

INFORMATION TO USERS

This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps.

Photographs included in the original manuscript have been reproduced xerographically in this copy. Higher quality 6" x 9" black and white photographic prints are available for any photographs or illustrations appearing in this copy for an additional charge. Contact UMI directly to order.

**Bell & Howell Information and Learning
300 North Zeeb Road, Ann Arbor, MI 48106-1346 USA**

UMI[®]
800-521-0600

UNIVERSITY OF OKLAHOMA
GRADUATE COLLEGE

MICROECONOMIC ADJUSTMENT DYNAMICS IN U.S. COAL MINING

A Dissertation
SUBMITTED TO THE GRADUATE FACULTY
in partial fulfillment of the requirements for the
degree of
Doctor of Philosophy

By
DAVID RAY MERRELL
Norman, Oklahoma
2000

UMI Number: 9962956

UMI[®]

UMI Microform 9962956

Copyright 2000 by Bell & Howell Information and Learning Company.

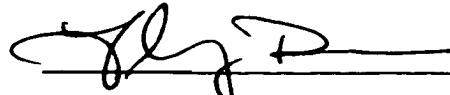
**All rights reserved. This microform edition is protected against
unauthorized copying under Title 17, United States Code.**

**Bell & Howell Information and Learning Company
300 North Zeeb Road
P.O. Box 1346
Ann Arbor, MI 48106-1346**

MICROECONOMIC ADJUSTMENT DYNAMICS IN U.S. COAL MINING

A Dissertation APPROVED FOR THE
DEPARTMENT OF ECONOMICS

BY



Bill Clark

James E. Hartigan

Scott C. L.

Zigler

Acknowledgements

I dedicate this dissertation to my mother, the late Barbara McMurry, and to my brother, the late Billy Merrell. Somehow, I know they would be proud.

By any metric that I can think of, graduate school is a painful experience—in both positive and negative ways. On the one hand, forcing oneself to think far more critically about any one issue or to think very carefully about how to express a concept in both spoken and written forms are painful yet also are beneficial. One could hardly make the case that developing those parts of one's intellect could be negative. On the other hand, graduate students they have to make certain choices about their lives that most people never have to make. They are transformed into people who are fed, by and large, by their own ambition and drive to finish school.

It is the latter dynamic that drives my thoughts in this section of my dissertation. It has been my experience that in the completion of this dissertation and even before that, the completion of my coursework, I have encountered a number of people who have made amazing contributions to my life and my career. Some of them deserve my gratitude, while others of them deserve my apology. This acknowledgement section is my (likely poor) attempt to recognize those whose contributions to my life and my study bear direct relation to my completing this work. I begin with the thanks.

First, I would like to thank my advisor and my friend, Timothy Dunne. Tim has the patience of Job and perhaps that is what makes him a very gifted

teacher. Tim has this way of taking an oftentimes arcane and faceless science and recasting what we observe in such a way that it seems real. Aside from academics, Tim has become my friend. He got me my first job, he laughs at my jokes, and he provides me with a model of what I would like to become.

Next, I thank my late grandfather, Robert Jett McMurry. Without a doubt, he is the wisest man that I have ever met. I recall back when I was about to leave Hollis, OK for West Point. He has asked me in a phone conversation what it was that I wanted to be my major. I said political science. His only response was, “What on earth would you ever do with that?” He then suggested economics. The suggestion seems to have stuck and at least in my mind, was one of the better pieces of advice that I have ever taken.

When I got to college, I waned in my commitment to be an economics major. All of my friends were engineering students, and I found their work to be more interesting than mine. At that point, I had only completed my first year of college. At about the same time that I nearly switched to become a systems engineering major, I enrolled in my first economics course: Principles of Macroeconomics. I had no idea what I was in for. The professor of this course opened my eyes to the most amazing view of the world. That you actually could couch a complex system of human behavior in terms of a few axioms and then explain how aggregate fluctuations are little more than people acting out those axioms stunned me. I was sold on being an economics major from the beginning of that course. So, I thank Will Clark for being the excellent teacher that he is and

for being the starting point of my academic career. Will taught me how to think about problems from an economic viewpoint, and later on when I started graduate school, he taught how to teach as well.

After my second year of college, I was still sold on the idea of being an economics major. However, I was sold not only because of the intellectual foundation of economics but also because economics majors always seemed to score better on LSAT exams. I was headed for law school. In the summer of 1992, I took Intermediate Microeconomics. The professor in this course tended to be pretty quantitatively rigorous—at least more so than any of the other economics faculty I had studied under before then. He was able to project what I had learned in principles of microeconomics onto a more formal plane, and from that, he was able to demonstrate the beauty of mathematical expressions for human behavior. At the same time, he was very supportive of me in terms of helping me become a better writer and of showing me that a lot about economic behavior still is unsolved. This very gifted teacher was Jim Hartigan, and I offer him my thanks for a number of things relating to my intellectual development but also for saving me from life as a lawyer.

While I mentioned before that I owe a great deal to Tim Dunne for shaping my economics and for trusting in my abilities to recommend me for the job that I have, I also owe a great deal of gratitude to Brad Jensen—my first boss. As much as Tim has shaped my economics, Brad has shaped my ability to be an economist. By force of personality and with a good deal of patience, Brad taught

me the vital lessons that we never get as graduate students—lessons that teach you not how to do economic analysis but rather how to work as an economist. From Brad, I learned management skills, marketing skills, and skills relating to how to interact with others in complex institutional environments. These lessons have served me well, and I would like to thank Brad for teaching me the importance of them.

I would like also to thank Shawn Klimek. Shawn and I met when I began working at Carnegie Mellon University in March 1997. The interesting thing about Shawn when we first met was that he was the student of my advisor's advisor. We quickly became friends and colleagues and now we are coauthors and collaborators on some interesting projects—as well as co-workers at the Center for Economic Studies. I thank Shawn for his intellectually stimulating presence, for being a sounding board for many of my ideas, and for being my friend.

Finally, I would like to thank Tami Kinsey, James Wilbanks, and Scott Chaput. All three of these people are wonderful friends of mine. Graduate school would have been literally intolerable without these three. Their friendship and mutual support were a clear blessing, and I want to thank them all for, if nothing else, being my friends. Tami's sense of humor, humanity, and reality continue to be a grounding force for a guy who spends a lot of time with his head far from the earth where he lives. James and Scott fast became my best friends, and we were able to commiserate through classes and dissertations. I am fortunate to know

them and thank them for their contributions not only to my intellectual development but also to my life as a whole.

There are a number of people who have contributed to my completing this degree but in ways that they probably did not want. The first and most important of these is T. Conner. When we first met as undergraduates, we fell deeply in love, and we stayed together until right before I graduated college. However, when I decided that graduate school was the right path for me, I developed a naked ambition to finish a doctorate at all cost. This ambition led me to give up our relationship in favor of my own self-interest. He was the first casualty of my ambition, and not a day goes by that I do not think of him and what mean continues to mean to me. To him, I owe the largest apology, and I say to him in this permanent record that I am deeply sorry.

When it comes to my godparents, Ron and Cheri Remington, I am torn about which class they fall into—the thanks or the apologies. Indeed, I owe them both. I thank them for their constant support, and I apologize to them for all too often putting my career before them. Given how I grew up, I have a lot of flexibility when it comes to thinking about family, and at the same time, I also have the benefit of being able to choose my family. I think I chose well, and I further think that if my mother were still alive, she would approve of the choice I made to become family with Ron and Cheri.

That “no man is an island” has never been more clear to me than it is these days. While sometimes I want to say that all of the contributions to my doctorate

are my own, I know very well that they are not. There are a number of people who have affected me in ways that have shaped who and what I am. They should not be forgotten nor should their contributions go unrecognized. To all of them, I owe my thanks, and to some of them, I owe my apologies. To be sure, though, to all of them goes credit that they deserve.

Table of Contents

1. INTRODUCTION	1
2. PRODUCTIVITY AND ACQUISITIONS IN U.S. COAL MINING	8
2.1 Introduction	8
2.2 Acquisitions and Productivity	15
2.3 Data and Measurement	20
2.3.1 The Data	21
2.3.2 Measuring Productivity	22
2.3.3 Identifying Acquisitions	24
2.4 Empirical Analysis	25
2.4.1 Pre-Acquisition Productivity Differences	29
2.4.2 Post-Acquisition Productivity Performance	29
2.4.3 Post-Acquisition Death	31
2.5 Conclusion	35
3. IMPLICATIONS OF RESOURCE EXHAUSTION ON EXIT PATTERNS IN U.S. COAL MINING	46
3.1 Introduction	46
3.2 Theoretical Issues Regarding Exit in Coal Mining	50
3.3 Data and Measurement	53
3.3.1 The Data	53
3.3.2 Measuring Productivity	54
3.3.3 Measuring Prices	55
3.3.4 Measuring Age	56
3.4 Empirical Analysis	57
3.4.1 The Proportional Hazard Model	57
3.4.2 The Choice of Covariates	58
3.4.3 The Role of Market Conditions	64
3.4.4 The Role of Mine Heterogeneity	65
3.4.5 Market Conditions and Mine Heterogeneity	67
3.4.6 The Case of Coal Prices	68
3.5 Future Directions for this Research	70
3.5.1 Unobserved Heterogeneity	70
3.5.2 Independence of Subsequent Spells	71
3.5.3 Modeling the Participation Decision	72
3.6 Conclusion	73
4. GROSS EMPLOYMENT FLOWS IN U.S. COAL MINING	86
4.1 Introduction	86
4.2 Measurement and Data Issues	91
4.3 Annual Employment Flows	94

4.3.1	Aggregate Trends	94
4.3.2	Mine Openings and Closings	98
4.3.3	Mine Heterogeneity and Employment Growth	100
4.4	Conclusion	102
5.	CONCLUSION	113
	References	117

1. Introduction

“The fundamental impulse that keeps the capitalist engine in motion comes from the new consumers’ goods, the new methods of production and transportation, the new markets...[The process] incessantly revolutionizes from within, incessantly destroying the old one, incessantly creating a new one. This process of Creative Destruction is the essential fact about capitalism.” (Schumpeter, 1942).

It has long been recognized that very dynamic processes underlie economic systems. Until recently, though, little was known about the engines that drive the dynamics of businesses and markets. Granted, there were a number of theories that proposed different paths of evolution, but in terms of empirical evidence, little was known beyond impressions about how markets, industries, and individual businesses grew. The recent availability of large microdata sets has created some focus on trying to measure the determinants of business growth and decline, success and failure; in short, this work seeks to uncover the dynamics inherent to the creative destruction process. This dissertation contributes to the empirical evidence on the dynamics of the creative destruction process by presenting evidence on patterns of microeconomic adjustment dynamics in the U.S. coal mining industry and by examining different mechanisms available to establishments for adjusting to changing economic environments.

The creative destruction process manifests itself in the turnover and growth of establishments—where decisions as to exiting or changing employment levels (size) are based on signals received regarding an establishment’s comparative advantage in the marketplace. Productivity dynamics play an

important role since productivity serves as a valuable signal as to the comparative advantage or disadvantage held by an establishment. Increasing (declining) productivity, *ceteris paribus*, signals that a firm has a relative advantage (disadvantage) in an industry. There are a number of theoretical models that describe how productivity can be a key component to firm/establishment dynamics. Jovanovic (1982), Dunne, Roberts, and Samuelson (1989), and Hopenhayn (1992) all develop models where firms are uncertain about their initial productivity relative to other establishments in the industry. Learning about relative productivity through market experience will provide managers with valuable information regarding the need to adjust to changing economic conditions. Based on received signals, firms have a number of tools available to them to make adjustments—should they be necessary.

Declining productivity, signaling a comparative disadvantage in the marketplace, can lead to the turnover of firms (in the entry/exit sense). The idea is simple: declining productivity (profitability) increases the likelihood of exit. Baldwin (1995) finds that exiting plants are significantly less productive than extant plants. Delving a bit deeper into productivity dynamics by looking at the contribution of within establishment productivity improvements of survivors to aggregate productivity growth, Haltiwanger (1997) attributes 54% of industry productivity growth in all of U.S. manufacturing to within plant productivity gains—suggesting in part that surviving establishments enjoy higher and higher levels of productivity than exiting plants. In Israeli manufacturing, Griliches and

Regev (1992) calculate even higher contributions of within plant productivity improvements to aggregate productivity growth. In short, establishment turnover, indeed the creative destruction process itself, is inextricably linked to changes in productivity as those changes signal the relative competence of an establishment in a given industry. The productivity of extant establishments will tend to grow rapidly as they enjoy and exploit their comparative advantages, and exiting establishments likely would have seen more declines in productivity had they not made the determination that exit was optimal.¹ For a very detailed review of this relationship, among others, see Caves (1998).

An alternative form of turnover would be to offer the business (or establishment) for sale. In such a case, the creative destruction process does force turnover but not in the same sense as in the case of exit; rather, turnover occurs through changing control of an establishment from poorly performing managers or owners to presumably better performing managers or owners. The notion here is that declining productivity need not necessarily imply that there is something fundamentally wrong with an establishment and its assembled assets (a case where turnover in the exit sense would occur) but that the comparative disadvantage could be a function of a poor match between owners (or managers) and assets. Offering an establishment for sale is a mechanism by which poorly performing assets can be allowed to migrate toward more productive control. Lichtenberg and Siegel (1987) and Maksimovic and Phillips (1999) find that

¹ Acs and Audretsch (1990) present this notion in the case of innovation and small firms. Additionally, Roberts and Tybout (1996) argue that industry level productivity is very much

establishments targeted for acquisition showed significantly lower productivity levels than non-targets before the acquisition. Additionally, both studies found evidence of significant productivity growth among extant acquired establishments. These studies suggest that an important part of the creative destruction process can be found in the turnover of management and owners of establishments (viz., changes in the control of an establishment).

