemic risk, all of which harm society *en route* to the theorized gains of a free market society. Conversely, economic contraction, whether it be due to exogenous shocks or predictable changes in population size or demographics, can lead to destabilizing declines in asset price or production.[219] An institutional mechanism capable of moderating these features of markets for products and services could, if done well, therefore contribute valuable stability to the broader economy.

The ability to manage changes in economic scale is thus a plausible desideratum for social institutions within an economic system. Ideally, such an institution would help smooth the volatility associated with both economic expansion and economic contraction and could both

[219] E.g., bankruptcy and firm failure, which can cause other firms' bankruptcy and failure, also causes unemployment, and so on.

dampen the formation of bubbles and mitigate the social costs of downturns. Whether an economy ought to grow or shrink, and by how much, in any particular period of time will depend on a range of factors including demographic trends, resources constraints, and medium-term expectations about production and consumption. What matters for present purposes is that complex societies benefit from institutional mechanisms that help navigate these changes in a more stable and predictable way.

Connected to the management of economic growth and shrinkage, is the management of what I will describe as a society's position along the range of possible risk. This range can be visualized as follows:

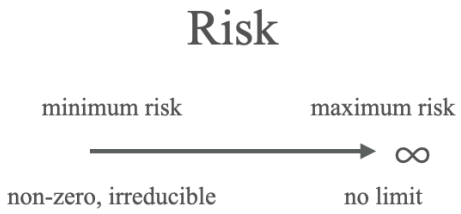

Figure 4: Visualization of risk that a society could take

If we consider the amount of risk present in a society, it is helpful to think of it as existing along a range. At one end of this range is the minimum level of risk a society could bear. That minimum is non-zero. Even the most cautious and conservative society (e.g., no interest or support for innovation, no investment capital for new business, and so on) cannot eliminate risk entirely due to the minimum level of products and services required for a population's survival. The project of producing sufficient food to sustain a population necessarily involves uncertainty regarding weather, disease, and the outcomes of harvests. Any system of production therefore

carries at least some unavoidable exposure to risk and uncertainty.[220] On the opposite end of the range, there is no upper bound: there can always be more risk added to the system. In modern economies it is possible to generate additional layers of risk through financial activities involving the derivative products discussed in Chapter II. Derivatives can create exposure that goes beyond the scale of the original economic activity upon which they are based. They can be constructed with increasing leverage, and so there is no straightforward limit to the amount of risk that can be created within an economic system. In our financial system, there are regulations that will end up constraining risk, like leverage ratio and capital requirements, but there are fewer regulations for private firms.[221] So while societies cannot reduce risk to zero, they can theoretically always accumulate more of it.

The asymmetry here has important implications for institutional design. If the lower bound of risk is necessarily greater than zero, and the upper bound is (absent intervention) unbounded and possibly infinite, society faces the problem of determining how much risk should exist within the system and how that risk should be distributed. Some risks are obviously productive and necessary (like those associated with growing crops) and thus clearly risks worth taking (e.g. the alternative to not subjecting ourselves to the risk of crops not growing is starvation– a 'guaranteeable' outcome with obvious downside). Other risks are not strictly necessary but seem to be required by even a conservative approach to unplanned economies or an economy that would respect the right for private businesses or banks to make loans. Those loans, even if conservative, have some risk of default associated with them that are justified by

[220] I mostly set uncertainty aside for this elemental description of my account because uncertainty is much more difficult to get a handle on than risk. This is not to obscure the influence of uncertainty on society, but rather to avoid biting off more than I can chew in this chapter.

[221] Here the free-market theorist will say that the free market will work in such a way as to regulate itself. This is plausible for the control of obvious risks between individual counterparties but is less plausible for systemic concerns.

their being necessary to achieve investment and growth. Riskier still would be those financial activities pursued for the sake of innovation or entrepreneurship. There is, by definition, some element of a new business venture or product that has 'never been done before' and thus carries with it various risks related to the failure of such a venture. Lots of investment into one kind of idea or new technology, even if on an individual basis the idea is good and there is a sound business plan, can collectively generate systemic vulnerabilities capable of destabilizing the broader economy. The widespread investments made into artificial intelligence, for example, seem to entail that the success or failure of AI firms will have a systemically important effect on the economy.[222] For this reason, it is plausible to view the management of the range of risk within an economy as another desideratum for some social institution. An institution capable of redistributing risk, when done well, helps ensure that necessary risks can be borne along with those within the desired spectrum of growth or development of the economy while preventing the accumulation of excessive or destabilizing downside exposure.

In teleological terms, then, managing fluctuations in the economic scale and the amount of risk borne by society can be understood as candidate functions of an institution which serves an economic and social system and so part of the *telos* of those institutions. If social institutions are justified in part by the collective goods they enable, then one such good could plausibly be the stabilization of economic activity in the face of risk and uncertainty, given inevitable change and an inability to predict such change with precision or accuracy. Financial mechanisms like credit markets, insurance, and derivatives, can, I suggest, allow societies to distribute and hedge risks associated with shocks. This has a stabilizing influence on the volatility that would

[222] Importantly, this systemic risk can plausibly emerge from everyone at the individual and firm level acting according to the most stringent standards of virtue and ethicality. We need not assume that anyone is doing anything wrong or illegal for this to emerge as a serious problem (to differentiate from systemic risks generated by, say, mortgage-backed securities, as was arguably the case there).

otherwise accompany innovation or disruption from exogenous shocks. It's not that economies should be insulated from all change or made static, because growth, contraction, and price adjustments are essential and inevitable features of a market system, but an institution that can mediate changes in ways that reduce negative externalities and so serve the people who inhabit this market system. I suggest that the management of economic expansion or contraction and amount of risk is a plausible dual-pronged purpose of a financial system. It seems clear, then, that society should want an institution to fulfill these functions. The question then becomes, what kind of institution is best suited to fulfill the *telos* of determining future markets for products and services by promoting stability and distributing risks in a manner that enables those markets to provide for the benefit of all.

Why the state is ill-suited to meet these objectives

If institutions are needed to manage the range of risk within an economy and to moderate fluctuations in economic scale, why shouldn't we think that these are functions that would be suitable for the state to fulfill? One difficulty here is the problem of pricing risk. The Hayekian information theorists would argue that a singular agent will be worse at pricing than a distributed network of agents with their own sets of information. Determining how much interest should be charged for a loan requires estimating the probability of details and the broader set of risks associated with the underlying activity. In contemporary market economies, these assessments are distributed across a large number of investors, lending, and financial intermediaries who serve to collectively evaluate riskiness, and prices emerge through decentralized processes in bond and credit markets.[223] Expected returns in these markets vary depending on the level of risk

[223] The best a government could do here would be to mimic that process. Oskar Lange thought this was possible, but Hayekians will argue that there is no way for the state to solve the knowledge problem (e.g. that knowledge is dispersed amongst agents). The socialist calculation debate is now mostly historical, but there could be a case to be

borne by investors. Again, the Hayekian will argue that this kind of price formation relies on information that is widely distributed and often difficult or impossible for any centralized authority to aggregate.

If the state were instead responsible for pricing risk directly, setting interest rates on loans or determining the returns associated with risky projects, it would be performing these functions from a disadvantageously centralized position. This should be thought of as conceptually distinct from the broader Hayekian critique of central planning or contemporary central banking. The narrower point is simply that assessing the value of risk across a complex economy requires the use of a lot of localized and rapidly changing information. Markets currently perform this task by way of the many participants whose decisions collectively generate prices that reflect varying risk assessments.

But let us suppose that the state takes on this role and does not change anything else.[224] There are two possibilities: First, state lending could prove profitable. Otherwise, the state incurs financial loss. In the first instance, where taxpayer funds are used for loans to new businesses, etc., and that turns out to be profitable, citizens may reasonably find objections to this.[225] If the state engaged in lending for things like education or housing, critics are likely to argue that the state is inappropriately treating citizens as sources of revenue rather than as beneficiaries of public policy. There will be obvious critiques to be made in how such an arrangement

made that prediction markets could provide a solution to the knowledge problem that would make socialist calculation possible. I discuss a different objection to prediction markets in Chapter V.

[224]It will be different if you also have state-controlled markets for products and services. This arrangement has well-known critiques: the informational burdens associated with large-scale economic coordination, the absence of price signals for guiding production decisions, and lack of incentives for innovation, efficiency, and responsiveness to consumer needs and wants.

[225] Fear of political influence being used for financial benefit e.g. only those businesses who are friendly to the regime will receive loans, loans being given to citizens for school (Student loans) or homes (mortgages) that citizens then find issue with the state profiting from them, the state circumventing market processes for ESG-related reasons and being subject to divestment protests from citizens.

incentivizes policy to be directed toward maximum educational attainment and maximal rates of homeownership or ownership of multiple homes.

If state investment proves unprofitable, the costs are directly borne by taxpayers, who will be critical of the state's loss of funds. Even if the state were to initially take on entrepreneurial risk (e.g., making loans to companies with new ideas or business plans), losses would be likely to instill an attitude of significant fiscal constraint and conservatism. This dynamic would be likely to push the state towards increasingly risk-adverse behavior. Risk-aversion at the institutional level would undermine the kinds of risk-taking and uncertainty that drive innovation and what is often necessary economic growth (due to increases in population size).

One might also worry about the political nature of state decision-making. Control of the state is often not fully democratic and can shift dramatically between administrations who have their own priorities. Many economic projects take more than four years to be complete or self-sustaining. If the state were the primary allocator of capital (e.g., if it functioned as the provider of venture investment), the direction and timeline for profitability would depend on the political aims of those in power.[226] Political leadership would be highly likely to favor projects that align with ideological goals and shorter-term electoral interests. This would introduce a high degree of systemic risk, because of the centralization of investment decisions within a political authority.

A natural response on behalf of a more social framework is that these concerns are not decisive: the state could, in principle, attempt to replicate the informational advantages of markets. It might do so by delegating decision-making authority across agencies, incorporating

[226] Where the state doesn't allocate capital but does manage risk is even worse– it's precisely the privatized upside, socialized downside' dynamic that people criticized in response to the government bailout and stimulus post '08 financial crisis.

local knowledge, and attempting to simulate decentralized price discovery. In this sense, the most plausible version of a centralized financial system would be one that attempts to replicate the epistemic structure of decentralized markets. So, to the extent that this kind of system would succeed, it would be doing so by reintroducing some market features (if not the distributional ones). To the extent that it fails to approximate these features, it will have issues with the mispricing of risk, politically distorted allocation, and lack of information integration into prices.

The state obviously plays an important regulatory and stabilizing role but seems ill-suited to serve as the primary institution responsible for pricing and allocating risk across an economy. Decentralized financial agents seem better positioned to perform this function because they incorporate dispersed information and distribute the consequences of risk-taking across multiple agents.

Democratic Deliberation

Although the state seems poorly suited to price and allocate risk, our intuition may still be that leaving these matters to the sole discretion of individual private firms is also a recipe for disaster. To be clear, I am not arguing that all questions about the overall level of risk within an economy should be left entirely to private market processes.[227] There are good reasons for our intuition that the state ought to play some role in fulfilling or sustaining these institutional functions. Decisions about how much total risk or systemic risk a society is willing to tolerate directly affect a wide range of public goods including economic stability, employment, and the ability of individuals to earn income. These consequences are broadly shared. Thus, it is reasonable to argue that the acceptable level of systemic risk is something that should, at least in

[227] For example, Hayekian information theory has no robust answer to questions of how systemic risk is supposed to be accounted for by individual market participants.

part, be subject to collective democratic deliberation, and not solely through decentralized market activity.

Rational market actors will likely evaluate risk using probabilistic models. These tools are useful for pricing but can fail to account for low-probability highly negative events (or "tail risks") that society will want to be adequately accounted for. Individual rational market actors can rationally pursue strategies that models predict will be profitable while also adding to systemic vulnerabilities that models do not see. In these kinds of cases, aggregate levels of risk (e.g., total investment in a an emerging sector with new technology) can become unbalanced and the financial system as a whole can become more precarious in ways that would seem impossible to hold individual decision makers accountable for.[228] Those individual actors may not have the necessary epistemic tools to make an analysis of economy-wide risks either.

For these reasons there is a plausible role for a form of democratic deliberation to evaluate the total risk profile of the economic system. Collective deliberation about societal priorities is something modern democracies embrace, for better and for worse. My aim in this project is not to offer a comprehensive model for how democratic deliberation about systemic risk ought to take place, but rather to make the more modest claim that questions concerning the acceptable level of risk within an economy involve normative judgements about collective welfare and stability. If that is the case, the process of answering such questions plausibly belongs (at least in some sense) within the domain of public deliberation, even if individual level

[228] For example, the economy is, as of Spring 2026, highly leveraged in the AI space. If we were to look at any individual-level AI-related loan or financial activity, they might all individually seem reasonable. Unreasonable, perhaps, may be the degree to which investment in AI 1) constrains investment in other categories (for example, because AI firms bid up the price of component that could be used in other products) and 2) generates systemic risk in the form of making the future health of the economy dependent on the success of a new kind of technology.

market actions (and thus pricing and risk allocation) remain within the domain of private individual market action.