Finally, the creative destruction process can manifest itself in the growth of an establishment—as measured by changes in employment levels. That is, when sufficient information arrives through productivity signals, managers of an establishment can make decisions to expand or contract in size—usually by adjusting levels of employment. Baily, Hulten, and Campbell (1992), Baily, Bartlesman, and Haltiwanger (1996), and Olley and Pakes (1996) find that the reallocation of jobs from less productive establishments to more productive establishments accounts for a large fraction of productivity growth. These findings suggest that the creative destruction process extends itself beyond the realm of turnover and into the spheres of job creation and job destruction.

The work in this dissertation aims toward gaining an understanding of the microeconomic dynamics associated with the process of creative destruction. I present additional evidence on all three of the mechanisms for microeconomic adjustment discussed above. In chapter two, I study the role of acquisitions as a corrective force for deteriorating productivity. Adapting a model of employer-employee job matching (similar to Jovanovic (1979)) to explain the match

affected by the turnover in industries—at least in the developing countries that they study.

between owners and coal mines, it is argued that less productive mines are more likely to be targeted than more productive mines. That is, if productivity is an index for the quality of a match between an owner and a coal mine, then deteriorating productivity signals a comparative disadvantage relative to other owners (viz., a relatively poor match). As a mechanism to correct this potential mismatch, the current owner can offer the mine for sale or close the mine—whichever presents the higher net payoff.

In reduced form regressions, it is found that acquired mines are between 5% and 12% less productive than non-acquired mines prior to having been acquired. This evidence supports the notion that acquisitions (as they embody the turnover of owners and managers) are a mechanism to correct poor productivity performance. Further, I find that extant acquired mines show faster productivity growth than non-acquired mines after having been acquired. Again, this is evidence that the acquisition process serves to correct poor owner-mine matches. Finally, it is estimated that acquired mines are more likely to fail than non-acquired mines—a finding somewhat at odds with the first two.

Chapter three continues the theme of microeconomic adjustment dynamics by examining the determinants of coal mine failure. Relying on a theory where productivity signals comparative advantage or disadvantage and by introducing the problem of resource exhaustion, I argue that there are competing hazards that affect failure probabilities. First, productivity should work to decrease the likelihood of failure—consistent with models of firm dynamics found in Ericson

and Pakes (1995), Jovanovic (1982), and Hopenhayn (1992). On the other hand, the implication of resource exhaustion has a diametric effect on the failure probability. That is, other things equal, an older mine would have a smaller remaining coal stock than a newer mine, and given that coal is non-renewable, this should raise the likelihood of failure since every ton of coal extracted from a mine reduces expected profitability for given productivity and energy market demand conditions.

Empirically, the problem is to separate the effects of productivity and resource depletion from one another. Using proportional hazard techniques, I estimate the effects of market conditions, mine heterogeneity, and resource depletion on the time to failure for a given mine. I find that the percent change in oil prices (believed to proxy for general energy market demand conditions) has a very slight positive effect on the time to failure, and the percent change in future oil prices significantly lengthens the time to failure. Further, I find that productivity lengthens the time to failure. Finally, controlling for productivity and other mine specific heterogeneous factors, I argue that the inclusion of a variable measuring the age of a mine will estimate the effect of resource depletion. I find that older mines, in fact, have higher probabilities of failure than younger mines, *ceteris paribus*. Again, the results in this chapter examine the establishment turnover component of the creative destruction process.

Finally, chapter four is an entirely empirical exercise that compares gross employment flows in the coal mining industry to the known patterns in

manufacturing industries. There are four principle findings. First, annual employment flows in coal mining are substantially higher than in manufacturing. Second, these high rates of job creation, destruction, and reallocation are in part attributable to mine openings and closings. Third, a significant amount of job creation and destruction is attributable to temporary openings and closings—a finding very different than what is found in manufacturing. Lastly, similar to the evidence from manufacturing, job destruction is counter-cyclical with respect to business cycles and is significantly more volatile than job creation. All of these findings serve to instruct one as to the patterns of employment adjustment dynamics in the coal mining industry.

This dissertation proceeds in the following way. Chapter two discusses the relationships between acquisitions and productivity. Chapter three addresses the implications of productivity and resource exhaustion. Chapter four measures gross job creation, destruction, and reallocation and makes comparisons between the patterns in coal mining to the patterns in manufacturing. Chapter five contains some concluding remarks.

2. Productivity and Acquisitions in U.S. Coal Mining²

Abstract: This paper extends the literature on the productivity incentives for mergers and acquisitions. We develop a stochastic matching model that describes the conditions under which a coal mine will change owners. This model suggests two empirically testable hypotheses: i. acquired mines will exhibit low productivity prior to being acquired relative to non-acquired mines and ii. extant acquired mines will show post-acquisition productivity improvements over their pre-acquisition productivity levels. Using a unique micro data set on the universe of U.S. coal mines observed from 1978 to 1996, it is estimated that acquired coal mines are significantly less productive than non-acquired mines prior to having been acquired. Additionally, there is observable and significant evidence of post-acquisition productivity improvements. Finally, it is found that having been acquired positively and significantly influences the likelihood that a coal mine fails.

2.1 Introduction

Firms regularly alter their physical and financial configurations as optimal responses to changing economic conditions. Depending on the prevailing circumstances, firms can open de novo facilities or scrap existing ones. They can expand into new product lines or exit current ones. Alternatively, mergers and acquisitions are an often used method for affecting the changes in firm configurations. In the United States from 1963 to 1997, the number of completed

² This paper is a part of my Ph.D. dissertation at the University of Oklahoma. I owe a great deal of gratitude to my advisor, Timothy Dunne—and not just in terms of this research. First, I would like to thank session participants at the 1999 Southern Economic Association Meetings in New Orleans, LA. I also thank Wendy Petropoulos, Jim Hartigan, Mark Roberts, and Dan Black for a number of very helpful suggestions. I also offer my thanks to seminar participants at Carnegie Mellon University and the Center for Economic Studies for their very constructive comments and to the Carnegie Mellon Census Research Data Center for financial support. Finally, I offer my thanks to Rhys Llewellyn and Harvey Padget, both of the Mine Safety and Health Administration, for a number of very helpful conversations and for providing the data for this analysis. All conclusions here are those of the author and do not represent the opinions or official findings of the U.S. Bureau of the Census or the U.S. Mine Safety and Health Administration.

acquisition transactions ranges from a low of 1,361 in 1963 to a high of 7,800 in 1997. Additionally, the nominal value of these transactions ranges from \$11.8 billion in 1975 to \$657.1 billion in 1997; from 1970 to 1997, the value of completed mergers and acquisitions increased 1407.11%—far outpacing any price index or even the growth in the S&P 500 index over the same interval of time.³

This seemingly increasing reliance on mergers and acquisitions to affect changes in firm structure has sparked debate over the motivations for and consequences of mergers and acquisitions. Much of the early concern emphasized market power and public interest issues (Stigler, 1950). While it is likely that the desire for market power represents some small part of the motivation for mergers and acquisitions, it is unclear in general that the anticipated gains have materialized as industrial concentration had not markedly increased during the two most recent merger waves. Still, as a strategic goal, one cannot discount entirely the search for market control as representing some part of the motivation behind mergers and acquisitions.

More recently, interest has focused on the implications which merger and acquisition activities have on the relationships between managers and owners. These concerns involve what may motivate managers to acquire whole or parts of other businesses. These motivations include strengthening managerial control over financial resources by siphoning off free cash flow from dividend payouts

³Source: Mergerstat Historical Trends. See the website <http://www.mergerstat.com/mod01/mod01-04.htm>. The number of completed mergers and acquisitions represents the number of completed merger and acquisition transactions representing at least one million dollars, and the values stated are for those transactions where a price was stated.

(Jensen, 1988; Roll, 1986), empire building (Baumol, 1987; Mueller, 1969 and 1993), and management entrenchment through maximizing objectives other than owner wealth (Shleifer and Vishny, 1989; Morck, Schleifer, and Vishny, 1990; Brandenburger and Polak, 1996). Common to all of these possible motivations for mergers and acquisitions is that they represent unchecked divergences between the interests of owners and managers.⁴

All of the above potential sources of the value gains represent uncompensated transfers of wealth from one group to another, and in this way, they represent potential sources of welfare loss. However, it is possible to have gains to mergers and acquisitions that represent true value creations. Jarrell, Brickley, and Netter (1988), Jensen and Ruback (1983) and Jensen (1988) argue that since there is no significant statistical evidence of transfer effects, the sources of the gains come from productivity windfalls resulting from freeing resources from poorly performing managers. To this end, there will be an active market among management teams for the control of corporate resources (Manne, 1965; Meade, 1968; Jensen and Ruback, 1983). Acquiring firms will target less productive firms or parts of firms, acquire them, replace the management structure, and institute programs to raise productivity.

⁴ Another direction the literature has taken is to argue that the gains from mergers and acquisitions could come from unfunded transfers from implicit labor contracts. See Summers and Schleifer (1987) and Ritter and Taylor (1999). Though an interesting claim, there is no statistical evidence that this sort of effect is present. Brown and Medoff (1987) find that employment and wages actually increase in acquired plants in Michigan. Additionally, McGuckin, Nguyen, and Reznick (1995) find that employment and wages increase in acquired manufacturing plants in the food and beverage industry. These findings are inconsistent with the notion that the gains to mergers and acquisitions come from violating implicit labor contracts.

The empirical literature on the productivity incentive for mergers and acquisitions is relatively sparse. Two general approaches have been taken. The first is to examine the pre- and post-acquisition productivity performance, and the second approach is to examine what affects the likelihood that an asset experiences an ownership change.

As an example of the pre- and post-acquisition event studies literature, Lichtenberg and Siegel (1987, 1990, 1992a, 1992b) examine the relationship between productivity and ownership change using a matching model that suggests that if productivity is a measure of the goodness-of-fit between management teams and assets, then low (high) productivity implies a poor (good) fit between management and a particular manufacturing plant, and thereby the probability of experiencing an ownership change rises (declines).

Using a balanced panel of manufacturing plants observed in the Census Bureau's Longitudinal Research Database, these authors look for productivity differences between acquired and non-acquired plants.⁵ Total Factor Productivity (TFP) is assumed to capture the quality of the match between owners and assets. In reduced form regressions, they find that acquired plants are less productive prior to being acquired than non-acquired plants—which is consistent with their matching story. Additionally, their panel exhibits post-acquisition productivity gains—to the extent that plants surviving seven years after having been acquired are not statistically different in terms of productivity than non-acquired plants;

⁵ Lichtenberg and Siegel use the Wall Street Journal index to identify manufacturing plants that have undergone an acquisition or a leveraged buy-out.

prior to being acquired, these plants performed significantly worse than non-acquired plants.

More recently, Maksimovic and Phillips (1999) use a simple neoclassical model of firm organization and profit maximization to examine the productivity-acquisition nexus. Using the Census Bureau's Longitudinal Research Database for the period 1974 to 1992, they find significant productivity gains in acquired assets in U.S. manufacturing plants—especially from assets moving from peripheral divisions of the selling firm to the main division of the purchasing firm. They find also that these productivity gains are significantly higher the more productive the acquiring firm.

The second general approach in examining the productivity incentive for mergers and acquisitions is to examine what influences the likelihood of an asset changing owners. McGuckin and Nguyen (1995) examine a sample of food and beverage plants observed in the Census Bureau's Longitudinal Research Database that change owners between 1977 and 1982. In probit regressions aimed at modeling the probability that a plant changes ownership, these authors find that there is a statistically significant positive relationship between productivity and the likelihood of being acquired⁶—suggesting in part that *high* productivity plants are *more* likely to be acquired than low productivity plants.⁷

⁶ Using financial data, Ravenscraft and Scherer (1987) and Matsusaka (1993a, 1993b) find that firms involved in mergers and acquisitions are highly profitable prior to the buyout and that there were little if any financially measured gains post-merger.

⁷ McGuckin and Nguyen (1995) find similar results to Lichtenberg and Siegel when using a balanced panel of plants constructed from the Annual Survey of Manufactures; when using their unbalanced panel, however, the result is reversed to suggest that higher productivity plants are

This paper extends the literature on the productivity incentive for mergers and acquisitions. The contributions here are twofold. First, productivity differences using microdata over time are examined in order to investigate whether the productivity differences between acquired and non-acquired assets are fundamentally related to the acquisition event. Virtually all other studies rely on balanced panels or cross-sectional observations. Second, the findings of Lichtenberg and Siegel and of Maksimovic and Phillips are corroborated in that acquired coal mines are between 5.23% and 12.46% less productive than non-acquired mines prior to the acquisition, and there are significant post-acquisition productivity gains.

In the empirical analysis, a data set on the U.S. coal mining industry containing observations on the statistical universe of coal mines from 1978 to 1996 is used. The benefits of these data are threefold. First, ownership changes of coal mines are observed at a number of points in time. Thus, is it possible able to examine whether the observed productivity differences between acquired and non-acquired coal mines manifest themselves repeatedly. Second, these data are not contained in the manufacturing universe. Virtually all of the empirical studies examining the relationships between ownership changes and productivity come from manufacturing data.

Third, the U.S. coal mining industry has undergone a good deal of acquisition activity over time. Between 5.8% and 12.2% of mines are involved

more likely to be acquired. This finding is interpreted as evidence that the Lichtenberg-Siegel estimates suffer from a sample biased in favor of large plants.

either in whole company acquisitions or partial company carve-outs.⁸ This activity is a product of a number of influences—not the least of which is the decline of the steel industry in the United States—mostly borne by integrated steel mills. U.S. iron, coke, and steel companies suffered a good deal during the recession of the early 1980s. As the production of coke and pig iron declined, companies needed less coal as a factor of production and at the same time had (generally) poor cash flows. Divesting of coal divisions is a natural mechanism to correct both problems. U.S. Steel, Republic Steel, ARMCO, LTV Corporation, and others divested much of their coal properties. For example, Inland Steel sold its coal assets to Consolidation Coal in 1986.