In Chapter III, I argue that 1) most discussions of risk either ignore it altogether or 2) treat it in a sense that is insufficient to grapple with the nature of risk in society. I identify what is missing from existing treatments and contend that risk is a topic worthy of additional attention. I hope it is clear that once we take risk seriously, we cannot straightforwardly classify economic growth as either good or bad in isolation. Growth is always accompanied by exposure to potential harms, and these harms are often unevenly distributed. They can also be difficult to predict, and the risks can be temporally displaced from the possible or likely benefits. As a result, any evaluation of growth depends on prior normative judgements about how much risk is acceptable and who ought to bear it. Without such judgements in place, the distribution between 'good' and 'bad' remains undetermined, and the only limits which exist are the bounds established by the law and whatever indirect political processes exist to change the law. My proposal is that we shift focus toward the normative and summative question of *how much* risk society ought to bear in pursuit of growth.

Chapter V: Applying the Teleological Framework to Contemporary Issues in Finance

This chapter applies my teleological framework to contemporary issues in finance in order to demonstrate the usefulness of the approach. There are reasons to believe that this approach is helpful and, in some cases, better than existing tools for normative analysis of issues in financial markets. I describe three kinds of issues that my view can give us conceptual clarity on. First, are those issues which existing frameworks give us insufficient analysis of. Second, are issues which intuitively strike us as troubling, but which existing frameworks offer no serious method of analysis. This seems to suggest that these cases are not being adequately theorized. Third, are those issues which existing frameworks which yield answers that seem wrong when we take into consideration risk and futurity. Each of these kinds illustrates ways in which my approach delivers more normatively adequate judgements than the standard approaches used for analysis of financial markets. I cover each of these types in turn, using examples from contemporary finance to demonstrate issues that seem inadequately explained by existing tools wielded by philosophers and those with good-faith interests in reforming finance for the better.

We might think about finance in a way analogous to how Harry Frankfurt understands the will.[229] Frankfurt distinguishes between first-order desires and second-order volitions, where first-order desires are those things an agent is inclined to do, and second-order volitions are those things they reflectively endorse wanting to do. In financial markets, it is often the case that individuals act under the influence of sentiment, herding behaviors, or what Keynes called 'animal spirits.'[230] But the presence of such first-order impulses does not prevent reflective, second-order evaluation about what financial market participants ought to be doing. Likewise,

[229] Frankfurt, Harry G. "Journal of Philosophy, Inc." *The Journal of Philosophy* 66, no. 23 (1969): 829-839.

[230] Keynes, John Maynard. "The general theory of employment, interest & money.", 1936." *Reprint. New York: Harbinger* (1964).

the fact that market participants sometimes behave irrationally does not preclude us from thinking coherently at a macro level about risk exposure and from deliberating about how much risk we, collectively, ought to be willing to bear. On this view, financial regulation can be understood as an institutional version of a second-order volition. It is a reflective attempt to shape the structure within which first-order risk-taking occurs. We can ask not only how much risk individuals and firms do in fact take, but how much risk a society should bear, and what first-order behaviors we should and should not condone.

This chapter addresses what, if we take on my approach, can be said downstream about regulation and secondly, what regulatory matters will necessarily be outside the scope of my project. After briefly situating the contemporary U.S. regulatory environment, I argue that attending explicitly to risk and futurity tends to better orient us toward a distinctive set of regulatory priorities. The considerations I emphasize push in the direction of policies that are more sensitive to the temporal distribution of risk, pay more attention to systemic vulnerability, and are less presumptive of status *ex ante* rule-setting. I also gesture to more general regulatory reforms and potential mechanisms for instantiating the democratic deliberation component of my approach.

I also take stock of the limits to my approach. Most importantly, even a well-developed account of financial and economic risk will not resolve the distinct problems of political risk, and the effects those realized risks have on economic life. Political uncertainty can serve as an exogenous shock to even the best-functioning economic systems, and although my approach can help articulate why these political goings-on have clear negative impacts on the economic and financial functioning of society, my proposed method for better understanding financial risk and

futurity won't touch questions of risk and futurity that are mediated by *political* decision-making, such as tariff policy.[231]

Finally, I comment on risks worth taking from society's perspective. A central challenge for the view is that any comparative evaluation of heterogenous risks requires further theoretical development establishing a hierarchy or ordering mechanism. Emphasizing *both* that some risks are worth taking (for their associated potential benefits) *and* establishing this ranking of risk are important to avoid a collapse into stringent risk aversion. The intended upshot of the view is not for society to eliminate risk but to enable better discrimination among risks in a way that is both normatively defensible and involves broader democratic deliberation.

Types of Issues

Finance exposes a category of moral problems that are not readily captured by standard moral frameworks focused on outcomes or individual actions. When our evaluative tools are limited to utilitarian calculations of aggregate welfare and its distribution or deontological assessments of discrete wrongdoing, some of the moral features of financial systems can be obscured from analysis. Establishing normative views on how we ought to distribute risk and uncertainty is an important component of designing financial systems to function in a manner that will contribute to society. I am not claiming that all existing moral theories fail when applied to any issue in finance. Instead, I claim that ethical theories fail in different ways when applied to many issues in finance, and those failures expose limitations in how existing theories can engage unintended non-intentional emergent action, institutional structure, and endogenous risk. When I

[231] If everyone (including politicians) were to be attuned to the purpose of financial markets, it is possible this would become less of an issue (in the case that it is not the politician's intention to harm society in the associated ways, but a simple lack of understanding which allows them to take these actions which have consequences unintended by them). To make this point, let us consider possible mechanisms for democratically deliberating on risk: it seems plausible that the polis could decide to have more or less financial risk, and that would have discrete metrics for things like assets under management and capital requirements. How would a polis deliberate or come to enforceable conclusions about political risk? The polis deliberation is the one very possible *source* for this kind of risk.

highlight ‘existing views’ I mean those views which avoid considering the unique features of finance. Additionally, I fear that treatment of these unique features will necessarily be inadequate when considered from either deontological or consequentialist lenses.

Type I: Insufficient Views

The second type of problem involves views which, when applied to issues within finance, are not obviously incorrect but are incomplete. In these cases, the theory being applied captures something of important moral relevance within finance but leaves unaddressed important features that exist within finance. This kind of analysis, although not incorrect, is too ‘thin’ to be of sufficient use in preventing the distinctive harms financial markets can cause.

Issues like this include the analysis and criticisms people make about cryptocurrencies. Existing critiques largely focus on two insufficient responses: One dismisses cryptocurrencies as effectively categorically fraudulent, inferring from the frequency of frauds and scams that the entire technology is deficient. Another critique takes no stance on the usefulness of cryptocurrencies but, so long as standard consumer protection constraints are satisfied, has no take on the normative status of cryptocurrencies, regarding them as simply a perhaps unusual preference of some individuals. Cryptocurrencies like Bitcoin are interesting, because their original proponents were explicitly making a risk-related argument for them. Traditional finance, they argued, exposes individuals to risks they may reasonably wish to avoid: the risk of bank failure or mismanagement, politically motivated restrictions on access to funds, and exposure to central bank monetary policy. In the aftermath of the 2008 financial crisis, these concerns were (and largely remain) legitimate worries. Early crypto enthusiasts were worried about systemic risk, which is indeed something we ought to worry about. This makes it ironic that the reality of

crypto adoption has thus far largely *concentrated* the very risks enthusiasts were originally worried about.

The crypto financial system has developed such that users typically rely on a single institution to perform multiple financial functions that are dispersed across the traditional financial system. Custody ("deposits" in traditional finance, your "wallet" in crypto), trading, and payments are often handled by the same platform within the crypto financial system. In the traditional financial system, individuals typically store money in banks (e.g., Wells Fargo, Chase, community credit unions), transact through payment networks (e.g., Visa, Mastercard, or cash), and invest using brokers (e.g., Charles Schwabb, ETrade, Robinhood). While none of these institutions are risk-free, their functional separation (and regulation) provides a degree of insulation against any single point of failure. So, in the (unlikely) event that your broker is hacked and your entire stock portfolio is suddenly wiped out, you would still have mechanisms to spend money that is stored elsewhere.

In this respect, the collapse of a firm like FTX was functionally equivalent, for its users, to the simultaneous failure of a major bank, a payment processor, and a brokerage firm. As a result, even users motivated by consequentialist concerns (beliefs about bad outcomes produced by traditional finance) or deontological concerns (beliefs about rights violations or coercion) will find that despite attending to risk in general, they have not meaningfully reduced the kinds of risks they were worried about (and have likely exacerbated such risks instead).[232]

Philosophical analysis of cryptocurrency often gets something right here: there are serious systemic risks in our existing financial system. Assessments of crypto as promoting

[232] This is to say nothing of the individual virtue theoretic concern which might go something like: traditional finance attracts or creates vicious sorts of people, who are people with bad character and thus ought not control global finance. It seems obvious in this case that crypto has absolutely attracted a kind of person with character traits that tend to include the propensity for manipulation and grift.

autonomy or reducing dependence on legacy institutions[233] may be partially accurate, just as criticisms emphasizing fraud and speculation capture genuine problems. What these analyses tend to miss, however, are the institutional features that (for both traditional finance and the crypto financial system) can serve to manufacture additional risk (in the case of crypto), concentrate or illegitimately control (in traditional finance, in the eyes of crypto enthusiasts) risk. As a result, existing views lack the conceptual resources to evaluate cryptocurrency as a proposal for institutional redesign, particularly with respect to how risk is authorized, distributed, and ultimately borne across society.

Type II: Non-Views

There are important developments in finance that are not meaningfully discussed within existing moral frameworks. These are not cases where a theory delivers the wrong answer, nor cases where it delivers a partial one. Rather, these are cases in which no answer is generated; where it seems ethical theories have nothing much to say, because the phenomena in question fall outside the bounds of what the theory can touch. Addressing these issues requires an understanding of the distinctive features of finance, particularly its relationship to futurity and risk.

Consider situations in which there is no principled account of how much risk ought to be taken on as a society. In such cases, investment decisions (e.g., large-scale investment in emerging technologies like AI) are constrained only by the willingness of individual counterparties to assume risk.[234] Consequentialist views have little to say here, because the

[233] Which are not problematic in themselves but are in virtue of their lack of democratic oversight, the argument goes.
[234] So long as they behave according to existing rules about fraud and consumer protection and so on. But none of those constrain an industry as a whole.

consequences of potentially transformative and industry-shaping technologies are deeply uncertain, and because technological change in general has often historically resulted in (both good and bad) unintended outcomes. Deontological views fare no better, because no individual actor has violated a right or acted impermissibly. From the perspective of both frameworks, nothing has gone wrong.

Deontological approaches that do address institutional or systemic concerns tend to do so retrospectively, through analyses of wrongdoing after the fact, or prospectively through notions of complicity or domination once a wrongful scheme is already underway.[235] Determining those standards appears to require a form of contextual judgment that's sensitive to environmental and circumstantial constraints which rule-based theories are ill-equipped to provide. That the amount of risk a society ought to take on is central to my account, and it should depend on background conditions.

The issues I describe as "non-views" may, at some future point, become insufficient views if and when theorists attempt to apply familiar moral frameworks to these cases without attending to the unique features of finance. Descriptively, this absence may simply reflect a lack of awareness of certain developments. Normatively, though, even if attention is drawn to these cases, existing theories struggle to evaluate them in a way that fully addresses what is at stake.

For example, the shift from defined-benefit pension plans to defined-contribution plans such as 401(k)s is an illustrative example: under this shift, the risk of ensuring adequate income in retirement was transferred from employers to individuals. This transformation did not initially register as a moral problem within dominant frameworks, but the consequences of it are

[235] Domination is going to be a flawed paradigm to apply to issues that relate to macroeconomics, because it would appear that some degree of domination is required to keep the system going (e.g. interest rate changes). See Petit, "The Domination Complaint." *Nomos*, Vol. 46, Political Exclusion and Domination (2005), pages 87-117.

becoming increasingly visible as members of Gen X approach retirement (without adequate savings). Early indications suggest that this risk transfer may result in widespread harms that were structurally embedded long before we could observe their effects. But if we had attended to risk and its distribution, we could have understood the significance of and critiqued this change beforehand. Without a robust understanding of risk, we see this as an issue of freedom, and the libertarians have a strong case for this transition. In some sense too, the domination theorists could see this transition as a good thing because employees are no longer being 'dominated' by their employer on how to save for retirement and have access and purchase power over more of their income.

This case illustrates how finance can generate significant risks without immediately producing either wrongdoing or harm. This case also illustrates why a framework attuned to the distinctive temporal and institutional features of finance is necessary to prevent harms and wrongdoing which would otherwise seem permissible.

Type III: Incorrect Views

The first kind of view to discuss is one which gives us a coherent account of an issue, but one that seems wrong when considering the futurity and risk characteristics of finance. The problem here is not merely that they reach undesirable conclusions, but that alternative views lack the resources to explain why they are wrong. For example, under a libertarian view (which I take to be a thin sort of deontological view), any voluntary exchange is presumptively permissible so long as the set of procedural rules governing individual behavior have been met.[236] If we accept this view, then there appears to be little in the way of justifying intense

[236] Mack, Eric. "In defense of 'Unbridled' freedom of contract." *American Journal of Economics and Sociology* 40, no. 1 (1981): 1-15.

regulation of financial products and services because their impact is only seen on a systemic scale.[237] This is what makes finance different from a libertarian approach to interpersonal markets.[238] It's unclear that a consequentialist will be able to convince the libertarian otherwise, other than to make exceptions from cases of historical turmoil. A libertarian response to such suggestions often takes the following contours: the undesirable features in question are ultimately a byproduct of securing the benefits of market systems (or of avoiding the moral costs associated with restricting them), or alternatively, that such outcomes are best understood as the result of prior distortions introduced in the system by state intervention. Appeals to past harms and poor outcomes such as we saw in the wake of the 2008 financial crisis, may only justify specific reforms and changes in regulation tied to mortgage lending or securitization. Further reforms are met with resistance on free market grounds and are largely unsuccessful when policymakers cannot point to past harms.[239]

It seems, then, that existing counters to the libertarian view can usually only motivate exceptions and will not quite justify regulation outside the bounds of standard consumer protection for seemingly irrational retail trading or similar financial activities.

My framework is useful here because it is only by attending to the unique features of finance that we can reveal where and why incorrect views go wrong. In theorizing about markets for products and services, shortcomings are often accessible via our institution based on our

---

[237] And looking to consequences to assess the permissibility of an action, product, or practice isn't part of a traditional libertarian view.

[238] The 'systemic risks' related to products and services are either diffuse (the 'risk' of Labubus becoming very popular) or are more related to production concerns (everyone wants computer chips, so the proliferation of AI investment risks making costs of production for iPhones more expensive) but don't seem to rise to the level of concern that one finds in financial markets.

[239] Wallison, Peter. "The Dodd-Frank Act: Creative Destruction, Destroyed." WSJ. Aug 31, 2010. A criticism of the Act from another angle questioned whether the oversight boards created would intervene on government action, if such action posed a risk to the financial system. See the section from Franklin Allen in "A Major Transformation: The Pros and Cons of the Dodd-Frank Act." Knowledge at Wharton. September 7, 2010.

experience as consumers. But many of the central features of finance are opaque without specialized knowledge about how finance works.

Cryptocurrencies

I have already described the motivations that led some to adopt cryptocurrencies and the kinds of risks they were originally intended to mitigate. Now I consider from first principles a best-case version of cryptocurrency. The aim here is not to defend existing crypto systems as they currently operate, but to show that even under ideal conditions, cryptocurrencies possess structural features that make them ill-suited as instruments of risk mitigation and economic stability (as currencies ought to be). It would also be inadequate to regard them as mere distractions or permissible odd preferences of market participants, because of the way they negatively impact the financial system with respect to its *telos*. The *telos* of financial markets is to determine future markets for products and services by promoting stability and distributing risks, enabling those markets to provide for the benefit of all. Cryptocurrencies undermine this purpose by introducing additional sources or risk without connection to identifiable benefits (either in terms of socially productive ends, or financial benefit to the individual), while siphoning capital away from the real economy. When someone buys crypto assets, that individual converts state-backed currency (like USD) into crypto assets that are largely disconnected from any productive economic activity associated with the real economy. In the counterfactual case, were the individual not to invest in crypto assets, at a minimum, their bank would have a higher depositor base from which to engage in lending. The individual is also clearly not allocating those dollars towards investment in projects tied to the production of goods and services. So capital is not only redirected away from investment purposes to fuel the real economy, it is also withdrawn from the traditional banking system, with no robust alternatives in

the crypto financial system to replace either of these practices. In addition, use of crypto platforms themselves has proven to be individually risky.[240] This should illustrate the value of my approach in critiquing something that strikes many as intuitively fraught but cannot be adequately theorized about using existing frameworks.

Critical literature on cryptocurrency tends to emphasize its most troubling features as instantiated thus far: fraud[241], speculative excess[242], and facilitation of criminal enterprise.[243] These critiques are descriptively accurate. It is undeniable that the crypto ecosystem has attracted actors who resemble snake-oil salesmen of yore. But these kinds of criticisms are unlikely to persuade crypto's staunchest proponents because they will dismiss these problems as typical of any early stage of adoption and temporary defects that can either be remediated with improved regulation or that gradually disappear as markets mature. So, it would be a mistake to evaluate cryptocurrency solely on the basis of its current manifestation. Even if we imagine a version of cryptocurrency that successfully deters fraud, prosecutes wrongdoing, and operates transparently, there will still be features that undermine its suitability for managing collective risk over time.

Many of the features that crypto proponents identify as flaws of traditional finance are, on my view, mechanisms that serve an important social function: enabling governance and distribution of collective risk. For example, the crypto proponent's criticism of central bank control over the money supply, which is framed as an instance of wrongful centralization, is used

[240]It is notable that crypto remains relatively disconnected from the traditional financial system, as the failure of large crypto firm FTX did not have a significant effect on the real economy. At the same time, this insulation is consistent with the claim that capital flowing into crypto is effectively diverted from uses tied to productive activity in the real economy.

[241] See reporting on the OneCoin scandal, including Voris and Yang, "'Cryptoqueen' Sidekick Gets 20 Years for $4 Billion Ponzi." Bloomberg. September 12, 2023. See also PlusToken and BitConnect. Exchanges people might also out in this category include Quadriga and FTX, led by Gerald Cotten and Sam Bankman-Fried, respectively.

[242] Haykir, Ozkan, and Ibrahim Yagli. "Speculative bubbles and herding in cryptocurrencies." *Financial innovation* 8, no. 1 (2022): 78.

[243] See the literature review Arnone, Gioia, Giovanni Scire', and Enzo Bivona. "The (mis) use of cryptocurrencies by criminal organizations: a systematic literature review." *Digital Finance* (2025): 1-37.

to mitigate the negative effects of economic downturns. From this starting point, decentralization of a monetary system appears to be an improvement precisely because it eliminates the discretion in monetary policy and large-scale interventions that my framework casts as a stabilizing influence.

Centralized control of the money supply has clear advantages in times of economic instability, although it is control which can be wielded in an improper way.[244] In a fully decentralized monetary system, there is no institution with the authority to respond to exogenous shocks, stabilize expectations, or dampen volatility. While crypto proponents celebrate the fact that such a system would preclude taxpayer-funded bailouts or inflationary monetary expansion, this view forgets that the socialization of downside risk would still occur, just in an emergent and uncontrolled way.

In the significant opposition to bank bailouts following the 2008 financial crisis, many were opposed on the grounds of fairness. Bailing out banks was seen as instantiating a socialization of downside risk: the costs that would have been felt only (or largely) by the bank, would instead be felt by the taxpayer. I do not mean to imply that this criticism is without merit, but the socialization of economic risk exists regardless of government intervention. The question of relevance then is not whether downside risk will be socialized, but how, when, and upon whom that risk will be concentrated.

From this perspective, the capacity of central banks to respond to economic shocks and stabilize expectations is something which can do significant good. The ability to distribute the costs of economic crises across time and people is of moral and economic value. This does not negate the fact that centralization can be politicized in ways that negate this good or can

[244] In this way it is much like 'centralized' control of a state– the structure itself is not enough to ensure good (or bad) governance, it is also about who rules.

plausibly cause even greater harm, but the centralization itself does not seem to be an inferior organizational structure morally speaking. And we do seem to have evidence that decentralization will not have the capacity to manage the temporal and distribute dimensions of risk in a satisfying way.

Crypto proponents critique central banking for a couple of core wrongs which relate to fairness and interference with markets. First, central banks obviously influence the money supply with multiple levers, each of which serves to increase liquidity. Critics can regard this action as a distortion of market processes and, importantly, as a primary driver of inflation. By this logic, what also straightforwardly seems to be a driver of an increase in the money supply is lending practices by expanding the amount of dollars consumers have access to. This private increase in the money supply is not something that would be addressed or fixed by cryptocurrencies whatsoever and does not seem to have any friction with the libertarian view. Consumer debt has increased steadily over time, and it is far from obvious how a crypto-based financial system would address this dynamic. So long as crypto systems rely on intermediaries that issue loans or extend credit (which is hard to imagine any financial system doing otherwise), similar inflationary and stability effects could arise, with few levers to alleviate such effects. Restricting loan or credit practices within crypto markets would itself seem to conflict with the market-based principles that many crypto proponents and libertarians endorse. In this sense, the crypto project largely fails to engage with the futurity-oriented dimensions of finance: the need to manage collective exposure to risk over time, to respond to exogenous shocks, and to govern macroeconomic stability. Even in its best-case form, cryptocurrency offers few tools for addressing inflation, systemic vulnerability, or coordinated economic response. It may be appealing to reject the legitimacy of existing institutions, but it does not appear to offer any

replacement of the valuable roles existing institutions play in managing volatility and risk (though we should concede that existing institutions could do this project better).

Meme Stocks, Retail Trading, Prediction Markets

The recent trend in 'democratized' finance, whereby more everyday people get involved in various aspects of the financial industry (and other aspects of previously outlawed behavior like sports gambling become financialized), is difficult to critique with the existing tools philosophers have from analysis of markets for goods and services. Retail stock trading, meme-stock speculation, and the financialization of activities such as sports gambling resist straightforward evaluation using familiar frameworks focused on consumer choice, preference satisfaction, or individual welfare.

The rise of retail trading on platforms like Robinhood has been documented in public media.[245] Much of the accompanying commentary focuses on consumer-protection or individual welfare concerns using a straightforwardly consequentialist diagnostic: individuals lose money engaging in these activities, and therefore the activities must be bad. To this, the libertarians predictably respond with a freedom argument: those retail traders have preferences to make these trades with their money, and whether they are harmed is beyond the scope of our normative concern. Libertarians who understand markets primarily as arenas for voluntary– and so morally legitimate– exchange preference satisfaction tend to accept these losses a matter of mere imprudence, and so as morally irrelevant. If individuals freely choose to engage in risky trading, there is no justification for interference (unless they are coerced or defrauded). This response becomes less plausible once we distinguish financial markets from markets for goods and

[245] Massa, Annie, and Sarah Ponczek. "Robinhood's addictive app made trading a pandemic pastime." *Bloomberg* (2020).

services. If retail trading were merely analogous to purchasing a discretionary consumer good, then the libertarian response would be more compelling. However, financial markets play a distinctive institutional role: they are not merely sites of exchange, but mechanisms for risk distribution and capital allocation that shape future markets for goods and services. If we accept this disjunction, we have explanatory power to distinguish concerns about widespread retail trading beyond those of an individual's potential financial losses.

Retail stock trading generates two distinguishable categories of risk: first, the aforementioned risk of harm to themselves, including financial losses and also possible behavioral harms like addiction. That kind of concern largely falls within the domain of consumer protection and is not impacted by my project. Second, is the risk posed to markets themselves, which is of central concern.

Prices in financial markets are supposed to serve an informational role as signaling mechanisms. Even if prices are never accurate, there are obvious better or worse conditions under which they approximate underlying economic realities. Professional traders and institutions are often justified in virtue of their role in facilitating price discovery and contributing to the efficient allocation of capital. These activities are resource-intensive and epistemically demanding, involving substantial investment in data collection and analysis by full-time employees. In a world where retail trading offers a predictable supply of irrational trades, hedge funds have less incentive to engage in these more resource-intensive methods for uncovering over- and under-priced assets. Instead, they can effectively do the opposite of what retail traders are doing. I characterize this dynamic as analogous to feeding bread to ducks: once the ducks are occupied with easily accessible, low-quality food, they no longer forage in ways that sustain the ecosystem. Similarly, when markets become saturated with irrational retail

activity, professional actors may rationally respond by arbitraging that behavior rather than contributing to meaningful epistemic investigation and price discovery. The resulting equilibrium undermines the institutional purpose of financial markets.

This concern has two possible regulatory responses: one would be to heavily regulate and restrict retail trading. I think this is not a good response, largely for freedom concerns: to prohibit an individual from investing in the stock market would be a violation of their freedom. Instead, I think meaningful improvements can be achieved with less restrictive measures. First, we could marginally increase the transaction costs, returning to industry norms prior to Robinhood's emergence as a broker in the mid-2010s. Robinhood's rapid adoption was driven in part by its industry-first elimination of per-trade fees. Prior to their existence, retail brokers commonly charged modest per-transaction fees in the range of ten to fifteen dollars. These fees do not meaningfully deter long-term investing but would impede the kind of retail trading I am worried about. Reintroducing nominal transaction costs could discourage behavior that undermines proper market functioning without preventing access to capital markets more generally.

Secondly, we can reshape the social narrative around retail trading. Presently, many participants do not regard their activity as a vice, instead seeing it as a value-neutral hobby or even as a part of engaging in responsible financial planning and management. If we articulated retail trading (of the sort exemplified by 'meme stock' trading based on hype or non-economic values) as a vice more akin to smoking, we could deter those who were not conceiving of the significance of their actions.

More broadly, many financial instruments involve risk-taking with a difficult to discern connection to social value or the underlying real economy. In some cases, it is unclear whether counterparties are merely speculating against one another or whether their transactions facilitate

the distribution of some ineliminable economic risk that arises from productive activity. Even when some speculators exist in a market, say, for commodities futures, they are still plausibly connected to the broader economic benefits of price stabilization given the two-party nature of financial transactions. Without an account of the institutional purpose of finance, we lack the conceptual resources to distinguish between risk that is connected (however distantly) from boots-on-the-ground economic activity and that done for the sake of non-economic reasons (like motivation to 'stick it to the man,' engage in political activism, or participate in an online community to achieve a sense of social inclusion).