Additionally, large oil and gas conglomerates sold many coal properties to concentrate on their “core” businesses. Houston Natural Gas sold Ziegler Coal Company to an investment group, Amoco spun off Cyprus Minerals, British Petroleum sold Old Ben Coal to Ziegler, and Eastern Gas and Fuel Associates sold its mines to Peabody—to name a few of these such transactions. Table 1 presents some selected acquisitions that occurred during the 1978 to 1992 period; to be sure, the transactions listed on Table 1 are separated into both whole company purchases and partial company “carve-outs.” For a very informative and more complete survey of these events, see The Changing Structure of the U.S. Coal Industry: An Update.

⁸ In the McGuckin-Nguyen study, about 8.4% of food and beverage plants changed ownership between 1977 and 1982.

In the next section, a stochastic matching model very similar to that used by Lichtenberg and Siegel is presented. Section 2.3 details the sources of data for the U.S. coal mining industry and also presents some interesting features of the productivity series in this industry. Section 2.4 details the empirical analysis. Section 2.5 concludes.

2.2 Acquisitions and Productivity

To organize the empirical agenda, a market search model similar to Jovanovic (1979) is adapted. This adaptation (which is very similar to the setup used by Lichtenberg and Siegel) implies that mergers and acquisitions are mechanisms to correct deteriorating productivity performance. Productivity performance provides owners with a valuable signal about the quality of the match between the owner and the property. If productivity is declining, then current owners infer that there is some intrinsic incompatibility between the owner and the coal mine. If an owner's comparative advantage with a given mine is unknown initially, then it is only through market tenure that true relative productivity is revealed. The effect is that a heterogeneous group of owners constantly re-examines the "fit" between an owner and a coal mine.

When deciding whether or not to purchase a coal mine, the purchaser has incomplete information about how well that operation can be managed, and it is reasonable to assume that purchasers are interested in maintaining control only over operations that can be managed effectively. Hence, a buyer constantly evaluates opening or acquiring decisions, and the longer a mine is operated, the

more information is gained about the quality of the match between owner and coal mine.

The process would proceed in the following way: mines and owners are matched initially. The quality of this match (assume to be indexed by productivity) varies randomly. Lower productivity provides a signal that the quality of the match between owner and mine is low. Further, lower productivity implies that the mine would be more likely to change owners—representing the desire of an owner to maintain control over operations that can be operated effectively. If some lower bound of productivity is reached, a current owner will divest or close any mine that cannot be operated effectively. A mine is sold or closed, and the same sort of constant evaluation and re-evaluation of the comparative advantage of operating a coal mine ensues with the new owner(s).

The theoretical considerations surrounding the merger and acquisition process can be expressed formally using simple stochastic dynamic programming arguments. The problem is twofold: to describe the decision process of the current owners and to describe the decision process of a potential purchaser of a mine, given that it is offered for sale. First assume that productivity evolves according to the following stochastic process:⁹

$$(1) \ x(t) = \alpha + \sigma z(t) \quad \forall t > 0$$

where α and σ are constants, and $\sigma > 0$. $z(t)$ is a standard Wiener process with time independent increments. Assume that σ is the same for each owner-mine

match and that in general α , which is learned over time, differs across owner-mine matches. In this way, α can be interpreted as an index of the quality of the match between the owner of a mine and the mine itself. High realizations of α denote relatively good match between owner and mine, while a low realization of α represents a relatively poor match. Let α be normally distributed and assume that changing owners involves drawing a new value of α from the distribution where successive draws are independent.¹⁰

Firms maximize the expectation of net revenues discounted by the rate, ρ . Let $\pi(x;u,t)$ denote the net revenues as a function of the random state variable x and a vector of exogenous parameters, u . Assume that $\pi(\cdot)$ increases in x and that x and x' (where $x'=x+dx$) are positively serially correlated such that x is first-order stochastically dominated by x' . Let $\Lambda(x'|x)$ represent the conditional cumulative density function of x . One should be clear that all heterogeneity is driven by different realizations of the productivity state variable, x , which in turn is a function of the realization of the goodness-of-fit between an owner and a mine, α .

Current owners compare the expected value of continuing control over a mine versus the expected value of the payoff from selling or scrapping a mine;

⁹ Dixit and Pindyck (1994) use a very similar model throughout their text. For a detailed discussion on the properties of these sorts of models, the reader is referred to that text.

¹⁰ Another possibility is that there could exist “bad” mines—mines that are located in places that are difficult to mine, that are plagued with unionization problems, etc. In cases such as this, there would be serial correlation among the draws on α . Although this could be a very real possibility, it does not present any implications for the empirical agenda below since all of the estimates are from reduced-form regressions.

denote the $\Omega(x)$ as the payoff from selling a mine and θ as the scrap value of a mine. If $\Omega(x) > \theta$, then a current owner who does not desire to continue with a mine would sell the mine to another owner, but if $\theta > \Omega(x)$, then the current owner would find it better to close a mine and recoup its exogenously determined scrap value. Assume that $\Omega(x)$ is known to all firms.

Formally, current owners make the intertemporal optimization calculation suggested by the following Bellman Equation (recalling that x is a function of the goodness-of fit parameter (α):

$$(2) \quad V(x) = \max\{\theta, \Omega(x), \pi(x; u, t) + (1 + \rho)^{-1} E[V(x'|x)]\}$$

where $x' = x + dx$.¹¹ This problem is essentially one of an optimal stopping calculation in the sense that an owner decides when to cease operating and/or owning a particular coal mine. The solution techniques to this class of problems are well known, and for ease of exposition, only the relevant decisions are discussed here. Suppose there is a single-valued threshold level, x^* , which demarcates the continuation region in (x, t) space from the stopping region. Realizations of $x > x^*$ will result in the present owner continuing ownership, while values of $x < x^*$ will result in the divestiture or closure of the mine. It is clear that

¹¹ Making the resale value of a mine a function of its productivity requires two additional technical assumptions. First, there must be a value-matching property to the boundary condition; that is, in the stopping region, we have $V(\cdot) = \Omega(x)$. By continuity, we can impose $V(x^*; \cdot) = \Omega(x^*)$ where the function $x^*(\cdot)$ represents a free-boundary. However, this formulation also implies curvature in $\Omega(x)$ —suggesting a potential continuity problem at the boundary. To avoid this problem, we assume that $V(\cdot)$ and $\Omega(x)$ meet tangentially at the boundary: $\partial V(x^*(t), t) / \partial x = \partial \Omega(x^*(t), t) / \partial x$ —which is known as the high-order contact property. Dixit and Pindyck (1994) have a detailed discussion of these properties.

the present owner will continue ownership and operation if the maximum is attained at the third argument of equation (2) (when $\Omega(x) > 0$); that is, if

$$(3) \pi(x; u, t) + (1 + \rho)^{-1} \int V(x') d\Lambda(x'|x) > \Omega(x)$$

is true. Optimal stopping occurs if the opposite inequality holds—which is to say that the maximum of (2) is obtained at either of the first two arguments.¹²

This suggests that low realizations of the random productivity state variable, to the extent that those low realizations are manifest in lower profitability, lead to the divestiture or scrapping of a mine. Here is our first empirically testable hypothesis: ownership change and exit are negatively related to productivity. Hence, taken to the data, one should observe that mines changing owners (or exiting) are less productive than those that do not.

The potential owner's problem also is simple. Keep in mind that the new owner takes a draw on α which is independent of previous draws, and now the evolution of the random state variable, x , begins anew. So, the problem for the new owner when deciding whether to purchase a mine is to compare the expected profitability with the sale price of the mine. Formally, if the expected profitability of the new owner is at least as great as the sale price of the mine, then the potential owner will purchase the mine; that is, if

$$(4) V(x) = \max \left\{ (1 + \rho)^{-1} E[V(x')] \right\} \geq \Omega(x) + c$$

¹² It is simple to show that x^* exists uniquely. For a very simple and intuitive proof, see Dixit and Pindyck (1994).

where c is a parametrically determined (possibly trivial) constant representing sunk transactions costs, etc, then the mine will be purchased. Though it does not follow immediately from (4), mine acquisitions ought to result in productivity improvements for those mines. Since productivity is assumed to be randomly distributed, the expected value of a new match is higher on average than the realized value of old matches—given that subsequent owners draw from the same distribution as previous owners.

Two testable implications arise from these theoretical considerations. First, poor matches induce ownership changes. Deteriorating productivity at a mine indicates that current owners possess a comparative disadvantage with that mine relative to other current owners, and the owner likely will divest of or close that mine. Hence, one ought to see that acquired mines have lower productivity prior to the acquisition than non-acquired mines. Second, changes in ownership should result in productivity improvements over pre-acquisition levels, other things equal. This prediction reflects the notion that the expected value of a new match is on average higher than the realized value of an old match.

2.3 Data and Measurement

This section outlines the data used to classify acquisitions in the coal mining industry as well as the data used to measure mine productivity. Additionally, some details are given that describe how productivity is measured and how productivity differs across a number of important dimensions.

2.3.1 The Data

The data used in this analysis come from the Mine Safety and Health Administration of the U.S. Department of Labor. These data contain the statistical universe of coal mines in a year and are collected under the regulatory and oversight authority of the Mine Safety and Health Administration. Among other things, these data contain information on employment, hours, production, the number of injuries at a mine, and certain descriptive/classificatory information for each mine. A mine is tracked using a unique mine identification number that allows intertemporal linkages of mine observations.

For present purposes, a sample of mines observed from 1978 to 1996 is used. Each mine must have a classification code indicating that it was active in a year and must have had positive employment, hours, and production; additionally, coal processing facilities and coal contractors are not included. This leaves a large number of coal mines in each sample year. This industry has undergone a number of very unique adjustments over time—some of which are detailed in Figure 1. Figure 1 documents the patterns of mine employment, production, and hours over time. Production has increased tremendously over the sample period. In 1978, the industry produced just around 600 million short tons of coal, and at the end of the sample in 1996, industry production was just under 1 billion short tons—a 98% increase in production over the 1978 level. One interesting aspect to this increase is that it happened while there was a general decline in the number of workers employed and hours worked; this equates to large gains in labor productivity at the industry level.

2.3.2 Measuring Productivity

Productivity is measured for each active mine in the industry for each year of the sample. Because of data limitations regarding the employment of non-labor factors, only labor productivity is observed—which is measured as short tons of coal produced per worker hour. Admittedly, this measure of productivity lacks the completeness of broader multi-factor productivity measures, but it is believed that labor productivity will serve as a good proxy for total factor productivity.

Labor represents the largest share of inputs in terms of output value. From 1948 to 1991, labor inputs accounted for approximately 40% of output value, materials about 30%, and capital and energy account for about 15% each.¹³ There has been a slight tendency for labor's share of output value to decline while there is a slight trend for material's share of output value to rise. Berndt and Ellerman (1997) document a significant labor-saving bias to technical change in the coal industry. This bias in technical change also could explain divergences between total factor productivity and labor productivity.¹⁴

Before turning to the empirical analysis, there are a few important observations to make about exogenous differences in productivity that are not necessarily related to acquisitions. First, Figure 2 shows that there are clear

¹³ I offer my thanks to Dale Jorgenson and Kevin Stiroh for making these industry aggregate data series available.

¹⁴ It should be noted that simple time series correlations between Jorgenson's total factor productivity series and the labor productivity series used here are estimated at +0.9377 ($p < 0.0001$).

differences in productivity stemming from differences in the type of mine.¹⁵ Ignoring the type of mine as an explanation for observable differences in labor productivity among coal mines could lead one to misstate the importance of ownership change to differences in labor productivity—representing an omitted variables bias. Second, Figure 2 also shows a clear upward trend in labor productivity over the sample period; this result is true both for the industry aggregate and for the mine type sub-aggregates. This fact suggests that year effects are important controls as well.

Third, Figure 3 plots the productivity series separately by geographic region.¹⁶ Generally speaking, there are three broadly defined coal producing regions: the Appalachian Region, the Interior Region, and the Western Region. Appalachian mines typically are smaller, underground, and more labor intensive. Interior mines generally are larger than the typical Appalachian mine—though smaller than the average Western Region mine. Interior mines are slightly less labor intensive and are divided between surface and underground mines. Finally, the Western Coal Region is populated by remarkably larger surface mines with thick coal seams located near the surface. Figure 3 makes clear that there are distinct productivity differences between coal producing regions, and in the productivity equation, it will be important to control for this regional effect.

¹⁵ The type of mine is thought of as representing an exogenous constraint on the type of technology used when mining coal. That is, given the geographic and geologic characteristics of mines, the type of coal extraction technique is at least partly determined. Underground mines because of their particular exogenous characteristics can be mined only in certain ways—irrespective of an owner's comparative advantage, and likewise for surface mines.

¹⁶ Joskow (1987) adds this control to his analysis of price contracts in the U.S. coal mining industry.

2.3.3 Identifying Acquisitions

Identifying acquisitions in the coal mining industry requires a second data source from the Mine Safety and Health Administration.¹⁷ The records in this file are identified with the same unique mine identification numbers mentioned above—making it possible to link acquisition indicators to the production and employment files. In addition to other information relevant to the assessment of fines and fees, this data set tracks the ownership of all coal mines by recording the beginning and ending dates of ownership regimes. Changes in ownership are indicated when there is an entry on the record listing an ending date to an ownership period. If there is a valid ending date (viz., an entry not showing a missing value code and an entry containing a real calendar value) to a regime (in year t) and a start date for a new regime (also in year t), then a mine is said to have been acquired in year t .¹⁸ Given that determination, a dichotomous variable is created indicating that a mine was acquired in that year.