Another recent development in financial markets has been the emergence of futures-contract based 'prediction' markets on platforms like Kalshi and Polymarket. These markets allow participants to buy and sell contracts that pay out based on the binary outcome of some future event.[246] Such events include whether inflation will exceed a certain threshold, the closing price of the S&P 500 index, whether the duration of a White House press briefing would be longer than 65 minutes, whether Cardi B would take part in Bad Bunny's Superbowl performance, and so on. For information theorists, the idea here is that contract prices are supposed to function as a probabilistic forecast: the higher the price of the YES contract, the likelier the market judges the event to be.[247] On a Hayekian picture, such markets improve

[246] Buying a YES contract is treated the same by Kalshi's system as selling a NO contract (e.g. the similarity between buying a call and selling a put). Each contract requires a counterparty, but Kalshi, as the market maker, steps in to facilitate each side, and then matches up buyers and sellers internally where possible (which is the same role that exchanges play in derivative markets). This commits Kalshi to managing the risk of trades that don't offset. So, if there are more YES buyers than NO buyers, Kalshi will be a net YES seller (which exposes them to negative risk if the event resolves to NO). In other words, each contract will have market participants on both sides, but Kalshi will likely fill the gap between the difference in demand.

[247] Suppose there is a Kalshi contract that resolves to YES if tariffs are upheld by the Supreme Court and NO if tariffs are struck down. For any trade to occur, there must be counterparties willing to take both sides of the contract (i.e., I am unlikely to find someone also desiring to trade a contract about what I'm going to eat for dinner, because no one else is interested in that). In this sense, positions are matched: one party will bet on the binary resolution of the contract to YES and the other party will bet on the binary resolution of the contract to NO. The contract pays $1 if the specified event occurs (resolves to YES) and $0 otherwise (resolves to NO). The prices of the YES and NO

information signaling by aggregating dispersed knowledge. Insider trading is regarded as normatively neutral or desirable, since privileged information helps move prices closer to the truth.[248]

Critics of prediction markets tend to view them as another form of speculative activity, sharing many of the objectionable features associated with gambling. Prediction markets share with apps like Robinhood the critique from cognitive psychology: that one objectionable part of this technology is the ease of use and propensity to nudge individuals to become addicted. What makes prediction markets unique, usually more objectionable than gambling, are the things being gambled upon. The moral status of a given contract on a prediction market is partly a function of what is being wagered upon: betting on a Superbowl performance might be relatively innocuous (and subject perhaps to the same permissibility standards as sports betting), while betting on events like whether Ayatollah Ali Khamenei would 'leave office' is of more serious concern.[249] One might worry not just that placing bets on things like assassination and war is vicious or noxious, but also that such practices might generate perverse incentives (e.g., incentivizing

positions should sum to $1 (less the spread between the positions; we should expect a greater spread the greater the asymmetry between YES and NO sentiment).

[248] If we take a fair coin flip as an example, imagine an event contract about whether the outcome of the coin flip will be heads. If the coin is indeed fair (or no one knows that it isn't), we should expect the probability of heads to be 50%. Accordingly, the YES contract (which pays $1 if the coin lands on heads) should trade at about $0.50, and the NO contract (which pays $1 if it lands tails) should also trade at about $0.50. Now suppose someone knows that the coin is heads on both sides. They possess insider information. From their perspective, the true probability of heads is 100% (barring violations to the laws of physics). But the market is pricing a heads outcome at $0.50! The informed trader can buy YES contracts at $0.50, and get paid $1 for each contract, making a profit of $0.50 for every contract they buy. But as the trader continues buying, the price of the contract goes up. If the contract on a coin flip rises, let's say from $0.50 to $0.51, and eventually to $0.60, the contract is seeming to come closer to reflecting reality. The insider, although the prices of the contracts are rising, still rationally prefers to purchase contracts as their price rises, because they know they will receive $1 from each one (they'll just receive, say $0.30 profit on a contract purchased for $0.70 and so on). The profit they receive from each contract will decrease as the price of each contract goes up, but they will still make a profit as the price of the contract approaches $1.

[249] To be fair to Kalshi, both Polymarket and Kalshi had contracts on whether Iran's leader would leave office. Kalshi's contract specified that his death or assassination would not satisfy the terms of the contract. So, when he involuntarily left office by being killed, Kalshi did not pay YES contracts out at $1 but rather at the price they were being bought at. Kalshi is now being sued for $54 million dollars in a class-action lawsuit over this. See Reuters, March 6, 2026. "Kalshi sued over ouster of Iran leader prediction market."

someone to carry out an assassination so that they can trade on such a contract beforehand and subsequently profit).

The worry here is more generally about the behavioral and social costs associated with widespread participation (and detractors might hold this view while granting the possible informational advantages).

Early empirical research suggests that prediction markets may in some cases possess modest information advantages over alternative forecasting mechanisms. Recent work comparing prices on Kalshi contracts to contemporaneous Consumer Price Index (CPI) predictions indicates that such markets may slightly outperform implied predictions from options markets and surveyed experts.[250] Even granting this information edge, however, it would be a mistake to evaluate prediction markets solely on the basis of predictive accuracy. Any assessment must also weigh the broader incentives and risks such markets generate. Prediction markets introduce perverse incentives in several realms of human behavior. We are inducing the insertion of unexpected (but actionable by some insider) actions to generate short-term volatility that can be anticipated and exploited.

This generalizes to incentivizing everyone to introduce volatility into any domain in which they exert influence.[251] Analogous surveillance akin to that required by professional traders to prevent insider trading is unrealistic to extend to the general public (not to mention a plausible invasion of privacy). Furthermore, it seems that those on board with prediction markets would explicitly reject the idea that we should deter insider trading. In fact, if information efficiency is all we are after, insider trading seems to be a benefit of such markets.

[250]Diercks and Wright, "Kalshi and the Rise of Macro Markets." NBER. DOI 10.3386/w34702

[251] The adverse selection point to be made here is that, if there is a market in something extremely niche and specific, it should occur to a would-be transactor in that market that the reason the market exists is because someone who knows more information than you about that thing instantiated the market.

This dynamic has similarities with earlier debates about insider trading in publicly traded securities. One argument in favor of allowing insider trading holds that it aligns managerial incentives with performance of their firm. Managers who can trade on insider information are motivated to improve their companies, since they can personally benefit when positive information becomes public. This incentive does seem to exist, but critics would argue that such incentive risks being overshadowed by a stronger and corrosive incentive to induce volatility rather than positive performance. Managers just as easily make choices which, when public, would predictably induce vast price swings. If the manager can strategically time announcements of information, they have a winning strategy they can trade on prior to such announcements which is more replicable than straightforward improvements to their business. It is cheaper to cause chaos you know about and can trade on beforehand than it is to improve a business.

Prediction markets have an analogous structure. Proponents of prediction markets argue that they incentivize the aggregation of useful information. They may, but they also create incentives for actors with influence over outcomes to manipulate the informational environment. The concern is not merely that participants may trade on privileged information, but that they may act so as to create markets they can predict and engineer volatility within. There is also a comparison to be made with a critique of high-frequency trading (HFT). HFT may improve price accuracy or accelerate the incorporation of information into market prices. But is the accuracy gained the sort of accuracy-improvement that helps, say, humans evaluate assets? If improvements in price accuracy or speed exceed the limits of human perception, their value (even in financial terms) seems pointless. The question for both HFT and prediction markets, then, is whether the informational improvement is worth the cost, where we both need to demonstrate the usefulness of the information improvement and take note of the downside costs.

Regulation

Financial regulation in the U.S. is administered through a network of interrelated institutions charged with enforcing legislation enacted by Congress. The Federal Reserve, the Securities and Exchange Commission (SEC), the Commodity Futures Trading Commission (CFTC), the Office of the Comptroller of the Currency (OCC), the Consumer Financial Protection Bureau (CFPB), and the Federal Deposit Insurance Corporation (FDIC): each has a role, and those roles sometimes overlap. The Federal Reserve, for example, has a mandate to do monetary policy, balance employment rates and inflation, and to oversee financial system stability.[252] The SEC oversees securities market activities and handles disclosures, the CFTC regulates derivatives markets, and the OCC and FDIC supervise and insure banks.[253] The contemporary regulatory landscape is shaped by the Dodd-Frank Act, which was passed in the wake of the 2008 financial crisis and substantively changed the rules governing financial activities. It introduced new capital and liquidity requirements and created new mechanisms for identifying and managing systemic risk, both of which are discussed further below.[254]

Normative analysis of financial regulation has largely been developed within legal scholarship rather than moral philosophy. For much of the twentieth century, the dominant method of evaluating regulation has been cost-benefit analysis.[255] Under this standard, regulatory action is justified when its expected benefits outweigh its expected costs. Critics of this approach

[252] They have a few tools to achieve these objectives. These include setting the federal bank interest rate, purchasing or selling government securities (e.g. treasury bonds), and managing the discount window, which is a facility where banks can borrow reserves directly from the central bank against some collateral.
[253] Congress.gov "Who Regulates Whom? An Overview of the U.S. Financial Regulatory Framework."
[254] The act created a couple of oversight organizations including the Financial Stability Oversight Council, the Office of Financial Research, and the Orderly Liquidation Authority (each designed with the intention to oversee, measure, and/or prevent some part of what led to the 2008 crisis). The Volcker Rule mandated a separation between 'proprietary' trading (trading for one's own bottom-line) and facilitation of customer orders.
[255] Sunstein, Cass R. "Some costs & benefits of cost-benefit analysis." *Daedalus* 150, no. 3 (2021): 208-219.

have argued that it systematically privileges short-term, quantifiable interests over long-term stability and resilience, which is especially problematic in domains like finance in which there exists significant uncertainty and systemic interdependence between institutions and agents. Having a cost-benefit standard requires strong evidence. Traditionally, regulators have been expected to justify new rules with concrete empirical evidence of harm. This, of course, is difficult to do if such harm has not yet materialized. Unsurprisingly, under this paradigm, regulation tends to be reactive to bad outcomes rather than proactive, because there is no credible reason (under this kind of standard) to justify an additional regulatory burden. This bears similar features to my critique of the pitfalls of a consequentialist approach to finance identified earlier in this chapter. In both cases, intervention is stalled until harms, or downside risks are realized, even if they could have been foretold.

In the regulatory landscape, these limitations were recognized and eventually overturned by a new regulatory paradigm based on the precautionary principle.[256] This was developed within the context of environmental and climate regulation and holds that the absence of full scientific certainty should not be used as a reason to postpone the implementation of regulation aimed at preventing or mitigating serious or irreversible harm. On this view, regulation ought to be forward-looking and cautious, and regulation is warranted where it serves to reduce futures risks rather than postpone regulation until such time that full empirical confirmation is available. Cass Sunstein argues that while attempting to right the wrongs of the strong evidentiary standard paradigm, the precautionary principle is in fact paralyzing due to its indeterminacy.[257] Both

[256] Hiliary Allen proposed specifically that the precautionary principle be adopted for financial regulation in A New Philosophy for Financial Stability Regulation (2013) for the same reasons it was proposed for environmental regulation: strong evidentiary standards require a level of certainty that is unattainable in most cases until after the damage has already been done.

[257] See Sunstein, "Beyond the Precautionary Principle" for a short version of this critique and the full-length project in *Laws of Fear: Beyond the Precautionary Principle.*

action and inaction can generate risks, so a strict precautionary standard would prohibit virtually *any* policy choices. Takeaways from the Precaution Principle suggest that we should reduce costs to participation and provide incentives for global cooperation, and that given its 'provisional character,' we should treat the precautionary principle as "a call for continuous learning and policy updating."[258]

The two sides in this debate both get something wrong about risks: on one hand, the stringent precautionary principlists seem to not account for the tradeoffs that ought to exist for risks taken. While we ought not be taking *any* risks without compensating rewards, everything else should be assessed in terms of risk versus reward judgments, ones that take into account the overall facts about the risk landscape in which any individual decision is being made. Risk aversion itself can have negative consequences. On the other hand, it is also reductive to see us as having our hands tied until the day some bad outcome occurs. The idea is that much progress can be made by getting better at deliberating about risk and the associated tradeoffs.

Regulatory Determinism

Debates about financial regulation are shaped by what I will call regulatory determinism: the assumption that regulations ought to be static, there is one 'right answer,' and that the regulator's primary objective is to discern what the correct regulatory form and requirements ought to be. On this view, regulation is successful insofar as it identifies the right number– capital requirements, leverage ratios, disclosure thresholds, and so on, which, once enacted into statute, should persist indefinitely.

[258] Jonathan Wiener, "Precaution and Climate Change" in The Oxford Handbook of International Climate Change Law.