2.4 Empirical Analysis

Recall that the matching model presents two broad empirical hypotheses. For convenience, they are as follows: 1) mines that are acquired should exhibit lower productivity relative to mines that are not acquired and 2) extant acquired mines should exhibit post-acquisition productivity gains. Each of these

¹⁷ Specifically, this file is the Coal Information File and is maintained by the Office of Assessments, U.S. Mine Safety and Health Administration. For a fee, the Office of Assessments will make various extracts of this file available. I offer my thanks to the Carnegie Mellon Census Research Data Center for providing the financial resources to acquire these data.

hypotheses are examined and discussed in terms of the productivity incentive for acquisitions. Additionally, the role of acquisitions in coal mine failure is examined.

2.4.1 Pre-Acquisition Productivity Differences

To examine differences in productivity prior to acquisition, total annual worker hours for all active coal mines and total annual short tons of coal produced are observed. Denote these quantities H_{it} and Q_{it} , respectively for $i=1,2,\dots,N$, $t=1,2,\dots,T$. These are combined to form an index of labor productivity: Q_{it}/H_{it} —short tons of coal per worker hour. Additionally, a mine may be acquired in period t . Define this event as a dichotomous indicator variable, x_{it} , which is governed by the following rule:

$$x_{it} = \begin{cases} 1 & \text{if mine } i \text{ is acquired in period } t \\ 0 & \text{otherwise.} \end{cases}$$

Recalling from the previous section that there are clear, observable differences in productivity attributable to the type of coal mine (i.e., underground or surface mine), time, and broadly defined coal producing regions, the following pooled ordinary least squares (OLS) model is estimated:

$$(1) \quad \ln\left(\frac{Q_{it-1}}{H_{it-1}}\right) = \alpha + \beta x_{it} + \delta m_i + \sum_{t=1}^{18} \gamma_t d_t + \sum_{s=1}^3 \phi_s r_{is} + \varepsilon_{it}$$

where x_{it} is the acquisition dummy, m_i is a dichotomous variable equally one if a mine is an underground mine, d_t are year effects, r_{is} are dichotomous variables

¹⁸ It is possible to determine the difference between changes in ownership and scrapping. Scrapped mines will not have a *valid* ending date listed in the sense that the ending date field for

with one for each of the coal producing regions, and ε_{it} is Gaussian error independent over time and across coal mines. From Section 2.2, it is expected that β ought to be negative—suggesting that acquired coal mines are less productive prior to being acquired. Additionally, δ ought to be negative representing that underground mines are less productive than surface coal mines. Finally, all of the γ_t and ϕ_s estimates ought to be negative (with the omitted classes being the last year (1996) and the Western coal region, respectively); this represents that the estimates of the remaining year and region effects are interpreted relative to the omitted class: earlier years have lower productivity than 1996, and the Interior and Appalachian mines are less productive (generally) than Western coal mines.

Column 1 of Table 2 presents the OLS estimates of this model. There are a number of things to note. First, controlling for the type of mine effect, year effects, and region specific effects, the estimates indicate that acquired coal mines are 12.46% less productive prior to being acquired than mines that were not acquired. This finding is consistent with the matching model of Section 2.2. Next, consistent with Figure 2, all of the estimates on the time dummies are negative and significant—indicating that productivity has risen almost monotonically over the entire period. The region specific effects also capture significant differences in productivity; these controls work as anticipated: relative to the Western Region mines, *ceteris paribus*, Appalachian mines are on average

these mines will contain a missing value code.

68.8% less productive while Interior mines are on average 54.3% less productive. Finally, note that the estimate on m_i is highly significant and indicates that on average underground mines are about 31% less productive than surface mines—also consistent with Figure 2.

Ellerman, Stoker, and Berndt (1998) find significant evidence that the scale of a coal mining operation is an important determinant of productivity growth in U.S. coal mining.¹⁹ Although it is unclear exactly how omitted size effects would bias the estimate of β , regressors are included to control for the size of a coal mine. Specifically, dummy variables for a mine's employment size quartile in a given year are created. The following pooled regression model is estimated by OLS:

$$(2) \quad \ln\left(\frac{Q_{it-1}}{H_{it-1}}\right) = \alpha + \beta x_{it} + \delta m_i + \sum_{t=1}^{18} \gamma_t d_t + \sum_{s=1}^3 \phi_s r_{is} + \sum_{j=1}^4 \psi_j s_{jit} + \varepsilon_{it}$$

All of the regressors are the same as before, and the s_{jit} are dummies representing a mine's employment quartile in year t . Non-singularity requires that one of these dummies be omitted, and the largest quartile is chosen; the interpretations of the ψ_j , then, are relative to the largest quartile. Column 2 of Table 2 lists the estimates of this model—again including the important type of mine effect, year effects, and region effects. Including these mine size effects does not qualitatively alter the conclusions of the base specification of Column 1. That is,

¹⁹ Bailey, Hulten, and Campbell (1992) find that in manufacturing data, the size of a plant is an important determinant of productivity growth. Jensen and McGuckin (1997) present a detailed discussion of the known empirical regularities of U.S. manufacturing microdata—including size effects.

even with these controls, it is estimated that acquired mines are 11.79% less productive than non-acquired mines. This finding also is consistent with the model in Section 2.2.

It is very likely that there are other important, but unobservable, mine idiosyncrasies that drive productivity differences—like capital intensity, union status, mine age, et cetera. To examine the importance of these omitted mine-specific characteristics, the following error-components model is estimated:

$$(3) \quad \ln\left(\frac{Q_{it-1}}{H_{it-1}}\right) = \beta x_{it} + \sum_{t=1}^{18} \gamma_t d_t + \eta_{it}$$

where η_{it} is an error term consisting of a mine-specific component and Gaussian error: $\eta_{it} = v_i + \varepsilon_{it}$. It is believed that v_i captures the mine-specific idiosyncrasies that could lead to differences in productivity but that are unobservable in practice. In this model, some observable mine characteristics are omitted since there is no time variation with which to identify them, e.g., mine type and coal producing region; they are, however, part of the mine-specific component of η_{it} . Year effects, however, can be identified and are included as controls in this model. Column 3 of Table 2 lists the estimates of this error-components model. Again, the omitted year effect is 1996, and the estimates on the year effects are interpreted relative to that year. The estimate of β , controlling for year effects and mine-specific idiosyncrasies, shows that acquired mines are 5.23% less productive than non-acquired mines. Even when controlling for mine fixed

effects and year effects, the empirical evidence is consistent with the theoretical predictions from Section 2.2.

To put all of this into perspective then, all of the estimates suggest that prior to having been acquired, acquired coal mines are between 5.23% and 12.46% less productive than non-acquired mines. Irrespective of the sets of controls that are used to capture differences in productivity that are not attributable to acquisitions, acquired mines are found to be less productive ex ante than non-acquired mines. These findings are consistent with the predictions of Section 2.2.²⁰

2.4.2 Post-Acquisition Productivity Performance

The second theoretical prediction of Section 2.2 is that acquired mines ought to exhibit post-acquisition productivity gains. This reflects the notion that the expected value of a new owner-mine match is higher than the realized value the old match. To examine this issue, the productivity growth equations of the general form below are estimated by OLS:

$$(4) \quad \% \Delta Y_{it} = \alpha + \beta x_{it} + \delta m_i + \sum_{t=1}^{18} \gamma_t d_t + \sum_{s=1}^3 \phi_s r_{is} + \varepsilon_{it}$$

$$(5) \quad \% \Delta Y_{it} = \alpha + \beta x_{it} + \delta m_i + \sum_{t=1}^{18} \gamma_t d_t + \sum_{s=1}^3 \phi_s r_{is} + \sum_{j=1}^4 \varphi_j s_{jit} + \varepsilon_{it}$$

$$(6) \quad \% \Delta Y_{it} = \beta x_{it} + \sum_{t=1}^{18} \gamma_t d_t + \eta_{it} \text{ where } \eta_{it} = \nu_i + \varepsilon_{it}$$

²⁰ The two OLS specifications also were estimated using more detailed geographic controls. That is, both models were estimated using state dummies and county dummies. In all four cases, the estimates on the change in ownership dummy increase in magnitude and are all significant at conventional levels.

where Y_{it} is labor productivity measured as short tons of coal per worker hour. All three specifications are the same as in Table 2; the difference between those models and these is that the dependent variable is average annual productivity growth between the year of the acquisition and $t+1$, $t+2$, and $t+3$. Note also that the samples for these models are different than those of Table 2 in the sense that these samples are conditioned on survival. That is, to be in the sample supporting the productivity growth equation for a one year horizon, a mine must have survived after the acquisition for at least one period. The same is true for $t+2$ ($t+3$): a mine must have survived at least two (three) periods after the acquisition.

Table 3 presents the estimates of β for all nine regressions. In every case, the effect of the acquisition on productivity growth is positive and significant. For extant mines surviving at least one year after the acquisition, acquired mines' productivity growth is between 5.8% and 6.5% higher than non-acquired mines. Even at the three-year horizon, acquired mines have productivity growth between 2.5% and 3.4% higher than non-acquired mines. This evidence supports the theoretical prediction of Section 2.2 that extant acquired mines ought to exhibit positive productivity gains.²¹

2.4.3 Post-Acquisition Death

While it appears that acquisitions work as a corrective force for extant mines that experienced deteriorating productivity (consistent with Section 2.2),

²¹ Comparing the different sample selection criteria, one observes that the estimates on the acquisition dummy decline by roughly half. In a very important paper on (among other things) the effects of sample selection bias on productivity estimates, Olley and Pakes (1996) find similar

the story does not end there. That is, acquired mines tend to have higher failure rates than non-acquired mines—a finding somewhat at odds with the predictions of Section 2.2. To address the issue of the role acquisitions play in mine closure, failure probits are estimated. One must be careful to control for effects that might influence the probability of failure that are not necessarily related to the acquisition event itself. To this end, the same controls are introduced as in the OLS regressions of Tables 2 and 3. Recall that these controls were important in determining productivity differences that were not related to acquisitions. Productivity also is included as a control to separate explicitly productivity effects from acquisition effects. Specifically, probit models of the following basic form are estimated:

$$(7) \quad \Pr\{Y_{it+1} = 1\} = f(\alpha + \beta x_{it} + \lambda p_{it} + \tau(x_{it} * p_{it}) + \delta m_i + \sum_{t=1}^{15} \gamma_t d_t + \sum_{s=1}^3 \phi_s r_{is} + \varepsilon_{it})$$

where Y_{it+1} is a dichotomous variable equal to one if a coal mine closed in year $t+1$ and zero otherwise, p_{it} is a three-year moving average of the log of labor productivity, and all other controls are the same as before.²² The estimate of τ on the interaction of the acquisition dummy and productivity gives some indication

results in the sense that they observe wild swings in the magnitudes of parameter estimates when moving from balanced panels to unbalanced panels.

²² To be sure, including a three-year moving average of labor productivity conditions the sample to those mines that had been active for three consecutive years. Mines that do not meet this criterion are excluded from the sample supporting Table 4. Additionally, the sample is limited to the years 1981 to 1995; this is because mines are not observed prior to 1978 or after 1996.

of whether productivity impacts on mine failure differently among acquired and non-acquired firms.²³

Table 4 presents the direct estimates of equation 7 as well as a number of alternative specifications.²⁴ There are a number of things to note. First, though not reported, all of the probit models of Table 4 include year and region controls. These controls work qualitatively the same as before. That is, Western mines are less likely to fail than Interior mines which are less likely to fail than Appalachian mines (the omitted class); the year controls likely capture business cycle effects and do not lend themselves to easy interpretation relative to the omitted year (1981). Second, consistent with traditional models of firm dynamics, productivity is negatively and significantly related to the probability of mine failure. Third, Table 4 shows that larger mines are less likely to fail than smaller mines—given that the omitted size class is the smallest class. Fourth, underground coal mines are significantly more likely to fail than surface mines—likely because

²³ The theoretical motivation for including this interaction comes from the idea that when a mine is acquired, the new owner takes a draw on α that is independent of previous draws, and the productivity process starts over—essentially implying that the mine is “new” at least from an information standpoint. Dunne, Roberts, and Samuelson (1989) argue that failure boundaries decline in age since older plants have more refined information about the distribution from which production cost expectations are drawn. This is an artifact of older plants having more observations on production costs than newer plants which tends to reduce the variance of the cost distribution: new information causes smaller revisions in cost expectations and hence reduces the exit threshold. As another alternative, Pakes and Ericson (1998) present a very simple example where hazard rates may rise and then fall in age; again, this suggests that there is reason to believe that age differences may have differential impacts on exit. Empirically, then, one needs to control for the interaction of the state variable and age when examining failure probabilities; the state variable (productivity) may impact differentially depending on the age of the mine (where “age” in this case is a function of having been acquired).

²⁴ The only differences between specifications 7 through 9 are the definitions of the acquisition event relative to the exit date. In specification 7, x_{it} is equal to one if a mine had been acquired in period t —recalling that the exit dummy is always dated in period $t+1$. Specification 8 has x_{it} equal to one if a mine had been acquired in period t or $t-1$ or $t-2$. Finally, specification 9 includes lags

underground mines are generally less productive than surface mines. Fifth, the estimate on the interaction term is negative but not statistically significant; this suggests that productivity has the same impact on coal mine failure irrespective of whether a mine had been acquired or not. Finally, having been acquired within one to three years prior to the failure year significantly raises the probability that a coal mine fails relative to mines that were not acquired. These findings are somewhat at odds with the predictions of Section 2.2.

To control for the presence of unobserved sources of serial correlation that might lead one to overstate the relevance of acquisitions to coal mine failure, specifications 7 through 9 were re-estimated as random effects probit models; see Butler and Moffitt (1982).²⁵ Referring to the columns in Table 4 labeled 10 through 12, this extra control does not qualitatively alter the results of the specifications that do not control for unobserved serial correlation. That is, it still is the case that acquisitions are significantly and positively related to the probability of mine failure—irrespective of the definition of the acquisition dummy. At the same time, the estimate of the serial correlation parameter (ρ) in a one-factorized multinomial probit model is significant at conventional levels indicating the presence of serial persistence among the participation patterns of

of x_{it} in order to show how the lagged effect of having been acquired influences period $t+1$ exit decisions.

²⁵ The rationale for controlling for unobserved serial correlation is that it is very unlikely that the effects of the measured regressors are independent over time. Failing to account for this serial correlation could overstate the importance of the acquisition “treatment” on the exit probability of a mine.

coal mines. Again, these findings, *prima facie*, are somewhat at odds with the theoretical predictions of Section 2.2.

That acquisition events are significantly and positively related to the probability of mine failure could be explained in a way that is not inconsistent with Section 2.2; this finding could be an artifact of data limitations in the sense that *firms* are not uniquely identifiable in the data. That is, it is likely the case that the merger decision occurs at the firm level rather than the mine level, viz., firms are targeting firms (or large parts of firms) and not individual mines. Certainly, Table 1 strongly suggests that this is true. In this scenario, a firm would buy all or part of another firm and then would operate some of the new mines and closes some others. If this were the case, then this finding would not necessarily be inconsistent with the theoretical predictions of Section 2.2 since firms would be making the same mine level participation decisions as before. Managers of firms would look at each mine owned by that firm and determine the comparative advantage of operating it; it would be the same optimal stopping problem described in Section 2.2.

Given that having been acquired significantly raises the likelihood of failure, these failure probit models also bring to mind the problem of selection bias in the post-acquisition productivity growth equations. By using the sample selection criterion of survival, it is likely the case that the estimate on the acquisition dummy is biased. Keeping this in mind, the most accurate statement that can be said about the effect of acquisitions on post-acquisition productivity

performance is that having been acquired significantly increases the productivity growth of extant mines. Beyond extant mines, it is unclear that one could make a principled statement about productivity growth. Still, referring again to Table 3, the effect on productivity growth is manifest in the shortest time interval after the acquisition—a period when selection should be less severe.

2.5 Conclusion

In this paper, the relationship between productivity and acquisition activity in the U.S. coal mining industry has been examined. Deriving from a stochastic matching model, there are two broad hypotheses that describe this relationship. First, acquired mines should exhibit lower productivity prior to having been acquired—representing an intrinsic incompatibility (poor match) between an owner and a coal mine. Second, acquired mines ought to exhibit gains in productivity after having been acquired—representing that the expected value of a new match is higher than the realized value of an old match.

It is found, consistent with this stochastic matching model, that acquired coal mines were between 5.23% and 12.46% less productive before being acquired than non-acquired coal mines. This comparative disadvantage is the impetus for the acquisition: current owners are willing to sell because of the substantially lower productivity and buyers are willing to buy in order to capture the productivity windfalls of mines which can be operated more efficiently.

Additionally, there is significant evidence of productivity improvements for acquired mines. In regressions of productivity growth on the acquisition

dummy and other controls, acquisitions are positively and significantly correlated with productivity growth. At all of the horizons examined (one, two, and three years post-acquisition), extant acquired mines have faster productivity growth than their non-acquired counterpart—between 5.6% and 6.5% faster in the year immediately after the acquisition. This evidence is consistent with the notion that acquisitions are corrective forces for poorly performing coal mines.

It also is found that having been acquired significantly and positively influences the likelihood of coal mine failure. Controlling for other factors that may contribute to mine failure (both observed and unobserved) and controlling directly for productivity, acquisition events significantly raise the probability of mine failure. This finding is somewhat at odds with the model of Section 2.2. However, it could be the case that this finding is a result of limitations in the data since only mines (and not firms) are identified uniquely.

In closing, this paper is an extension on the literature that examines the productivity incentive for mergers and acquisitions. This paper corroborates the findings of Lichtenberg and Siegel and Maksimovic and Phillips by finding that acquired mines are less productive prior to being acquired and that acquired mines exhibit persistent post-acquisition productivity gains. These findings are consistent with a stochastic matching model that suggests that acquisitions are corrective forces in the evolution of the U.S. coal mining industry—at least in the sense that acquisitions are corrections for mines exhibiting relatively poor productivity. These findings are confirmed using data outside the manufacturing

universe and with a number of acquisition events occurring at different points in time—where virtually all other work has focused on manufacturing and on cross-sectional datasets. Altogether, these findings suggest that acquisitions promote the reallocation of resources from firms less able to exploit them to firms more able to profit from them.

Table 1. Selected Whole and Partial Company Acquisitions

Whole Company Acquisitions	
Aquirer	Seller
Bow Valleys Industries, Ltd.	Coal Reserves Group
Patrick Petroleum Corp.	Belibe Coal
Sun Company, Inc.	Elk River Resources
Trafalgar Industries	Avery Coal Co.
Gulf Resources and Chemical Corp.	R. D. Baughman Coal Co.
Chevron Corp.	Pittsburgh and Midway Coal
Drummond Coal Co.	Alabama By-Products Corp.
DuPont (Consolidation Coal Co.)	Inland Steel Coal Co.
Investor Group	Ziegler Coal Co.
Arch Minerals	Diamond Shamrock Coal
AOI Coal Co.	Kitanning Coal Co.
Hanson PLC	Peabody Holding Company
Ziegler Coal Holding Co.	Franklin Coal
Ziegler Coal Holding Co.	Old Ben Coal
Drummond Coal, Inc.	Mobil Coal Producing, Inc.
Carve-Out Acquisitions	
Consolidation Coal Co.	Exxon Coal and Minerals Company
Drummond Coal Co.	ARMCO
Peabody Holding Group	Arch Minerals Corporation
Mitsubishi Corporation	Cyprus Minerals Company
AMVEST Corporation	Bethlehem Steel Corp.
Arch Minerals Corporation	Quaker State Corporation
Ashland Coal Inc.	Bethlehem Steel Corp.
A. T. Massey Coal Company, Inc.	Bethlehem Steel Corp.
Montana Coal Company	A. T. Massey Coal Company, Inc.
Great Northern Properties LP	Burlington Resources, Inc.

Source: The Changing Structure of the U.S. Coal Industry: An Update, Energy Information Administration, July 1993.

Table 2. Productivity Differences Between Acquired and Non-Acquired Coal
Mines
(Student's t)

Regressor	Base Model	Size Effects	Fixed Effects
	(1)	(2)	(3)
Intercept	-1.9185 (77.61)	1.9271 (77.86)	---
Changed Owner	-0.1246 (-10.21)	-0.1179 (-9.67)	-0.0523 (-5.37)
Underground	-0.3110 (-47.00)	-0.3380 (-49.28)	---
Year78	-0.6257 (-28.94)	-0.6270 (-29.10)	-0.2990 (-16.78)
Year79	-0.6191 (-28.46)	-0.6200 (-28.61)	-0.3179 (-18.00)
Year80	-0.5261 (-24.09)	-0.5263 (-24.19)	-0.2606 (-14.87)
Year81	-0.4934 (-22.67)	-0.4941 (-22.79)	-0.2447 (-14.07)
Year82	-0.5233 (-23.58)	-0.5241 (-23.70)	-0.2927 (-16.68)
Year83	-0.4451 (-19.91)	-0.4460 (-20.02)	-0.2374 (-13.55)
Year84	-0.3946 (-17.96)	-0.3950 (-18.04)	-0.1921 (-11.09)
Year85	-0.3908 (-17.37)	-0.3913 (-17.46)	-0.2126 (-12.15)
Year86	-0.3413 (-15.02)	-0.3416 (-15.08)	-0.1721 (-9.82)
Year87	-0.2944 (-12.87)	-0.2942 (-12.91)	-0.1277 (-7.29)
Year88	-0.2462 (-10.69)	-0.2465 (-10.75)	-0.0893 (-5.10)
Year89	-0.2277 (-9.80)	-0.2273 (-9.82)	-0.0810 (-4.63)
Year90	-0.2111 (-9.00)	-0.2105 (-9.01)	-0.0888 (-5.08)
Year91	-0.1654 (-6.95)	-0.1647 (-6.95)	-0.0728 (-4.14)
Year92	-0.1076	-0.1076	-0.0374

	(-4.43)	(-4.45)	(-2.12)
Year93	-0.0834 (-3.37)	-0.0837 (-3.39)	-0.0192 (-1.08)
Year94	-0.0502 (-1.99)	-0.0505 (-2.00)	-0.0338 (-1.92)
Year95	---	---	---
Appalachia	-0.6725 (-38.76)	-0.6567 (-36.73)	---
Interior	-0.5320 (-26.17)	-0.5346 (-26.28)	---
Western	---	---	---
First Quartile	---	-0.1009 (-10.35)	---
Second Quartile	---	-0.0121 (-1.28)	---
Third Quartile	---	0.0755 (8.04)	---
Fourth Quartile	---	---	---
N=	50,549	50,549	12,255
R ²	0.1301	0.1363	0.7240

**Table 3. The Effect of Acquisition on Productivity Growth:
Estimates of β at Various Horizons
(Student's t)**

Average Productivity Growth	Base Model (4)	Size Effects (5)	Fixed Effects (6)
One Year	0.0586 (5.41) N=36,304	0.0632 (5.80) N=36,304	0.0650 (4.81) N=8,417
Two Years	0.0367 (5.31) N=26,784	0.0401 (5.78) N=26,784	0.0285 (3.53) N=5,889
Three Years	0.0292 (5.08) N=20,403	0.0339 (5.86) N=20,403	0.0252 (3.82) N=4,450

**Table 4. Post-Acquisition Failure
(standard errors)**

Regressor	Probits			Random Effects Probits		
	(7)	(8)	(9)	(10)	(11)	(12)
Intercept	-0.5352 (0.0353)	-0.5410 (0.0356)	-0.5384 (0.0356)	-0.4956 (0.0390)	-0.5029 (0.0392)	-0.5017 (0.0391)
Changed Owners (t)	0.2517 (0.0556)	---	0.2100 (0.0566)	0.2080 (0.0599)	---	0.1855 (0.0602)
Changed Owners in Last Three Years	---	0.2250 (0.0389)	---	---	0.1930 (0.0429)	---
Changed Owners (t-1)	---	---	0.2293 (0.0583)	---	---	0.2223 (0.0618)
Changed Owners (t-2)	---	---	0.0943 (0.0585)	---	---	0.0626 (0.0620)
Productivity	-0.1309 (0.0164)	-0.1298 (0.0171)	-0.1315 (0.0170)	-0.1848 (0.0198)	-0.1804 (0.0205)	-0.1810 (0.0203)
Changed Owners * Productivity	-0.0654 (0.0624)	-0.0281 (0.0414)	-0.0490 (0.0632)	-0.0617 (0.0671)	-0.0311 (0.0453)	-0.0483 (0.0673)
Changed Owners (t-1) * Productivity	---	---	-0.0567 (0.0646)	---	---	-0.0652 (0.0686)
Changed Owners (t-2) * Productivity	---	---	0.0846 (0.0623)	---	---	0.0794 (0.0662)
Underground Mine	0.2777 (0.0208)	0.2589 (0.0211)	0.2540 (0.0211)	0.2741 (0.0251)	0.2580 (0.0252)	0.2526 (0.0252)
First Size Quartile	---	---	---	---	---	---
Second Size Quartile	-0.1491 (0.0254)	-0.1566 (0.0255)	-0.1552 (0.0255)	-0.1894 (0.0292)	-0.1938 (0.0431)	-0.1916 (0.0290)
Third Size Quartile	-0.3074 (0.0263)	-0.3163 (0.0264)	-0.3129 (0.0264)	-0.3787 (0.0313)	-0.3823 (0.0311)	-0.3781 (0.0310)
Fourth Size Quartile	-0.8615 (0.0317)	-0.8569 (0.0317)	-0.8496 (0.0318)	-0.9780 (0.0389)	-0.9685 (0.0387)	-0.9597 (0.0387)
Rho	---	---	---	0.1310 (0.0130)	0.1243 (0.0130)	0.1226 (0.0129)
N=	24,027	24,027	24,027	24,027	24,027	24,027
Pseudo R ²	0.0606	0.0620	0.0627	0.0560	0.0567	0.0573
Log Likelihood (absolute value)	11,860.1	11,942.8	11,833.8	11,789.2	11,779.3	11,772.0

Note: All of these regressions also have region and year controls.

Figure 1. Total Hours, Employment, and Production: 1978 to 1996




Figure 2. Short Tons of Coal by Type of Mine





Figure 3. Short Tons of Coal per Worker Hour by Coal Region



3. Implications of Resource Exhaustion on Exit Patterns in U.S. Coal Mining²⁶

Abstract: This paper examines the determinants of mine closure in the U.S. coal mining industry. Based on a notion of firm dynamics that incorporates energy market demand conditions, productivity differences, and resource exhaustion, Cox proportional hazard models are estimated to examine the competing nature of market and productivity effects versus resource depletion effects. These models are estimated using a unique panel data set with multiple failures containing the statistical universe of coal mines observed from 1974 to 1995. It is found, consistent with the theoretical structure, that favorable energy market demand conditions as well as productivity differences tend to lower the time to failure, while older mines tend to have shorter time to failure, ceteris paribus.

3.1 Introduction

That some businesses fail while others do not is hardly surprising. Because of a variety of reasons, shocks to factor and product markets impact differentially on establishments in any given industry. Understanding the nature and influences of business failure is an important topic both from the policy and management viewpoints. To that end, this paper returns to the issue of business failure and examines the influences of exit behavior in the U.S. coal mining industry—an industry somewhat different than those traditionally under study.