This attitude is one well-suited to markets for goods and services, and many regulatory prohibitions in both markets for products and services and financial markets ought to exist in perpetuity. The rule that imported goods should not be produced using slave labor is not something we expect to need revisiting in response to new technology or changing market conditions.[259] Consumer protection law too exhibits a high degree of stability, with only the need to perhaps adopt new specific rules for new technologies. While new technologies require regulatory adaptation, this adaptation is largely translational rather than substantively changing the principles in play. Whether the product for sale is snake oil claiming to cure mumps or a digital product that obscures fees, the regulatory options are relatively similar: we can opt to allow sale, prohibit sale, or allow sale subject to disclosure, approval, or oversight to varying degrees. What changes over time are the forms deception might take, not the governing principle, which in each case is something like 'prevent deception.'

The success of this model can create the impression that by default, good regulation is that which can be fixed in advance. But finance is a domain in which static rules are a poor fit if we accept that finance has a function broader than being an arena for human behavior. On the other hand, readers will likely be acquainted with other kinds of policies that are explicitly dynamic in nature. Many institutional rules are designed to adjust with changing environmental conditions: Only a fixed percentile of test takers will pass certain standardized exams. Home insurance premiums rise as replacement costs (and other costs) increase. Rent stabilization regimes often index allowable rent increases to the Consumer Price Index (CPI). Congestion pricing varies tolls in response to traffic patterns. Title I funding adjusts with poverty rates and census data. And so on.

[259] 19 U.S.C. § 1307. (Sec. 307 of the Tariff Act of 1930).

These policies differ in their normative aims (constrain the number of candidates considered exemplary; keep insurance companies from going bankrupt; allow landlords to continue operating; reduce traffic congestion; educate students in low-income communities, respectively). But they share a structural feature in that they govern behavior in ways that respond to changes in the environment rather than remaining static. The motivation is straightforward: if rent stabilization were not indexed to inflation, it would gradually fail in some manner. This would occur not because the policy was poorly designed, but because the world to which it was addressed had changed.

This suggests an important distinction: Some forms of flexibility are rule-bound and parameterized. The relevant variables are identified in advance, and the relationship between inputs and outputs is fixed. We know, before any change occurs, how we expect a policy to respond. Other forms of flexibility introduce an additional layer of mediation. Economic indicators like unemployment or inflation are first observed, and then interpreted by designated expert agents (i.e., central bankers). These experts exercise judgement in determining how policy outputs should change (e.g., the Federal Funds interest rates) and the flexibility at play here is codified so as not to generate surprising (e.g. volatility-inducing) outcomes.

Regulatory determinism treats financial regulation as if it belongs to the first category and implies that the task at hand is identifying the correct static rules to govern financial markets. But financial markets are not merely sites of exchange, that are systems for allocating risk, often in ways that can change the conditions under which the regulatory regime operates. Static rules in this kind of environment risks failure in achieving the aims of said regulation by stabilizing the wrong features of an environment which may change. I will argue that more things should behave like interests do: fluctuating as needed and suggested by changes in the environment.

But this degree of determinism is not universal, and there are other arenas of human behavior where we have accepted flexible or variable standards or regulations that are contingent on certain environmental features (which themselves are subject to change). Only the top percentile scorers of exams will pass them. Home insurance, which is a controversial number in a place like Oklahoma, goes up when prices rise.[260] Rent stabilization in cities that have it is usually tied to CPI. Congestion pricing adjusts a bridge's tolls based on traffic patterns. Title I funding for schools adjusts with poverty rates and census data. Some of these are market mechanisms that act like self-regulating systems, or, like congestion pricing, are regulations based on a hypothetical self-regulating system. Many of them have encoded variation in their metrics because of the purpose the thing in question is trying to serve: exams are often trying to identify the best test-takers, relative to each test-taker's competition. The purpose of home insurance is for homeowners to have an insurer to give them money in case their home incurs damages, something an insurance company can only stay solvent doing by scaling insurance premiums relative to their underlying value as opposed to all insurers paying a flat fee. What they all have in common is that they are all examples of institutional policies or rules governing behavior which respond dynamically to conditions rather than being static. The idea behind these sorts of policies, to take one example, is that if we do not set rent stabilization to respond to changes in CPI, we might end up failing to meet our objective because our intervention/solution is insufficient for reasons relating to a changing environment. There is this kind of 'flexibility' in which the cause of the flexibility is quantified or identified: before any change occurs, we know what will happen to the output given some change in the input. Another mechanism of flexibility

[260] Oklahoma had the 3rd highest home insurance rates in 2025. This is partly due to intrinsic risks associated with our variable climate and tendency for tornadoes and ice storms, and also because it is the fall guy for other states where insurance companies aren't allowed to raise rates more than a nominal percentage YoY.

can add in another layer: the thing that changes (data on rates of unemployment and inflation, for example) is then 'processed' by some set of agents, who in some fashion decide how to change the output. The meddling done by central bankers can be justified in virtue of a recognition of risk's role in maintaining or expanding the size of the economy. We all benefit from distribution of risk and (disproportionately) by growth. We are also all (disproportionately) harmed by volatility and shrinkage of the economy.

Still, the downstream problem will be: which speculation should we keep, and which should we prevent? The criticism that finance should only do the good things and avoid doing the bad things implies both that we can tell good from bad, and that the good could go on without the bad. In reality, both of these are flawed assumptions– the former for reasons of unintended consequences, and the latter for more basic mechanistic reasons.

As I have argued in earlier chapters, it is impossible to implement a 'no speculation' regulation or otherwise regulate financial markets in such a way as to clearly delineate between actors only good 'good things' in the 'real economy.' We are left with two options: either we continue with this approach despite its metaphysical limitations, much like playing a game of whack-a-mole in harm reduction, or we try to establish some kind of metric to set a ceiling. This is a common practice in banking regulation, where certain banks are held to specific reporting standards based on metrics like assets under management. The logic is that, by drawing a line, however arbitrary, we can separate institutions into tiers and apply different expectations to each. 'Big banks' have a higher regulatory compliance burden than small banks. Of course, mistakes can be made in where we draw this line or in the standards we impose on each tier. Nonetheless, this remains our best option, making it crucial that the lines are drawn correctly.

Possible Regulatory Changes: Cryptocurrencies

Now I will demonstrate using two examples discussed earlier, how my approach would suggest we regulate two issues. Cryptocurrency is an example of the pacing problem[261] where technological innovation outpaces the ability for legal and regulatory oversight to respond.

So, let us consider our options: One approach is to regulate cryptocurrencies as securities. In many cases, crypto tokens plausibly satisfy the Howey test: they involve i) an investment of money, ii) in a common enterprise, iii) with an expectation of profits, iv) derived primarily from the efforts of others.[262] Treating them as securities would impose disclosure and compliance obligations and would likely be an improvement (from a harm-reduction standpoint) over the current lack of enforcement. However, regulating cryptocurrencies also gives them legitimacy. It places them in the same category as publicly traded firms. A more selective version of this would include regulating some cryptocurrencies as securities (conferring those legitimacy) while declining to enforce securities regulation for other cryptocurrencies.[263] In either of these cases, though, it seems even less plausible to say that the cryptocurrency is functioning as a currency at all.[264]

Some proponents might argue there is no necessary tension between something is being both an asset and a currency, and gold seems to be both an asset and something that at least historically served as a currency. The deeper issue here is not the metaphysical status of

---

[261] See Marchant et. al. *The Growing Gap Between Emerging Technologies and Legal-Ethical Oversight.* 2011 which is an edited volume, definition and outline of the pacing problem in Gary E. Marchant's chapter and Joseph R. Herkert chapter.
[262] Henderson, M. Todd, and Max Raskin. "A regulatory classification of digital assets: toward an operational Howey test for cryptocurrencies, ICOs, and other digital assets." *Colum. Bus. L. Rev.* (2019): 443.
[263] "SEC Announces Dismissal of Civil Enforcement Action Against Coinbase." SEC. 2025.
[264] This may not be a problem for those who support crypto because of its usefulness to people in situations like those experienced by Afghan women. For people like them crypto represents a way for them to store the value earned by their labor while existing in an environment which does not make such a storage of value readily available. In these cases, this seems to be of use for them even if the value storage mechanism is a 'volatile asset' rather than a 'liquid and stable currency.'

something as asset or currency, but rather that while assets can fall along a wide spectrum of volatility (with some assets maintaining stable prices over time and other moving dramatically) without that undermining their status *qua* asset, currency must necessarily be stable to maintain its existence *qua* currency. Gold and silver were once currencies and remain assets whose prices fluctuate, but not with the kind of persistent volatility characteristic of cryptocurrencies. For something to be the thing one is paid in exchange for labor and used at predictable intervals to purchase goods and services, it must be, in some sense, boring.

Boringness matters because currency systems exhibit strong equilibrium effects. We are in a stable equilibrium with respect to the small set of state-issued currencies which dominate payment, pricing, and taxation globally (and specifically in whichever jurisdiction we are subject to laws and taxation within). For individuals or firms to switch to an alternative currency, there must be a compelling reason to do so. Someone must be excited about cryptocurrency (perhaps its novelty, ideological commitments, wealth-generating potential) to use their national currency to buy it. But 'excitement' in financial markets tends to manifest in the form of price volatility, and volatility is precisely what undermines a currency's usability. This makes currency equilibria difficult to disrupt in the sense that the forces which drive adoption (that instantiate some change) will simultaneously undermine a requirement for success.

If we suppose, however, that we do reach a stable equilibrium where cryptocurrencies are widely used for transactions, while state-issued currencies remain the default for wages and taxation, this equilibrium will make us worse-off in risk mitigation terms. This resulting world would be one with parallel currencies used for different economic activities. Individuals would need to choose which currency to use for long-term planning in the form of savings and retirement. Given the unpredictability of token protocols and a demonstrated history of changing

them, these individual choices would vary widely in outcome (even assuming one dominant crypto alternative). While crypto already appears in some investment portfolios of major investment managers, this is categorically distinct from denominating one's own portfolio in a cryptocurrency or crypto-denominated asset class.

Large-scale investment in cryptocurrencies would also likely have destabilizing system-level effects. Capital allocated to crypto is, necessarily, capital not allocated elsewhere. As of 2026, cryptocurrencies are not meaningfully used to finance productive projects in the real economy. Dollars converted to Bitcoin are, in an important sense, walled off from mechanisms of productive investment.[265]

My broader project also does not necessarily justify the outright banning of cryptocurrencies, although given the constraints my project suggests, perhaps it is reduced to something only valuable for criminal enterprises and thus vulnerable to justified bans on related grounds. Again, I want to suggest that cryptocurrency activity be relegated in the minds of the general public to something more akin to a vice like smoking or gambling, as this reframing has many of the same benefits of outright bans without relying on paternalism. Undoubtedly, there should be ample consumer protection legislation and money for enforcement against outright scams. Those cryptocurrencies which are, legitimately, trying to generate their own financial system may need to be tolerated in the same way smoking is, relegated to a vice behavior not endorsed by professionals but not outlawed altogether either.

[265] This is beneficial in events of crypto collapse or sudden downturns. In the wake of FTX's downfall, despite cryptocurrencies across the board losing much of their value, there were relatively few noticeable effects from this crash in the wider economy. If we imagined a world where lots of people were investing in crypto and no one was cashing out of it (in "rug pull" situations or otherwise), we would reasonably fear potential deflation of this investment– whereby liquid capital is being funneled into the crypto financial system.

Possible Regulatory Changes: Meme Stocks

I want to present a structurally similar solution for meme stock trading as that I suggested for cryptocurrencies. If such trading behavior is better understood as a vice, then regulatory strategy should focus less on outright prohibition (which, like prohibitions on 'speculation' would be virtually impossible to enforce), and more on adding friction and changing public perception. The core concern here for my project again is not about the possibilities for individual harm (e.g. financial loss). It is rather that meme stock trading is driven largely by non-economically rational values: nostalgia, political motivations, or mere ignorance. Widespread trading like this degrades the informational function of prices and introduces volatility into markets, and this distorts price discovery, making capital allocation less tethered to underlying realities of economic activity.

Here I think a couple of regulatory solutions would be helpful. First, a return to the norms of retail trading prior to the emergence of the app Robinhood. Prior to Robinhood, there were nominal fees each time someone placed a trade.[266] Robinhood pioneered 'commission-free' trading, and this was purported to be for the benefit of consumers. Indeed, for each transaction a user makes, they are theoretically saving somewhere between ten and fifteen dollars (which was approximately the fees charged per trade prior to Robinhood's innovation). But if a lack of fees allows users to engage in behavior which is, on the whole, fiscally unsound, such savings mean little in the face of money-losing trading activity.[267]

The reinstatement of norms of having a nominal fee each time one places a trade has knock-on effects and would make day trading (trading in and out of securities over the course of

[266] This also needn't be a regulatorily instantiated rule. If we established the vice narrative, brokers should have an incentive (e.g., profit) for reinstating nominal transaction fees.
[267] An analogy: Having a realtor waive their fees on your purchase of a home is on net, not a benefit to you if you're buying a structurally unsound home which will, unbeknownst to you, end up being a money-losing investment.

a day) far less profitable. A fee of say ten or fifteen dollars per trade is unlikely to deter a customer legitimately striving to save for retirement.