²⁶ This paper is a part of my Ph.D. dissertation at the University of Oklahoma. I owe a great deal of gratitude to my advisor, Timothy Dunne. I offer my thanks to John Engberg for some very useful conversations and to seminar participants at Carnegie Mellon University. Additionally, I would like to thank conference participants at the 2000 Winter Meetings of the Econometric Society in Boston, MA and to Mark Roberts for some very detailed comments and suggestions. I also thank Rhys Llewellyn and Harvey Padget both of the U.S. Mine Safety and Health Administration for some very helpful conversations and for providing the data for this analysis. All conclusions here are those of the author and do not necessarily represent the opinions or official findings of the U.S. Bureau of the Census or the U.S. Mine Safety and Health Administration.

The theoretical literature on firm dynamics typically uses differences in learning processes to generate observable patterns of industry behavior. These theoretical models can be classified by the type of learning that is believed to exist in an industry; generally, there are active learning models and passive learning models. First, passive learning models posit that managers are uncertain about their productivity relative to other establishments in an industry. Only through market experience are managers able to learn relative productivity. Exit occurs if managers learn that they are relatively inefficient producers. That is, if managers witness productivity declining below some threshold level, then an establishment will exit the industry. As establishments age, the likelihood of failure declines since older, extant establishments will have more refined information about their relative productivity. Models of this type can be found in Jovanovic (1982), Dunne, Roberts, and Samuelson (1989), and Hopenhayn (1992).

On the other hand, active learning models posit that while managers are initially uncertain about relative productivity, which is learned over time, they can make investments that influence an establishment's position in the industry's productivity distribution. That is, managers can invest in research and development, product development, et cetera, to influence relative productivity—the returns to which also are stochastic. In these models, the probability of failure initially may rise in age (resulting from the stochastic nature of the returns to these productivity enhancing investments) but eventually declines—representing the notion that managers have more refined information as an establishment ages.

Models of this type can be found in Ericson and Pakes (1992) and Pakes and Ericson (1998).

A common feature of either type of model is the presumption that irrespective of the type of learning process, an establishment could continue indefinitely in the industry—if conditions in product and factor markets permit. However, there are types of industries where this presumption does not necessarily fit—coal mining being a good example. That is, there are industries in which establishments *must* exit regardless of how productive they are. These establishments face exogenously determined constraints on lifetime production. Extractive industries, e.g. coal mining, logging, or oil wells have a finite amount of resource to extract, and once a mine exhausts its coal reserves or an oil well runs dry, it must exit. Further, as long as a mine or well continues to produce, then the likelihood of failing approaches certainty because of resource exhaustion. So, the first clear implication of resource fixity is that, *ceteris paribus*, failure probabilities ought to increase in age—representing the notion that every unit produced reduces expected profitability for given factor and product market conditions and for given initial endowments of coal.

Figure 1 presents Kaplan-Meier hazard estimates for the U.S. coal mining industry.²⁷ This graph points out an interesting feature about the coal mining

²⁷ More formally, these estimates are called product limit estimators of the hazard. Specifically, this hazard estimator has the following form: $\hat{h}(t) = d_t / r_t$ where d_t is the number of failures at date t (the death set), and r_t (the risk set) is the number of subjects at risk at date t immediately prior to the failures, d_t . This hazard is a step function with steps at each observed failure time; additionally, Kaplan-Meier hazards are fully non-parametric estimators of failure probabilities. See Lancaster (1990).

industry. The upward slope of the empirical hazard is a clear indication of positive duration dependence—that the failure rate generally increases in time (age); see Heckman and Singer (1984).²⁸ This finding makes sense when thinking about extractive industries where every unit produced will reduce expected future profitability, *ceteris paribus*. In contrast, Pakes and Ericson (1998), using a panel of establishments in Wisconsin in both the manufacturing and the retail sectors, find that both the mortality and hazard rates generally decline in age (time)—that those sectors tend to show negative duration dependence. These divergent findings present an interesting problem of trying to figure out the determinants of business failure in these different types of industries.

The finding of positive duration dependence will serve as the point of departure for this paper. That is, this paper seeks to uncover the determinants of business failure in the U.S. coal mining industry. It is argued that there are three effects that convolve to determine the exit decisions of coal mines. Not only are there product market differences and productivity differences (where productivity serves as a (noisy) signal regarding an underlying cost shift parameter that managers learn over time) that drive exit behavior but also there is the competing effect of resource exhaustion that impacts upon the *ability* of a mine to continue in the industry. Empirically speaking, in order to determine how these three forces combine to determine coal mine failure, it will be important to disentangle them

²⁸ One should be careful to note that the notion of age and time used here will be somewhat different later in this paper. That is, at this point, in a rather general discussion, the notions of time and age are used interchangeably. However, in the later sections of this paper, time will refer to

from one another, to separate the effects of these competing hazards. In short, it is found that controlling for energy market demand conditions, mine characteristics, and regional characteristics, older mines are more likely to fail than younger mines. This finding is interpreted as the impact of resource exhaustion dominating the effect of learning at some point in the life of the mine.

This paper proceeds in the following way. Section 3.2 of this paper discusses a theoretical framework that organizes the empirical agenda. Section 3.3 discusses the data. Section 3.4 presents the empirical analysis of exit. Section 3.5 discusses extensions and improvements to this project. Finally, Section 3.6 concludes.

3.2 Theoretical Issues Regarding Exit in Coal Mining

Most models of firm dynamics make the implicit assumption that an establishment could continue forever if factor and product market conditions permit. The idea is that if there is continual access to factor and product markets, an establishment could continue to produce indefinitely. This assumption works well when thinking about industries in say manufacturing or services. However, when it is the case that the output of an establishment is a natural resource or at least some sort of product with exogenously imposed fixity, then the life of the establishment is predetermined—though not necessarily in a strictly deterministic

the duration effects of mines and will be common across all mines. On the other hand, age will become a mine specific trait and will vary cross-sectionally over the risk set.

sense. This section discusses the determinants of coal mine dynamics—an industry characterized by exogenously determined limits to lifetime production.²⁹

Assume that there are three state variables: price, productivity, and the remaining coal stock. Price and productivity clearly are stochastic state variables, and assume that they are realized in period t and that previous realizations of both are known.³⁰ Let P_t and γ_t denote period t prices and productivity, respectively.³¹ Remaining coal stock is defined as the endowment of coal less cumulative production from the de novo period to the first previous period from date t (i.e., t_0 to $t-1$ where t_0 is the date of de novo entry and the current period is t); hence, remaining coal stock is a deterministic state variable.

The decision calculus works in the following way: a coal mine realizes price, P_t (a random state variable), and productivity, γ_t (a random state variable), and then chooses output (q_t) to maximize current and expected future profits. This is the same as most models. What makes the story here different from most models is that it also has to be determined if there is sufficient reserves remaining to supply q_t optimally. Exit behavior, then, is driven by two components. First, if the realization of P_t and/or γ_t were sufficiently low, such that current and expected future profits (net of sunk exit costs) are below some minimum level, a mine

²⁹ The concept of firm dynamics here is akin to Pindyck (1980) where a mine is able to optimize over time but without a finite terminal state.

³⁰ This is a common feature of models of firm dynamics. The only structure that needs to be imposed on this assumption is that the processes of the random state variables have the Markov property. Ericson and Pakes (1995) show that the Markov property is sufficient to generate a full set of firm and industry dynamics in the Markov Perfect sense.

³¹ In most models of firm dynamics, underlying this productivity index is a parameter that shifts the cost function up or down. Firms learn about whether their draw on this parameter shifts the

would exit. On the other hand, if sufficient reserves do not remain to supply optimal q_t (irrespective of product and factor market realizations) then the coal mine would find it optimal to exit as well. This last notion represents the fact that every unit of coal extracted previously from a mine site reduces expected profitability by reducing the amount of remaining reserves—a state variable for the next period's optimization problem. The implication of this is that older mines ought to have a higher likelihood of failure than younger mines, *ceteris paribus*.

So, three effects determine the exit process: product market conditions, learning about relative productivity, and resource exhaustion. The empirical problem to follow will need to be able to mete out each of the three effects to gain some understanding of how all three convolve to generate exit dynamics in U.S. coal mining. Nevertheless, the point to take from this section is that irrespective of how productive a mine may be or even how favorable product market conditions may be, it is a certainty case that the mine some day must fail. At different points in the life of the mine, productivity or market effects may dominate the exhaustion effect, but eventually the exhaustion effect will overtake the learning process and favorable factor and product market conditions.³²

function in either direction. In this work, however, only productivity is observed which provides a noisy signal about the underlying cost parameter.

³² To be sure, this concept of production in an extractive industry is different from what typically is found in the natural resources literature. That is, the natural resources literature focuses nearly entirely on the "optimal extraction" policy for a natural resource establishment; see Krautkraemer (1998) for a very detailed discussion of this literature. The concept here does not rely on any sort of optimal extraction policy over the estimated life of the coal mine. Rather, by allowing the coal mine to adjust to stochastic product market conditions, the coal mine is able to change each period the expected life of the mine. In period of high (low) coal prices, a mine may find it optimal in

3.3 The Data and Measurement

This section outlines the data used in the following empirical analysis. Additionally, some details are given that describe how productivity, prices, and age are measured in the empirical analysis that follows.

3.3.1 The Data

The data used in this analysis come from the Mine Safety and Health Administration of the U.S. Department of Labor. These data contain the statistical universe of coal mines in a year and are collected under the regulatory and oversight authority of the Mine Safety and Health Administration. Among other things, these data contain information on employment, hours, production, the number of injuries at a mine, and certain descriptive/classificatory information for each mine. A mine is tracked using a unique mine identification number that allows intertemporal linkages of mine observations.

For present purposes, a sample of mines observed from 1974 to 1995 is used. Each mine must have a classification code indicating that it was active in a year and must have had positive employment, hours, and production; additionally, coal processing facilities and coal contractors are excluded. Mines are selected such that the de novo year is 1974 or after; this is to cure problems of left censoring. One final sample selection criterion is that each mine must have been active for two consecutive years; this restriction allows for the inclusion of lagged productivity values in the analysis that follows. A large number of coal mines are

that period to increase (decrease) production above what an optimal extraction policy (relying on a

left in each sample year—between 371 and 2069; see Table 1. This industry has undergone a number of very unique adjustments over time—some of which are detailed in Figure 2. Figure 2 documents the patterns of mine employment, production, and hours over time. Production has increased tremendously over time. For example, in 1972, the industry produced just around 500 million short tons of coal, and at the end of the sample in 1996, industry production was just under 1 billion short tons—a 93% increase in production over the 1972 level. One interesting aspect is that this increase happened while there was a general decline in the number of workers employed and hours worked; this will equate to large gains in labor productivity at the industry level.

3.3.2 Measuring Productivity

Productivity is measured for each active mine in the industry for each year of the sample. Because of data limitations regarding the employment of non-labor factors, only labor productivity is observed—which is measured as short tons of coal produced per worker hour. Admittedly, this measure of productivity lacks the completeness of broader multi-factor productivity measures, but it is believed that labor productivity will serve as a good proxy for total factor productivity. One benefit of this measure is that it is expressed as a physical quantity of output per physical quantity of input. Most other studies rely on measures calculated from the nominal prices values of either the input or the output to generate measures of labor productivity.

given price trajectory) might dictate. The effect then is to change the estimated life of the mine.

Labor productivity should serve as a good proxy for a couple of reasons. First, labor represents the largest share of inputs in terms of output value. From 1948 to 1991, labor inputs accounted for approximately 40% of output value, materials about 30%, and capital and energy account for about 15% each.³³ Second, there has been a slight tendency for labor's share of output value to decline while there is a slight trend for material's share of output value to rise. Berndt and Ellerman (1997) document a significant labor-saving bias to technical change in the coal industry. This bias in technical change also could explain divergences between total factor productivity and labor productivity.³⁴

3.3.3 Measuring Prices³⁵

Information on real oil and coal prices comes from the Department of Energy's Energy Information Administration. Coal prices are measured at the industry level as the average annual price of coal expressed in 1992 dollars per short ton. Oil prices are measured in terms of 1992 dollars per barrel of crude. Note that the oil and coal price series chosen do not control for quality differences. That is, for coal, there is no distinction between anthracite or bituminous, and for oil, there is no direct control for Alaska North Slope, Texas, or California petroleum ranks.

³³I thank Dale Jorgenson and Kevin Stiroh for making these industry aggregate data series available.

³⁴It should be noted that simple time series correlations between Jorgenson's total factor productivity series and the labor productivity series used here are estimated at +0.9377 ($p < 0.0001$).

³⁵All of the price data were gathered from the website of the Energy Information Administration (U.S. Department of Energy). See www.eia.doe.gov/emeu/aer/contents.html and choose the appropriate energy category.

Figure 3 presents time series plots for these two series. The price of coal series begins at the coal boom and then shows general decline thereafter—showing little period-to-period variation.³⁶ On the other hand, petroleum prices show much more time variation—likely a function of a more active spot market compared to coal as well as a function of political considerations among petroleum producing nations. A simple time series correlation is estimated at about +0.48 ($p < 0.0001$). For reasons discussed below, both price series should be instructive when looking at the exit thresholds of coal mines. However, given the general contractual nature of coal markets, oil prices should portray conditions in energy markets more accurately than coal prices themselves.

3.3.4 Measuring Age

The age of a coal mine is measured as the cumulative years that a mine has been active—where active is defined above. That is, to be active in a given year, a mine must have had positive hours, production, and employment. If these criteria are satisfied in a year, a mine is given a status flag equal to one. The age of the mine, then, is the sum of all of those status flags from the year of de novo entry to year t . There are, however, mines that are in for some time, exit, and then re-enter. In these cases, the mine's age is the same for those inactive years as the last previous year that it was deemed active; that is, those inactive years have no contribution to the calculation of age. Age is simply the number of years that a

³⁶ At the height of the coal boom, real coal prices increased 122% over their levels immediately prior to the boom.

mine has positive employment, production, and hours since its date of de novo entry.