Another regulatory mechanism that would make retail traders better off without stopping them from trading entirely would be to instantiate higher educational requirements for trading, at least for complex financial products like derivatives. Right now, platforms like Robinhood have a simple survey that users are asked to complete before being allowed to trade complex derivatives. Instantiating a kind of 'licensing' procedure similar to the exams that professional traders have to take would both level the playing field intellectually and deter some folks from getting involved in more complex financial products.[268]

Other Regulatory Considerations

If we understand the purpose of financial markets to be the distribution of risk for the sake of the future economic picture, the amount of risk any given society will need to take on will depend on the medium- and long-term projections of their population. For example, a low growth or shrinking population like Japan and Korea will face very different risk prospects than a high growth population like India. In fact, there are no playbooks for how asset prices are likely to respond when more people are aging and selling out of assets to fund retirement and there are fewer young people investing in them for their own retirement. For Korea and Japan, growth of the economy is secondary to managing the risks that a shrinking population and thus asset holders will have. In contrast, India will need to solve real material economic problems like building future infrastructure, houses, roads, for a growing and industrializing population.

[268] To be clear, I would not advocate for this licensing requirement for the simple trading of individual securities or index funds.

These descriptive facts about predictions of a future population and their needs are necessary to discern the normative role that finance ought to play. In each kind of situation, the worth of taking on risk for the sake of growth is very different.

Capital Requirements

One pervasive kind of regulation is capital requirement: the mandated amount of capital that must be kept inside of some institutions (e.g., banks) at all times.[269] As mentioned in Chapter II, for each dollar of deposits a bank stores for depositors, they can only deploy some fraction of that dollar for lending. This is usually done as a percentage of an institution's balance sheet. The lower required threshold is supposed to protect the bank from going bankrupt in case some of its loans default. This admits an irreducible epistemic uncertainty: if a depositor deposits $100 into the bank, it makes sense for the bank to use some amount of that $100 to issue a loan to someone looking to, let us say, start a business. But how much should it lend? To lend $100 would be to assume that the depositor will not request back any of their money until the bank receives money back from the loan. But no matter how reliable the loan may be, there is always some chance it will not be repaid, or that payments will be made late. Capital requirements are supposed to help smooth over these kinds of losses.

Capital requirement regulation requires that banks fund themselves with a minimum percentage of equity and equity-like capital[270] relative to their risk-weighted[271] assets, rather than

[269] Labonte, Marc; Scott, Andrew P. *Bank Capital Requirements: A Primer and Policy Issues*. 2023.
[270] Equity-like capital includes preferred stock and convertible bonds. The important categorical feature of these is that they are supposed to be loss-absorbing in the sense that investors know in the event of bankruptcy they are among the last to be paid out.
[271] The risk-weighing of assets is a method used by regulators to try and categorize different types of assets into levels of riskiness, which is then used to differentiate how much capital is required based on assets of different types. So, the idea here is that the less risky the kind of asset, the less capital is required on-hand (and for more risky kinds of assets, more capital is required on-hand) as a percentage of the asset's value.

relying on deposits (which can be taken out of the bank suddenly by depositors) or debt (which can fail to be repaid). These requirements scale based on the size of the bank, their risk profile[272], and 'systemic importance.'[273] Systemically important banks (or at least those predetermined to be systemically important) and large banks have higher requirements than others.

Capital requirement regulation and liquidity regulation address two distinct but related sources of bank fragility. Capital requirements are designed to reduce insolvency risk by requiring banks to fund a minimum portion of their balance sheets with equity or equity-like instruments capable of absorbing losses. By ensuring that banks maintain a certain level of capital, they have a buffer relative to their risk-weighted assets that are supposed to prevent institutions from failing if asset values fall.

Liquidity regulation, on the other hand, is designed to reduce the risk that a bank fails not because the things it has on its balance sheet become less valuable, but because it cannot meet short-term cash demands from people withdrawing (and to fund normal operations). Even a solvent bank can collapse if it lacks sufficiently liquid assets to meet a sudden surge in withdrawals or funding pressures. Liquidity rules therefore require banks to hold assets that can be quickly converted into cash with minimal loss of value.

Two core post-crisis liquidity ratios are illustrative here: The liquidity coverage ratio (LCR) requires banks to hold enough high-quality liquid assets[274] to cover projected net cash

---

[272] 'Systemically important financial institutions' is a category usually applied to large banks that are interconnected with other banks and non-bank institutions. The idea here is that if such an institution were to fail, that failure would be likely to cause cascading failures at other institutions. This determination is made prior to possible failure, and triggers additional regulation aimed at preventing the failure of such institutions (as a post-'08 reform aimed at not being in a position to need to bail out banks again).

[273] This is a term of art, and in the wake of the Silicon Valley Bank crisis of 2023, it is fair to wonder whether nearly any sufficiently connected institution would, after the fact, be deemed 'systemically important' and thus have the benefits those institutions receive from the central bank and the associated moral hazard.

[274] A critique of this policy is that banks are required to keep this buffer even in times of stress, meaning they effectively have to have a much higher buffer. See Marco Macchiavelli, "The Upside-Down World of the Liquidity Coverage Ratio: Some Unpleasant Arithmetic." Bank Policy Institute. February 11, 2026.

outflows over a 30-day period under regulator-specified stress scenarios. The ratio requires that the stock of HQLA divided by estimated stressed net outflows over 30 days be at least 100 percent. The net stable funding ratio (NFSR) requires banks to fund longer-term and less liquid assets (e.g., mortgages and other long-duration loans) with sufficiently stable sources of funding.[275] This is supposed to limit reliance on volatile short-term financing.

Taken together, these two ratios are supposed to reduce solvency- and run-risk, respectively. That these are ratios means they adjust to the underlying features of the bank. But the ratios themselves stay the same over time. These two ratios are examples of regulations that could fluctuate depending on the underlying circumstances of the societies in which they exist. For societies with less estimated medium- and long-term growth, perhaps the requisite ratios should be far above 100% (this would be a more conservative and restrictive regulation than what exists now). Perhaps for countries with a lot of predicted population growth, those ratios should be relaxed to less than 100%.[276] These are just two concrete examples, and I do not mean to suggest either that these are policies I think are bad in any particular way in their current form. Rather, it is meant to illustrate the point that our current regulatory environment is flexible and adaptive to different institutions. My approach would suggest we push this further. If the purpose of finance is diachronic– to think about and manage the future economy which will exist in the medium- and long-term, then the amount of risk we should allow banks to take on (here, we are talking about risk in terms of percentage of deposits lent out, not about the riskiness of the individual lending) should depend on future expectations. For example, a country which is

[275] Basel Committee on Banking Supervision, *Basel III: The Net Stable Funding Ratio* (Basel: Bank for International Settlements, 2014).

[276] It is possible that this kind of regulatory dynamic, if adopted around the world, would lead to 'winners' with less regulation and thus subsequent increase in financial and business activity. This would be a benefit for the sake of the economic questions raised here, but possible political consequences would need to be examined.

forecasted to grow in population should likely allow more flexibility with capital requirements than a country which is projected to decrease in population.

Limits to My Approach

My approach will not tell us what the proper orientation is in financial markets with respect to questions of international security and geopolitical strategy. There might be instances where, for reasons of national security, we wish to establish various industries and production in-house instead of abroad. And doing this may have some effect on the riskiness of production on net. It could be said that in the case of COVID-19, we saw how distributed supply chains materializes in greater 'risk' for the production of goods: if we have products for manufacturing being imported from China into the Philippines, where a widget is constructed before being shipped to the United States to be sold, we have a situation with risk complexity: In one sense, this distribution of materials and production away from final consumer product markets could be seen as risk-reducing (if something happens in the Philippines which makes them unable to construct the widget, another manufacturing partner may be found) relative to a situation where the manufacturing of a product is 'vertically integrated' with the supplier, manufacturer, and consumer market located in the same place. On the other hand, this distributed supply chain dynamic will be more sensitive to risks related to the transportation of goods (which can be disrupted by events like pandemics).

What this example demonstrates is that there are several different dimensions to risk. Although I am pointing to aspects of risk-taking that go beyond expected value metrics, we are unlikely to be able to perfectly reduce any individual investment opportunity to a one-dimensional risk assessment. The dominant single-sum account which reduces risk to dollar-denominated expected value clearly obscures morally salient distinctions between kinds of risk,

and risks that might be socially worth taking. My approach would suggest that we ought to be sensitive to the social value (or disvalue) of financial transactions which necessarily generate both risk and return. If we understand finance to be for the purpose of distributing risk and shaping the future contours of the economy, considerations beyond monetary payoff become more obviously relevant, even to individual decision-making. But it does not tell us how we ought to order those priorities or interests. Profit maximization will still be a permissible motive at the level of individual agents, but the governance of financial markets and the norms that exist within them should track the other interests and in generally the same priority-order we determine them to fall in. This should include both risks borne at the individual interpersonal level (e.g., the risk of an individual having inadequate access to food in the event of disruptions to the food supply) and those borne at the financing level for individuals (e.g. the risk of a first time homebuyer defaulting on their mortgage). In order to pick between alternatives, we will likely have to have some priorities about different kinds of risks: when faced with, let's suppose, two options with equal absolute values of risk, but risks of different kinds, which do we value more (in terms of minimization) more?

Another limit: It is obvious that a given society's economic and political systems interact, with feedback from either set of institutions and agents having the potential to impact others. There are many ways that these effects can be negative: people highlight the bad role lobbying plays in supporting the existing economic status quo at the expense of the political system. My concern is not with the permissibility of politically motivated economic behavior within markets for goods and services. Practices like consumer "dollar voting," boycotts, and other forms of politically motivated behavior expressed in markets can be valuable mechanisms for signaling preferences and as a method of exerting political accountability on economic actors. These

actions are consistent with the role of markets for products and services. The problem, on my view, arises when analogous forms of behavior motivate actions which undermine the role of prices in financial markets and can serve to misallocate risk, capital, and thus the future of markets.[277]

My approach also will not help us with those economic risks or shocks that may come as the result of political volatility. As of this writing in 2026, there are a litany of potential actions and events that fit into this category: sudden tariff announcements or reversals, threats to independent central banking, and sudden increases in the net cost of employing talented people born abroad are three contemporary examples. It is not obvious from what I have argued here that any particular policy should or will necessarily imbue more risk or volatility into the economic system than any other, but rather that any sudden change (catalyzed from the political institutions in society) will do this. In other words, the concern for the financial system's purpose, on my account, attaches virtually nothing to the substantive content of any given policy change, and only touches the degree of unpredictability and abruptness in how policy changes are introduced.[278]

[277] Take, for example, the tin price-signaling example in Friedrich, Hayek. "The use of knowledge in society." American Economic Review 35, no. 4 (1945): 519-530: The price of tin going up is supposed to signal to entrepreneurs that there's an incentive for them to increase supply or find a substitute for tin. When the price of GameStop skyrockets, as it did during the meme stock craze of 2021, it doesn't seem clear that an entrepreneur should interpret that signal as a sign that physical retail stores which sell videogames are a good business venture to get into. Meme stock behaviors incentivize companies to try and 'memeify' themselves (i.e. make themselves attractive to retail traders). This seems to induce a behavior which steers us even further from the purpose of financial markets (because those companies focusing on appealing to online traders are doing so plausibly at the expense of engaging in actual economic activity).

[278] There are various ways this certainly connects to the content of his policies and personal character, but making such connections would require further argumentation than what I provided in Chapter IV.

What my approach can help us with, however, is clarifying how the realization of such politically generated risks bears on the functioning of financial markets, like how sudden and unpredictable shocks can disrupt their capacity to reliably price and distribute risk over time.[279]

There are also goods for the sake of which taking on financial risk can be justified. In the period leading up to the 2008 financial crisis, there was a substantial policy push to increase rates of homeownership. This seems to be a straightforwardly valuable aim: expanding access to housing can serve important social purposes and make individual homeowners better off. A purely conservative regulatory attitude might suggest that we should avoid risky financial activity altogether. This is too limiting. Consider two alternative financial activities a bank might engage in with different risk profiles: issuing a very low-risk, heavily collateralized loan to a wealthy borrower with many existing liquid assets (which the bank can seize and subsequently sell in the case of default), or issuing a mortgage to a working-class household with few assets, but for whom homeownership would be welfare-enhancing.[280] If we attend to the asymmetric distribution of social benefit across these cases, we can see the contrast between the bank's preference for the lower risk loan (in terms of expected value) of the wealthy borrower and what society's preference might be. Although the working-class household with few assets is riskier in terms of expected value, I would consider it to be highly preferable from a social perspective as it serves to expand access to housing, which is a good necessary for human welfare.

Framed purely in risk terms, the first loan is highly preferable. Framed in social terms, the second loan is preferable. There are clearly genuine tradeoffs between risk and social benefit.

[279] Not, for example, because we should necessarily feel sorry for investors who lose money due to some exogenous shock.

[280] We can imagine this is someone who wants to live in this community long-term, and even if the mortgage payments are no lower than what they would otherwise pay in rent, they are building equity over time which can be transferred to their children or relatives. Although the house itself may depreciate in value, the land upon which the house was built is likely to appreciate.

If we treat risk reduction as a regulatory objective, we overlook an important dimension of collectively choosing the kinds of economic risk-taking we want to support. Determining which risks are worth taking and their limits is beyond the scope of this project, but it is properly part of normative finance and should be subject to democratic deliberation.

In the wake of the 2008 financial crisis, both policymakers and the public developed an understandable aversion to financial risk writ large. One unintended consequence of this shift, however, has been the migration of certain forms of risk away from heavily regulated publicly traded institutions and into less heavily regulated private markets.[281] Publicly traded banks are subject to extensive disclosure requirements and oversight, which make at least some of their risk exposures visible on an ongoing basis. Private equity firms, hedge funds, and other private market actors face far fewer transparency obligations. As bank regulation has tightened, banks have indeed curtailed some higher-risk activities, but many of those activities have not disappeared. Instead, they have increasingly been taken up by non-bank financial institutions.

The result is not necessarily that aggregate risk has clearly decreased, but that risk has been redistributed into parts of the financial system where it is less transparent and often less directly regulated.[282] This opacity creates a measurement and governance problem. Policymakers and market participants have a less complete picture of where vulnerabilities are. This suggests that regulation requires a clearer institutional sense of purpose. Descriptively, post-2008 reforms were likely focused on banks because banks where the epicenter of the crisis. But again, mirroring my critique of consequentialist approaches to finance, regulating the actors who precipitated one crisis sets the stage for a game of whack-a-mole crisis causality: the market will

[281] Smith and Chafkin, "What We Do in the Shadows: Will Private Credit Take Down the Economy." Bloomberg Businessweek. March 27, 2026.
[282] Levine, Matt. "Pod Shops Are the New Banks." Bloomberg. August 28, 2025.

move the behaviors from the now-heavily-regulated market participants to the less-heavily-regulated participants absent a systems-level sense of what constraints should exist for financial markets overall.

Deliberation

To say that we ought to democratically deliberate about how much risk we take on immediately suggests a follow-up question: *how* should such deliberation be conducted? This project will not fully develop an institutional design, but I can identify several principles and structural features that would distinguish this kind of deliberation from more familiar forms of democratic decision-making. As I describe the empirical evidence which paints a negative view of individual aptitude for risk-laden decision-making below, I note again the analogy to Frankfort's theory of free will: Although we may, in a given moment, be unable to avoid certain psychological biases about risk, this needn't prevent us from endorsing either our own individual behavior (or not) or that of institutions (or not).[283]

In standard models of democratic deliberation about political and social questions, it is often assumed that citizens possess (or can reasonably be expected to acquire) the basic competencies required for participation. Civic education, access to information, and protections of political rights are typically regarded as sufficient background conditions for the kind of democratic deliberation liberal democracies do for many issues of governance. While inequalities in knowledge and reasoning ability exist, democratic theory (in the political context) generally presumes that deliberation about public values and political choices does not require highly specialized technical expertise.

[283] This is another way of saying that the adage "rules for thee but not for me" is a coherent attitude to take towards the regulation and development of a normative stance about risk in the financial system.

Deliberation about risk exposure at the macro-financial level is different. A large body of work in psychology and behavioral economics suggests that human beings are systematically prone to error in reasoning about probability, uncertainty, and tail risk.[284] Individuals tend to overweight losses relative to gains, misjudge low-probability events, exhibit framing effects, and respond strongly to recent salient shocks regardless of their probability. These tendencies are not limited to everyday people without experience. Professionals in finance can exhibit these behaviors too, although domain expertise can mitigate some of these behaviors.[285] The result is that judgements about aggregate risk tolerance are especially difficult for humans to rationalize about.

The design of an appropriate deliberative process itself should likely be a subject for deliberation, since the central normative question of this project (how much aggregate financial risk should a society bear?) is not one that existing democratic institutions are well-suited to answer. One possible model would be to adapt structures similar to existing deliberative bodies like central banking committees (perhaps with a broader mandate than currently held by the Fed, or more explicit democratic accountability). We also have reason to expect that uninformed polling on acceptable levels of risk would produce highly risk-adverse outcomes, especially in the wake of realized downside risk. A well-designed deliberative process would therefore need an educational component which would ideally equip deliberators with a working understanding of uncertainty prior to their weighing in.

[284] See Savage, Sam L., and Harry M. Markowitz. *The flaw of averages: Why we underestimate risk in the face of uncertainty*. John Wiley & Sons, 2009., Kay, John, and Mervyn King. *Radical uncertainty: Decision-making beyond the numbers*. WW Norton & Company, 2020., Thaler, Richard H. *Misbehaving*.
[285] Bikhchandani, Sushil, and Sunil Sharma. "Herd behavior in financial markets." *IMF Staff papers* 47, no. 3 (2000): 279-310.

Conclusion

This dissertation has argued that prevailing approaches to financial markets treat finance as a kind of evaluative "black box": something whose outputs we judge (as good or bad, efficientb or inefficient, etc.) without adequately theorizing its internal structure. Theorists often use the normative tools developed for use on markets for products and services, and apply them to financial markets, yielding extensive work on distributional issues and individual ethics, but less than sufficient analysis of financial markets as a whole. Against this, I have developed a teleological account of financial markets that allows us to evaluate them not merely by their outcomes or by rule compliance, but by how well they fulfill their proper social function.

The central claim is that financial markets are not best understood as sites of preference satisfaction analogous to markets for goods and services. Rather, they are institutions oriented toward the future: their role is to allocate and distribute risk in ways that sustain the conditions for future production and exchange. When we fail to recognize this diachronic and risk-oriented function, we systematically misunderstand both what finance is doing and how it should be evaluated.

This mischaracterization has a negative impact on our normative landscape. Much of the philosophical and economic literature on markets focuses on the conditions and outcomes of exchange, leaving the systemic role of risk undertheorized. As a result, normative analysis tends to default either to rule-based constraints or to evaluations of welfare *ex post*. Both approaches are limited. Rule-based views risk treating financial activity as permissible so long as it complies with existing or theorized regulations, while outcome-based views struggle to assess systems whose most impactful effects occur over time and under uncertainty.

In contrast to those types of views, a teleological framework allows us to distinguish between functional and dysfunctional financial systems by asking whether they successfully allocate risk in ways that support productive enterprise and social stability. This shift also clarifies a structural difference between financial markets and markets for products and services. While markets for products and services are largely driven by preference heterogeneity (e.g., I prefer bread, you prefer money) and operate synchronically, financial markets are driven by belief heterogeneity (i.e., we both want more money, but we disagree about which investments are good ones) and operate diachronically, shaping the range of future economic possibilities. Without this distinction, we are prone to reduce questions about finance to distributive concerns rather than structural ones about the size and trajectory of the economy. Existing theorizing helps us think about how to divide the fruits of our economic system, but my approach equips us with the tools to better understand what should justify the pursuit of expanding or shrinking the economic 'pie.'

On this view, risk is a central object of collective concern. Every society occupies a position along a spectrum of risk: the lower bound is constrained by irreducible naturally occurring uncertainty, while the upper bound is effectively unbounded, as financial instruments can generate exposures far exceeding the real economy. Because the level of systemic risk affects the stability and future prospects of all members of society, determining where to locate ourselves along this spectrum (i.e., how much risk our financial system should take) is a normative and inherently political question.

This has important implications for regulation. It undermines what I have called 'regulatory determinism'— the idea that there exists a fixed set of optimal rules that can govern financial markets across time. If we accept my approach, there is reason to think that regulatory

tools such as capital requirements and leverage ratios should be understood as dynamic instruments that ought to change given changing demographic and economic conditions. Just as monetary policy adjusts interest rates in response to shifting environments, perhaps so too should financial regulation adapt to the level and distribution of risk a society is willing to bear.

More broadly, this project shows that finance should be treated as a distinct social institution. Once we consider its systemic and future-oriented effects, it seems appropriate to subject it to the kinds of normative analysis typically conducted for political institutions. Once we see finance in this way, we have new normative questions to ask.

Finally, this framework helps explain why existing normative tools have struggled to address contemporary financial phenomena. Part of the difficulty lies in a disciplinary separation: philosophers have often lacked detailed engagement with the mechanics of finance, and economists have not always framed their analyses in explicitly normative terms. By integrating a teleological account of institutions inspired by the work of Monte Johnson (on Aristotle) and Seumas Miller (on social institutions), this dissertation aims to provide the normatively minded with, minimally, enough information on how finance operates to persuade the reader that there is much more of interest than has currently been theorized about.

In this dissertation, I have aimed to provide a set of conceptual tools suited to diagnosing and evaluating the distinctive features of modern financial markets. If successful, my approach helps clarify why financial markets should exist in our society, as well as what they should look like.

Bibliography

Allen, Franklin. “A Major Transformation: The Pros and Cons of the Dodd-Frank Act.” *Knowledge at Wharton*, September 7, 2010.

Allen, Hiliary. *A New Philosophy for Financial Stability Regulation*. 2013.

Apple Inc. “Stock Price.” *Investor Relations*. Accessed May 24, 2026. https://investor.apple.com/stock-price/default.aspx.

Arnone, Gioia, Giovanni Scirè, and Enzo Bivona. “The (Mis)use of Cryptocurrencies by Criminal Organizations: A Systematic Literature Review.” *Digital Finance* (2025): 1–37.

Banner, Stuart. *Speculation: A History of the Fine Line Between Gambling and Investing*. Oxford: Oxford University Press, 2017.

Baron, Dennis. *What’s Your Pronoun?: Beyond He and She*. New York: Liveright Publishing, 2020.

Bartram, Söhnke M., Gregory W. Brown, and Frank R. Fehle. “International Evidence on Financial Derivatives Usage.” *Financial Management* 38, no. 1 (2009): 185–206. https://doi.org/10.1111/j.1755-053X.2009.01033.x

Basel Committee on Banking Supervision. *Basel III: The Net Stable Funding Ratio*. Basel: Bank for International Settlements, 2014.

Bell, Derek R. “Carbon Justice? The Case Against a Universal Right to Equal Carbon Emissions.” In *Seeking Environmental Justice*, 2008.

Bikhchandani, Sushil, and Sunil Sharma. “Herd Behavior in Financial Markets.” *IMF Staff Papers* 47, no. 3 (2000): 279–310.

Bird, M. Y. “What We Want to Be Called: Indigenous Peoples’ Perspectives on Racial and Ethnic Identity Labels.” *American Indian Quarterly* 23, no. 2 (1999): 1–21. https://doi.org/10.2307/1185964

Boatright, John R. *Finance Ethics: Critical Issues in Theory and Practice*. Hoboken, NJ: John Wiley & Sons, 2010.

Brennan, Jason, and Peter Jaworski. *Markets Without Limits: Moral Virtues and Commercial Interests*. New York: Routledge, 2022.

Bruni, Luigino, and Robert Sugden. “Reclaiming Virtue Ethics for Economics.” *Journal of Economic Perspectives* 27, no. 4 (2013): 141–164. https://doi.org/10.2307/1185964

Bueno, José Luis Cendejas. “Economics, Chrematistics, Oikos, and Polis in Aristotle and St. Thomas Aquinas.” *Journal of Philosophical Economics* 10 (2017).

Buffett, Warren. “Berkshire’s Performance vs. the S&P 500.” In *Berkshire Hathaway Inc. 2014 Annual Report*, 2015.

Camerer, Colin, and Richard H. Thaler. “Anomalies: Ultimatums, Dictators and Manners.” *Journal of Economic Perspectives* 9, no. 2 (1995): 209–219. https://doi.org/10.1257/jep.9.2.209

Chamon, Marcos D., and Eswar S. Prasad. “Why Are Saving Rates of Urban Households in China Rising?” *American Economic Journal: Macroeconomics* 2, no. 1 (2010): 93–130. https://doi.org/10.1257/mac.2.1.93

Commodities Futures Trading Commission. “Swap Execution Facilities.”

Congress.gov. “Who Regulates Whom? An Overview of the U.S. Financial Regulatory Framework.”

CFA Institute. *CFA Curriculum 2025, Level I*. 2025.

Davis, Maggie, and Dan Shepard. “State of Home Insurance: 2025.” *LendingTree*, June 9, 2025.

De Swaan, Jean-Christophe. *Seeking Virtue in Finance: Contributing to Society in a Conflicted Industry*. Cambridge: Cambridge University Press, 2020.

Diercks, Luke, and Joshua Wright. “Kalshi and the Rise of Macro Markets.” National Bureau of Economic Research, 2024.

Donaldson, Thomas. *Corporations and Morality*. Englewood Cliffs, NJ: Prentice-Hall, 1982.