3.4 Empirical Analysis

In this section, attention is focused on what influences the probability of coal mine failure. First, the estimation of a Cox proportional hazard model is discussed. Second, the choice of covariates is discussed. Finally, the estimates of the Cox models are presented and discussed.

3.4.1 The Proportional Hazard Model

The data discussed in the previous section allow for multiple failures; that is, a mine may be observed failing more than one time throughout the sample period. In fact, the sample contains a total of 8,520 subject mines (the risk set) and 8,961 (the death set) observed failures. It is assumed that subsequent failures are independent. Finally, a “spell” is defined as “time to failure.” That is, a mine is observed as active and then inactive; the time between these two determinations is the “spell.”³⁷

It is assumed that the hazard function for a given coal mine is parameterized as the following proportional hazard:

$$(1) \quad \lambda(t; x) = \lambda_0(t) \exp(x\beta)$$

where $\lambda_0(t)$ is an arbitrary and unspecified base-line hazard function. x is a vector of m observable covariates (both time variant and time invariant), β is a vector of

³⁷To address the problem of left censoring, the sample is selected such that all mines enter de novo in 1974 or after.

m parameter estimates, $t \in T$ is the time index with T as failure time. This model will have a density function of the following form:

$$(2) \quad f(t; x) = \lambda(t; x) \exp \left[- \int_0^t \lambda_0(u) \exp(x\beta) du \right]$$

where the last term is the survivor function associated with the proportional hazard model of equation (1). Kalbfleisch and Prentice (1980) show that the corresponding partial likelihood function takes the following form:

$$(3) \quad L = \prod_{i=1}^n \left[\frac{\sum_{k \in D_i} \exp(x_k \beta)}{\left[\sum_{j \in R_i} \exp(x_j \beta) \right]^{d_i}} \right]$$

where i indexes ordered failure times $t_{(i)} \forall i=(1,2,\dots,n)$, D_i is the set of observations k that fail at time $t_{(i)}$, and R_i is the set of observations j that are at risk at time $t_{(i)}$. Finally, d_i is the number of coal mines that fail at time $t_{(i)}$.³⁸

3.4.2 The Choice of Covariates

Table 2 lists the set of covariates believed to influence the probability of coal mine failure. Among other things, Table 2 shows the variables thought to capture the state variables of Section 3.2. Namely, these are prices, productivity, and age. First, real coal and real oil prices will be measured as the percent change between period t and period $t-1$. The information contained in this percent change

³⁸One of the benefits of choosing a proportional hazard function ($\lambda(t;x)$) and estimating the effects of its covariates using (3) is that the base-line hazard function need not have parametric restrictions on its functional form, yet an empirical estimate of the base-line can still be estimated once the parameter vector is calculated. See Kalbfleisch and Prentice (1980) for details on obtaining an estimate of the base-line hazard.

gives the managers of a mine information about prevailing energy market demand conditions. Positive (negative) percent changes indicate higher (lower) present and expected future profitability, other things equal.

Since the mid-1950s, the generation of electricity has accounted for far and away the largest share of coal consumption—ranging from about 15% of total production in 1955 to 80% of total production in 1994; see Pierce (1996).³⁹ Additionally, most coal is sold under long-term contracts. Joskow (1987) finds that about 15% of coal consumed in electric generation comes from integrated mines, about 15% is traded on a spot market, and the other 70% is purchased by electric utilities under contracts ranging from one to fifty years. The lack of an active market for coal may lead one to question the use of coal prices as a good measure of energy market demand conditions. To address this potential measurement error, real oil prices are used as a covariate since oil is traded very actively in spot and futures markets. It is believed that oil prices are a better measure of energy demand since oil and coal are very highly substitutable in the generation of electric power. Also, since the importance of price as a state variable stems from its role in capturing general energy market conditions, one needs to find a measure that captures the period-to-period fluctuations in energy demand. The prevalence of long-term contracts in coal markets suggests that coal

³⁹ There are two other broad uses of coal other than electric generation. First, coal is used largely for metallurgical coking. However, the decline of the steel industry in the U.S. and the increased use of pig iron in coking have reduced the use of coal for this purpose (Pierce (1996)). Second, coal is exported to more than 30 countries. While this continues to be a growing use of coal, it still pales in comparison to the use of coal in domestic electric power generation; see *Coal Data: A Reference* (1995).

prices are not likely to perform well in capturing energy market demand conditions. For this reason and since oil and coal are easily substitutable in the generation of electric power, oil prices will be the main focus of the hazard models that are to follow. To be sure, increases (decreases) in oil prices should be accompanied by increases (decreases) in coal prices. Hazard ratios of the percent change in price ought to be less than one—representing the idea that periods of higher prices tend to lengthen the time to failure.

Additionally, indicator variables for a mine's labor productivity quartile in year t are included to capture cross-sectional differences in productivity. The notion here is that mines learn about their productivity relative to other mines in the industry. Relative productivity signals to managers that they face comparative advantages or disadvantages in producing coal at a given mine. High (low) realizations tend to lengthen (shorten) the time to failure. Hence, productivity moves inversely with the baseline hazard in the sense that higher productivity mines should have a lower failure threshold (baseline hazard) than lower productivity mines. Figure 4 plots the Kaplan-Meier failure functions by productivity quartile (prod1 represents the least productive quartile while prod4 represents the most productive quartile).⁴⁰ What is interesting to note is that there

⁴⁰ To be sure, a Kaplan-Meier failure function is not the same as the Kaplan-Meier hazard presented in Figure 1; though, they are fundamentally related. A failure function is one minus the survivor function which has the following form: $\hat{F}(t) = \prod_{t_j < t} (1 - \hat{h}_t)$ $\forall t \geq 0$ where \hat{h}_t is the

Kaplan-Meier product limit estimator. Hence, the failure function is $1 - \hat{F}(t)$. See Lancaster (1990). Level differences in the plotted failure functions represent differences in baseline hazards associated with each treatment group. Hence, when looking at Figure 4, for example, the curves

is not a lot of difference between the failure functions of the first and second quartiles, but the differences tend to grow larger as one moves to higher productivity quartiles. Accordingly, in the following empirical models, it is expected that the hazard ratios on all of these quartile dummies should be less than one—indicating that higher productivity mines have longer time to failure than lower productivity mines. It also is expected from Figure 4 that the hazard ratio ought to decline as the productivity class moves up in the productivity distribution.

The age of the coal mine is believed to encapsulate two competing effects. First, age will capture the effects of learning—as argued by the models discussed in Section 3.1. Other things equal, older mines should have lower exit probabilities than younger mines—representing the notion that older establishments have more refined information about relative productivity than younger establishments. However, as suggested by Section 3.2, it also is true that older mines have less coal reserves remaining than younger mines, *ceteris paribus*; this would tend to increase the likelihood of failure. Given these competing effects contained in the age measure, it is difficult at best to say what affect age should have on the time to failure. However, if productivity is included as a direct control for learning, then age (now encapsulating only the depletion effect) should tend to increase the hazard.

themselves are failure functions, but the shifts among them represent changes in the baseline hazards for each treatment group.

Size effects also are included—where size is measure by a coal mine's employment size quartile in a given year. Figure 5 shows the Kaplan-Meier failure functions by employment size class. What this diagram makes clear is that larger mines have higher survival probabilities (viz., longer time to failure) than smaller coal mines. This likely is a result of the fact that larger mines tend to be more productive than smaller coal mines. In manufacturing data, size effects are important sources of productivity growth and hence are important determinants of survival; see Jensen and McGuckin (1997) and Bailey, Hulten, and Campbell (1992). Deily (1991) also finds that firm size effects are important in the exit strategies of steel plants. Following the guidance of the literature on manufacturing and also of Figure 5, it is believed that size effects are important controls, and further they should have hazard ratios less than one.

Mine type is included to control for technological differences among coal mines. Figure 6 shows the Kaplan-Meier failure functions by mine type. Clearly, there are differences in the non-parametric hazards associated with each type of coal mine (viz., surface (12) or underground (11)). Mine type also will control for different sunk costs related to exit, e.g. reclamation costs; high sunk costs tend to reduce turnover in an industry.⁴¹ This control is included in the regression analysis to follow since it appears to capture material differences on the failure thresholds of coal mines—effects that are unrelated to age. In the models to follow, a dummy variable equal to one if a mine is a surface mine is included as a covariate. It is expected that relative to underground mines, surface mines are

more likely to exit, cf. Figure 6. Hence, the hazard ratio associated with this covariate should be greater than one.

The coal producing region is included to control for differences in local factor markets and to control for different coal mine characteristics that vary by location.⁴² Generally, there are three broadly defined coal producing regions: Appalachia, Interior, and the Western Region. Appalachian mines typically are smaller, underground, and more labor intensive than mines in the other two regions. Interior mines generally are larger than the typical Appalachian mine—though smaller than the average Western mine. Interior mines are slightly less labor intensive and are divided between surface and underground mines. Finally, Western coal mines usually are very large, highly productive surface mines that do not employ as many workers as Interior or Appalachian mines. Further, the coal seams of Western mines are very thick and located near the surface—making them relatively inexpensive operations. Figure 7 shows the Kaplan-Meier failure functions by coal producing region. This diagram shows that there is not a lot of difference in the survivor functions of Interior (Region 2) and Appalachian (Region 1) mines, but that both are more likely to fail than Western (Region 3) mines. Hence, the coal producing region should serve as an important control in the following analysis. It is expected from Figure 7 that relative to Appalachian mines (the omitted class), Western mines should have lower hazard ratios than the

⁴¹ See Hopenhayn (1992) and Lambson (1992) for a discussion of this point.

⁴² See Merrell (1999) for a discussion of some of these observable differences. Joskow (1987) also finds this to be an important control when examining long term coal price contracts in the U.S.

other two regions. However, it is difficult to say much about any expectations on the Interior Region control.

3.4.3 The Role of Market Conditions

The first column of Table 3 reports hazard ratios that illustrate the effects of energy market conditions (as captured by oil prices) on the failure patterns of coal mines. Recall from Section 3.2 that one of the random state variables affecting the participation decision of coal mines is a measure of energy market demand conditions. The idea here is clear. Realizations of the price state variable have a direct effect on the current profitability of a given mine: high (low) realizations tend to promote continuation (exit) for a given remaining stock of coal and for a given relative productivity. Prices also enter the problem through the formation of expectations surrounding future profitability. Specification one in Table 3 shows that the percent change oil prices between periods $t-1$ and t has a very slight positive influence on the time to failure. That is, the estimate is statistically significant and is estimated to be 1.0086 (0.0006). Strictly speaking, this estimate does suggest that higher prices tend to shorten the time to failure. However, the effect, while statistically significant, is so slight that it does not appear to have material importance on the time to failure. Also, to capture expectations of future energy market conditions, the percent change in real oil prices between period $t+1$ and $t+2$ is included. The hazard ratio estimate of this effect suggests that higher (lower) future market conditions significantly lengthens (shortens) the time to failure.

3.4.4 The Role of Mine Heterogeneity

The second column of Table 3 presents the effects on the baseline hazard from including factors that describe exogenous differences in productivity as well as cross-sectional differences in labor productivity. Specifically, these controls include a mine's coal producing region (a time invariant covariate), employment size quartile (a time variant covariate), and mine type (a time invariant covariate). The inclusion of these covariates is believed to capture variations in productivity not related to the learning effect captured by the cross-sectional differences in labor productivity. Since these heterogeneous factors are qualitative variables, it should be noted that the omitted classes are the smallest productivity quartile, the smallest size quartile, underground mines, and the Appalachian coal producing region. Hazard ratios, then, are interpreted relative to those omitted classes. Recall the discussion regarding the expected effects of each of these covariates on the baseline hazard: that underground mines, more productive mines, larger mines, and Western mines should have lower time to failure relative to surface mines, less productive mines, smaller mines, and Appalachian mines.

The second specification of Table 3 shows the effects of these exogenous heterogeneity effects along with the cross-sectional differences in mine labor productivity. First, relative to less productive mines, higher productivity mines tend to have longer time to failure. Additionally, relative to the first quartile, the second quartile group does not seem to have statistically different time to failure than the first group; this is consistent with Figure 4. Finally, with respect to

productivity, when comparing across productivity quartiles, the hazard ratio declines monotonically in the size of the class; this finding also is consistent with Figure 4. In short, high (low) productivity tends to lengthen (shorten) time to failure. Second, relative to smaller mines, larger mines have significantly longer time to failure; another interesting thing to note is that the proportional effect on the hazard declines monotonically in the size class—consistent with Figure 5. Third, the estimate of the effect of mine type on the hazard is significant and greater than one—suggesting that relative to underground mines, surface mines tend to have shorter time to failure. This finding is consistent with Figure 6 that shows that surface mines are more likely to fail relative to underground mines. Finally, relative to Appalachian mines, Interior mines have a shorter time to failure, while Western mines have a statistically longer time to failure.

Recall from the discussion of Section 3.2 that the role of age is believed to encapsulate two effects: learning and resource exhaustion—each of which should have different (in fact, opposite) effects on the hazard. Learning provides managers with information about relative productivity while age proxies for the resource depletion effect. The empirical problem, then, is to disentangle these two effects contained in the age covariate. By controlling for productivity explicitly, the only effect left in the age covariate should be resource exhaustion. By controlling for observable differences in productivity that are not related to learning (e.g., region, size, and mine type) or to resource exhaustion, the primary effect left in the productivity class variables should be the learning effect, while

the remaining effect in the age covariate is resource exhaustion. It is expected that productivity will lower the hazard (viz., lengthen the time to failure), while age ought to increase the baseline hazard (viz., shorten the time to failure). The third column of hazard ratios in Table 3 shows the effects of the exogenous productivity differences, the cross-sectional differences in productivity, and age on the baseline hazard. It is important at this point to keep in mind the competing nature of the hazards at work: market conditions and productivity will tend to lengthen the time to failure while exhaustion tends to shorten time to failure.