Donaldson, Thomas, and Lee E. Preston. “The Stakeholder Theory of the Corporation: Concepts, Evidence, and Implications.” *Academy of Management Review* 20, no. 1 (1995): 65–91. https://doi.org/10.2307/258887

Donaldson, Thomas, and James P. Walsh. “Toward a Theory of Business.” *Research in Organizational Behavior* 35 (2015): 181–207. https://doi.org/10.1016/j.riob.2015.10.002

Donohoe, Michael P. “Financial Derivatives in Corporate Tax Avoidance: A Conceptual Perspective.” *The Journal of the American Taxation Association* 37, no. 1 (2015): 37–68. https://doi.org/10.2308/atax-50907

Ellis, B. “Pareto Optimality.” In *Encyclopedia of Philosophy and the Social Sciences*, 2013. https://doi.org/10.4135/9781452276052

Ellis, Stephen. “Market Hegemony and Economic Theory.” *Philosophy of the Social Sciences* 28, no. 4 (1998): 513–523. https://doi.org/10.1177/004839319802800404

Fama, Eugene F. “Efficient Capital Markets: A Review of Theory and Empirical Work.” *The Journal of Finance* 25, no. 2 (1970): 383–417. https://doi.org/10.1111/j.1540-6261.1970.tb00518.x

Frankfurt, Harry G. “Freedom of the Will and the Concept of a Person.” *The Journal of Philosophy* 66, no. 23 (1969): 829–839. https://doi.org/10.2307/2023833

Fraser, Nancy. “Can Society Be Commodities All the Way Down? Post-Polanyian Reflections on Capitalist Crisis.” *Economy and Society* 43, no. 4 (2014): 541–558. https://doi.org/10.1080/03085147.2014.898822

Freeman, R. Edward. “A Stakeholder Theory of the Modern Corporation.” In *The Corporation and Its Stakeholders*, 1998. https://doi.org/10.3138/9781442673496-009

Friedman, Milton, and Anna J. Schwartz. *A Monetary History of the United States, 1867–1960*. Princeton: Princeton University Press, 1963.

Fritsch, Matthias. “Climate Ethics and Intergenerational Reciprocity in Indigenous Philosophies.” In *Intercultural Philosophy and Environmental Justice Between Generations*, 2024. https://doi.org/10.1017/9781009343756.004

Froot, Kenneth A., David S. Scharfstein, and Jeremy C. Stein. “Risk Management: Coordinating Corporate Investment and Financing Policies.” *The Journal of Finance* 48, no. 5 (1993): 1629–1658. https://doi.org/10.1111/j.1540-6261.1993.tb05123.x

FTC Staff Report. *Consumer Fraud in the United States*. 2004.

Gauthier, David. Morals by Agreement. Oxford: Oxford University Press, 1986.

Geanakoplos, John. “Promises Promises.” In The Economy as an Evolving Complex System II, 285-321. Boca Raton, FL: CRC Press, 1997.Georgakopoulos, Nicholas L. The Logic of Securities Law. Cambridge: Cambridge University Press, 2017.

Gideon, Michael, and Joshua Mitchell. “Which Employers Sponsor Defined Contribution Plans?” U.S. Census Bureau, 2016.

Graham, John R., and Daniel A. Rogers. “Do Firms Hedge in Response to Tax Incentives?” *The Journal of Finance* 57, no. 2 (2002): 815–839. https://doi.org/10.1111/1540-6261.00443.

Graeber, David. *Debt: The First 5,000 Years*. Brooklyn: Melville House, 2014.

Guo, Xu, Chen Gu, Chengping Zhang, and Shenru Li. “Institutional Herding and Investor Sentiment.” *Journal of Financial Markets* 68 (2024). https://doi.org/10.1016/j.finmar.2024.100891

Hausman, Daniel M., and Michael S. McPherson. *Economic Analysis, Moral Philosophy, and Public Policy*. Cambridge: Cambridge University Press, 2016.

Hayek, Friedrich A. “The Use of Knowledge in Society.” *American Economic Review* 35, no. 4 (1945): 519–530.

Hayden, Grant M., and Stephen E. Ellis. “Law and Economics After Behavioral Economics.” *University of Kansas Law Review* 55 (2006): 629–688.

Heath, Joseph. *Morality, Competition, and the Firm: The Market Failures Approach to Business Ethics*. Oxford: Oxford University Press, 2014.

Herkert, Joseph R. “Ethical Challenges of Emerging Technologies.” In *The Growing Gap Between Emerging Technologies and Legal-Ethical Oversight*, 2011.

Henderson, M. Todd, and Max Raskin. “A Regulatory Classification of Digital Assets.” *Columbia Business Law Review* (2019): 443–500.

Hiller, Avram. “Climate Change and Individual Responsibility.” *The Monist* 94, no. 3 (2011): 349–368. https://doi.org/10.1093/monist/94.3.349

Hockett, Robert, and Aaron James. *Money from Nothing: Or, Why We Should Stop Worrying About Debt and Learn to Love the Federal Reserve*. Brooklyn: Melville House, 2020.

Innes, A. Mitchell. “The Credit Theory of Money.” *Banking Law Journal* 31, no. 2 (1914): 151–168.

Irwin, Terence. *Aristotle’s First Principles*. Oxford: Oxford University Press, 1990.

Jevons, William Stanley. *Money and the Mechanism of Exchange*. London, 1875.

Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler. “Fairness and the Assumptions of Economics.” *Journal of Business* (1986): S285–S300. https://doi.org/10.1086/296367

Katona, George. “Psychology and Consumer Economics.” *Journal of Consumer Research* 1, no. 1 (1974): 1–8. https://doi.org/10.1086/208575

Kay, John, and Mervyn King. *Radical Uncertainty: Decision-Making Beyond the Numbers*. New York: W. W. Norton & Company, 2020.

Keynes, John Maynard. *The General Theory of Employment, Interest and Money*. London: Macmillan, 1936.

Kirk, Jeff. “Subprime Auto Loans: *The Rising Menace of Wall Street's Latest Darling." Journal of Consumer & Commercial Law* 18, no. 2 (2014): 72–75.

Labonté, Marc, and Andrew P. Scott. *Bank Capital Requirements: A Primer and Policy Issues*. 2023.

Lange, Oskar. *On the Economic Theory of Socialism*. Edited by Benjamin E. Lippincott. Minneapolis: University of Minnesota Press, 1938. Reprint, New York: McGraw-Hill, 1964.

Labonté, Marc, and Andrew P. Scott. *Bank Capital Requirements: A Primer and Policy Issues*. 2023.

Levine, Matt. “Everything, Everywhere Is Securities Fraud.” *Bloomberg*, June 26, 2019.

Levine, Matt. “Pod Shops Are the New Banks.” *Bloomberg*, August 28, 2025.

Levine, Matt. “Prediction Markets Are a Thing Now.” *Bloomberg*, November 7, 2024.

Liu, Glory M. *Adam Smith’s America: How a Scottish Philosopher Became an Icon of American Capitalism*. Princeton, NJ: Princeton University Press, 2022.

Longo, Matthew, and Bernardo Zacka. “Political Theory in an Ethnographic Key.” *American Political Science Review*113, no. 4 (2019): 1066–1070. http://dx.doi.org/10.1017/s0003055419000431

Mack, Eric. “In Defense of ‘Unbridled’ Freedom of Contract.” *American Journal of Economics and Sociology* 40, no. 1 (1981): 1–15. https://doi.org/10.1111/j.1536-7150.1981.tb01366.x

Macchiavelli, Marco. “The Upside-Down World of the Liquidity Coverage Ratio: Some Unpleasant Arithmetic.” Bank Policy Institute, February 11, 2026.

Maheshwari, Kritika. “A Unificationist Approach to Wrongful Pure Risking.” *Inquiry* (2024): 1–28. https://doi.org/10.1080/0020174X.2024.2339512

Massa, Annie, and Sarah Ponczek. “Robinhood’s Addictive App Made Trading a Pandemic Pastime.” *Bloomberg*, 2020.

Mauss, Marcel. *The Gift: Expanded Edition*. Hau Books, 2016.

McConnell-Ginet, Sally. *Words Matter: Meaning and Power*. Cambridge: Cambridge University Press, 2020.

Meikle, Scott. *Aristotle’s Economic Thought*. Oxford: Oxford University Press, 1995.

Mehrling, Perry. “Financialization and Its Discontents.” *Finance and Society* 3, no. 1 (2017): 1–10.

Metrick, Andrew. “Basel III A: Regulatory History.” *Journal of Financial Crises* 1, no. 4 (2019): 45–58.

Milken Institute. “Deriving the Economic Impact of Derivatives: Growth Through Risk Management.” 2014.

Miller, Seumas. *Social Action: A Teleological Account*. Cambridge: Cambridge University Press, 2001.

Miller, Seumas. *The Moral Foundations of Social Institutions: A Philosophical Study*. Cambridge: Cambridge University Press, 2010.

Moellendorf, Darrel. “Justice and the Assignment of the Intergenerational Costs of Climate Change.” *Journal of Social Philosophy* 40, no. 2 (2009): 204-224.

Murphy, Liam, and Thomas Nagel. *The Myth of Ownership: Taxes and Justice*. Oxford: Oxford University Press, 2002.

Nozick, Robert. *Anarchy, State, and Utopia*. New York: Basic Books, 1974.

Odd Lots. “Vineer Bhansali on Losing Fed Independence as the Biggest Tail Risk Right Now.” *Bloomberg Podcast*, September 13, 2025.

Petit, Philip. “The Domination Complaint.” *Nomos* 46 (2005): 87–117.

Philippe, and William N. Goetzmann. “Global Stock Markets in the Twentieth Century.” *The Journal of Finance* 54, no. 3 (1999): 953–980. https://doi.org/10.1111/0022-1082.00133

Polanyi, Karl. *The Great Transformation: The Political and Economic Origins of Our Time*. New York: Farrar & Rinehart, 1944.

Public Law 111–203. *Dodd-Frank Wall Street Reform and Consumer Protection Act*.

Rawls, John. *A Theory of Justice*. Cambridge, MA: Harvard University Press, 1971.

Reuters. “Kalshi Sued Over Ouster of Iran Leader Prediction Market.” March 6, 2026.

Ramsay, John. “Remarks on the Volcker Rule’s Market Making Exemption.” Speech, New York City Bar Association, February 4, 2014.

S&P Dow Jones Indices LLC. “S&P 500.” *Federal Reserve Bank of St. Louis, FRED*. Accessed April 9, 2026.

Satz, Debra. *Why Some Things Should Not Be for Sale: The Moral Limits of Markets*. Oxford: Oxford University Press, 2010.

Savage, Sam L., and Harry M. Markowitz. *The Flaw of Averages: Why We Underestimate Risk in the Face of Uncertainty*. Hoboken, NJ: John Wiley & Sons, 2009.

Sen, Amartya. “Rational Fools: A Critique of the Behavioral Foundations of Economic Theory.” *Philosophy & Public Affairs* (1977): 317–344. https://www.jstor.org/stable/2264946

Sen, Amartya. *On Ethics and Economics*. Oxford: Basil Blackwell, 1987.

Sen, Amartya. *Development as Freedom*. New York: Alfred A. Knopf, 1999.

Shiller, Robert J. *Finance and the Good Society*. Princeton, NJ: Princeton University Press, 2013.

Smith, Adam. *An Inquiry into the Nature and Causes of the Wealth of Nations*. Indianapolis: Liberty Fund, 1982.

Smith, Adam. *The Theory of Moral Sentiments*. Indianapolis: Liberty Fund, 1982.

Smith, Annie, Colleen Poon, Garrett Jones, and Zainab Thawer. *Understanding the Odds: Gambling Among BC Youth Aged 12–18*. Burnaby, BC: McCreary Centre Society, 2021.

Smith, Clifford W., and René M. Stulz. "The Determinants of Firms' Hedging Policies." *Journal of Financial and Quantitative Analysis* 20, no. 4 (1985): 391–405. https://doi.org/10.2307/2330757

Stiglitz, Joseph E. *The Roaring Nineties: A New History of the World's Most Prosperous Decade*. New York: W. W. Norton & Company, 2003.

Stiglitz, Joseph E. *Freefall: America, Free Markets, and the Sinking of the World Economy*. New York: W. W. Norton & Company, 2010.

Sunstein, Cass R. "Some Costs & Benefits of Cost-Benefit Analysis." *Daedalus* 150, no. 3 (2021): 208–219.

Sunstein, Cass R. "Beyond the Precautionary Principle." *University of Pennsylvania Law Review* 151 (2002): 1003–1058.

Sunstein, Cass R. *Laws of Fear: Beyond the Precautionary Principle*. Cambridge: Cambridge University Press, 2005.

Thaler, Richard H. "Toward a Positive Theory of Consumer Choice." *Journal of Economic Behavior & Organization* 1, no. 1 (1980): 39–60. https://doi.org/10.1016/0167-2681(80)90051-7

Thaler, Richard H. *Misbehaving: The Making of Behavioral Economics*. New York: W. W. Norton & Company, 2015.

Trachtenberg, Zev. "The Necessity of Convenience: Adam Smith's Conjectural History of the Human Niche." In *The Adam Smith Review*, 166–182. London: Routledge, 2020.

Tufano, Peter. "Who Manages Risk? An Empirical Examination of Risk Management Practices in the Gold Mining Industry." *The Journal of Finance* 51, no. 4 (1996): 1097–1137. https://doi.org/10.1111/j.1540-6261.1996.tb04064.x

Voris, and Yang. "'Cryptoqueen' Sidekick Gets 20 Years for $4 Billion Ponzi." *Bloomberg*, September 12, 2023.

Walzer, Michael. *Spheres of Justice: A Defense of Pluralism and Equality*. New York: Basic Books, 1983.

Wallison, Peter. "The Dodd-Frank Act: Creative Destruction, Destroyed." *Wall Street Journal*, August 31, 2010.

Weisenthal, Joe, and Tracy Alloway. “Why Mortgage Rates Went Up After the Fed’s Big Cut.” *Bloomberg*, October 21, 2024.

Wiener, Jonathan. “Precaution and Climate Change.” In *The Oxford Handbook of International Climate Change Law*. Oxford: Oxford University Press.