Column three of Table 3 shows first that productivity does in fact tend to lengthen the time to failure; all of the estimates on the productivity quartile dummies have the same qualitative effects discussed in the previous specification. This is consistent with the notion that higher productivity mines are less likely to fail than lower productivity mines. Second, column three of Table 3 reports that the hazard ratio associated with the age covariate is significantly greater than one. This suggests that older coal mines have shorter time to failure—a finding consistent with the argument of Section 3.2. That is, the competing hazards associated with productivity differences and with aging can be separated such that productivity lengthens time to failure while age shortens time to failure—this latter finding representing the notion that, *ceteris paribus*, every short ton of coal extracted from a mine reduces expected future profitability since coal is non-renewable.

3.4.5 Market Conditions and Mine Heterogeneity

Now that there is some understanding of how each of the three state variables affects the time to failure, the problem is to examine all of these effects together in a single specification. Column four of Table 3 shows hazard ratios of the effects of the percent change in real oil prices, the percent change in future real oil prices, the cross-sectional differences in labor productivity, age, and a set of exogenous controls that are believed to capture differences in productivity not attributable to learning or aging. For the most part, the predictions on the state variables hold in this specification. First, the percent change in real oil prices retains its value and significance—again suggesting that this covariate has little material importance to the time to failure, and the percent change in future oil prices still significantly reduces the time to failure. Second, the cross-sectional productivity differences still significantly raise the time to failure. Third, age significantly shortens the time to failure suggesting, *ceteris paribus*, that older mines suffer from resource exhaustion and hence are more likely to fail than newer mines. Finally, all of the controls have the same impact as before—except that the surface mine dummy loses statistical significance. On balance, these covariate estimates support the notion of coal mine exit dynamics outlined in Section 3.2.

3.4.6 The Case of Coal Prices

Ideally, one would repeat the above exercises by including coal prices as covariates to determine the direct effect of changes in these prices on the exit patterns of mines. There are a couple of problems that have been encountered

with this approach. First, because of the pervasive use of long-term price contracts in coal markets, it is likely that coal prices capture variations in energy market conditions less accurately than oil prices. It is believed that oil prices do a better job than coal prices when attempting to explain general energy market demand effects.

There is another, more technical, problem with using coal prices as a covariate in the Cox model. The optimization routine used to estimate the Cox proportional hazard model with real coal prices as covariates has very significant problems converging on a consistent basis to an optimum on the likelihood function. In fact, for those times when the algorithm does converge, subsequent estimates on the coal price covariates vary markedly across different runs of the model. It is unclear at this point why the algorithm does not seem to behave well with coal prices; though, the issue is being studied. In addition to the idea that coal prices probably do not perform as well as petroleum prices in capturing energy market conditions, it is unclear how much trust one should place on covariate estimates on coal prices when the estimation algorithm has problems converging to an optimum value on the likelihood function.⁴³

⁴³ An alternative approach can be used to obtain non-proportional covariate estimates of the effects of coal prices on the time to failure. That is, using event history analysis, one can non-parametrically control for the baseline hazard while estimating the effects of covariates on this baseline hazard. The story proceeds in the following way. For any coal mine i , the odds of failing at each discrete point in time, $t_i=1, 2, \dots$, are a function of the odds of failing for a group of mines that represent the baseline states of covariates such that

$$\frac{\lambda(t_i; X)}{1 - \lambda(t_i; X)} = \frac{\lambda_o(t_i)}{1 - \lambda_o(t_i)} \exp\left(\sum_k \beta_k x_k\right)$$

where $\lambda(t_i; X)$ is the conditional probability of failing at time t_i for a given vector of k covariates, $X=(x_1, x_2, \dots, x_k)$, and a given set of k parameters, β_k . To be sure, note that the baseline hazard

3.5 Future Directions for this Research

The work in this paper is somewhat preliminary in the sense that it relies on some assumptions that may overly constrain the flexibility of the analysis. This section discusses ways in which these assumptions can be relaxed and how those relaxations can improve the quality of this project. Additionally, this section details a complementary methodology that could be employed to examine the participation decisions of coal mines.

3.5.1 Unobserved Heterogeneity

function, $\lambda_0(t_i)$, is characterized by the conditional probability for cases where the covariate vector is equal to zero; this is the baseline state of the covariates. Taking the natural logarithm of this equation allows one to estimate this model as a discrete-time logit regression of the following form:

$$\ln \left\{ \frac{\lambda(t_i; X)}{1 - \lambda(t_i; X)} \right\} = \alpha_i + \sum_k \beta_k x_k$$

where $\alpha_i = \ln(\lambda_0(t_i)/1 - \lambda_0(t_i))$, or the log-odds for the baseline group. Including dummy variables for each of the failure (event) times will create a separate intercept for each discrete point in time along the baseline hazard—allowing one to trace out (non-parametrically) the baseline hazard. The final thing to note is that if all of the covariates are time-invariant, then this model is a proportional odds model. However, if some or all of the covariates are time-variant, then the model is non-proportional—the effect being that the duration effects and time effects cannot be separated from one another. For very detailed discussions of this sort of estimation strategy see Allison (1984) or Yamaguchi (1991).

Using this approach, the parameter estimate (not the hazard ratio), with standard errors in parentheses, on the percent change in real coal prices is -0.0046 (0.0017), which is significant at the 1% level. Additionally, the effect of the log lead real price of coal is -0.4041 (0.0802), which also is significant at the 1% level. This model has a log likelihood value of $-19,283.865$. These findings suggest the high (low) prices of coal tend to lengthen (shorten) the time to failure. This is consistent with the findings of the Cox proportional hazard model when oil prices are used to capture energy market conditions as well as the argument of Section 3.2.

Finally, this approach itself has some limitations. Namely, the loss of the proportionality feature makes it difficult to interpret the findings. For example, when more complete specifications (such as column four in Table 3) are estimated, the effect on the age variable now tends to *shorten* the time to failure. This is an artifact of the non-proportionality of this type of hazard model. The inability of being able to separate duration effects from time effects will tend to understate the age variable since it increments over time just like the baseline hazard. To be able to understand fully the effect of aging (resource depletion) one must be able to mete out the time effects from the duration effects. So, while this approach does get more believable estimates for coal prices, those estimates come at the cost of losing the proportionality feature of the Cox

The resource endowments that coal mines face are not observable econometrically. However, these endowments represent sources of heterogeneity that could account for differences in baseline hazards. Heckman and Singer (1984a, 1984b, 1984c) show that failing to account for unobserved heterogeneity tends to bias estimates toward negation duration dependence. The results shown here find strong and significant patterns of positive duration dependence; so, failing to correct of unobservable heterogeneity will not reverse the results found thus far. However, more precise estimates could be obtained by controlling for unobserved heterogeneity. Examples of correcting for unobserved heterogeneity using the Heckman-Singer approach can be found in Trussell and Richards (1985) and Meyer (1990).

One concern that surrounds this sort of idiosyncratic effects model is that the endowment of coal faced by mines in each period is itself a function of time. The Heckman-Singer approach relies on this unobservable idiosyncratic effect being time invariant. To be able to integrate over the full idiosyncratic effect (viz., the cross-sectional and time series effects) which is the essence of the Heckman-Singer correction, then another dimension is needed in the data. Current efforts in this work focus on, among other things, finding a defensible “third” dimension—perhaps firm-level or regional effects.

3.5.2 Independence of Subsequent Spells

model. (Note: Any time a proportional hazard is run (with or without parametric restrictions on the baseline hazard), the effect of age is to reduce the time to failure.)

The empirical analysis in this paper is based on multiple failure observations and also relies fundamentally on the assumption that subsequent spells are independent. This places a good deal of a priori structure on what is believed to be the decision process of coal mine managers. It likely is the case that spells, in fact, are not independent, and failing to control for this serial dependence would tend to misstate the effects of the covariates on the time to failure. Future work also will focus on the effects of lagged spells on current spells—though it is not clear whether failing to control for lagged durations will under- or overstate the covariate effects already examined. Heckman and Walker (1987) present a model that allows for lagged durations to influence current spells.

3.5.3 Modeling the Participation Decision

Finally, since the death set is large relative to the risk set (viz., there is a good deal of multiple spells), another methodology would be to model the participation decision of a coal mine. The idea is to examine what influences the hit-and-run nature of the participation patterns of coal mines. Using a latent variable approach in a dynamic setting, it is possible to estimate the determinants of market participation. This sort of setting controls for duration dependence by using k-factorized random effects probit model with a lagged dependent variable (which itself is a latent variable); see Heckman (1981a, 1981b). Such an approach has been used to analyze the decision to export in Colombian manufacturing (Roberts and Tybout, 1997) and U.S. manufacturing (Bernard and

Jensen, 1997). Similar methods could be applied to coal mining in order to describe the frequent exit, re-entry, and re-exit behavior of a large number of coal mines.

3.6 Conclusion

This paper examines the determinants of failure in the U.S. coal mining industry. Relying on the predictions of a theory of firm dynamics that incorporates product market conditions, learning through productivity signals, and the effects of resource depletion, this paper seeks to mete out the competing effects of product market conditions and productivity from the diametric effect of resource depletion. That is, the first two effects should tend to lengthen the time to failure, while the resource depletion effect should tend to shorten the time to failure, *ceteris paribus*. Using a unique microdata set containing the statistical universe of U.S. coal mines observed from 1974 to 1995 and also observed failing multiple times, Cox proportional hazard models are estimated to determine the influence of favorable energy market and productivity conditions as well as the effects of aging on the exit patterns of U.S. coal mines.

It is found that the percent change in oil prices between periods $t-1$ and t has a very slight tendency to shorten the time to failure; however, the slight magnitude of this covariate suggests that it has little material importance on the failure threshold for coal mines. It also is found that the percent change in real oil prices between periods $t+1$ and $t+2$ significantly lengthens the time to failure—suggesting that high (low) future oil prices tends to lengthen (shorten) the time to

failure for a coal mine. Additionally, it is found that, after controlling for exogenous differences in productivity that are not attributable to learning, higher productivity coal mines tend to have longer time to failure relative to lower productivity mines. Again, this finding is consistent with the notion of firm dynamics discussed in this paper. It is found, further, that, after controlling for exogenous differences in productivity as well as cross-sectional differences in productivity, older mines tend to have shorter time to failure than younger mine, *ceteris paribus*. Finally, when examining all three effects in a single specification of the proportional hazard model, all of the covariate effects on the time to failure hold their values and significance—except that the price covariate loses statistical significance.

In closing, this paper revisits the determinants of business failure but in the special case of an industry where establishments have exogenously imposed constraints on lifetime production. The effect of these constraints is to introduce a competing dynamic to the exit or continuation decisions of mines. The findings presented in this paper are consistent with the predictions surrounding the convolution of all three of these forces. These findings are estimated from an interesting microdata set not within the manufacturing universe—as is the case of most empirical studies of business failure. Finally, improvements to this work could include controlling for unobservable heterogeneity and the effects of lagged durations on current durations as well as proposing a complementary methodology to describe the large number of multiple failure events found in the

data. Altogether, these findings suggest that the microeconomic adjustment dynamics of extractive industries is complicated substantially by the inclusion of the resource depletion dynamic.

Table 1. Number of Mines by Year: 1974 to 1995

Year	Number of Mines	Change
1974	371	N/A
1975	789	+418
1976	1,212	+423
1977	1,629	+417
1978	1,831	+202
1979	2,064	+233
1980	1,896	-168
1981	2,057	+161
1982	1,924	-133
1983	1,683	-241
1984	2,069	+386
1985	1,858	-211
1986	1,622	-236
1987	1,656	+34
1988	1,643	-13
1989	1,598	-45
1990	1,548	-50
1991	1,536	-12
1992	1,337	-199
1993	1,264	-73
1994	1,157	-107
1995	1,065	-92
Total	33,809	+694

Table 2. List of Covariates and Definitions

Covariate	Definition
Coal Price	Industry Average Annual Price Measured as Dollars per Short Ton. Source: <i>Minerals Yearbook</i> and <i>Coal Industry Annual</i> .
Oil Price	Industry Average Annual Price Measured as Dollars per Barrel. Source www.eia.doe.gov
Age	Total number of years a mine is active. Source: Author's calculations.
Labor Productivity Quartile	Quartile values of the natural logarithm of short tons of coal per worker hour. Source: Author's calculations.
Mine Type	Underground or Surface Mine. Source: Coal Address and Employment Files, U.S. Mine Safety and Health Administration.
Size Quartile	Employment Size Quartile. Source: Author's calculations.
Region	Coal Producing Region. Source: Coal Address and Employment Files, U.S. Mine Safety and Health Administration.

**Table 3. Proportional Hazard Ratios for the Effects of Oil
Market Conditions and Mine Heterogeneity
(Standard Errors)**

Covariate	Market Effects (1)	Heterogeneity Effects without Age (2)	Heterogeneity Effects with Age (3)	Market Effects and Heterogeneity Effects (4)
%Change in Oil Prices	1.0086* (0.0006)			1.0086* (0.0006)
%Change in Lead Oil Prices	0.6464* (0.0331)			0.6207* (0.0320)
Age Since De Novo Entry			1.0736* (0.0032)	1.0747* (0.0032)
Productivity Quartile 1		-----	-----	-----
Productivity Quartile 2		0.9707 (0.0286)	0.9874 (0.0291)	0.9754 (0.0287)
Productivity Quartile 3		0.8127* (0.0245)	0.8483* (0.0257)	0.8421* (0.0255)
Productivity Quartile 4		0.6799* (0.0212)	0.7241* (0.0227)	0.7178* (0.0225)
Mine Type (Surface)		1.0627* (0.0250)	0.9838 (0.0233)	0.9799 (0.0233)
Size Quartile 1		-----	-----	-----
Size Quartile 2		0.8532* (0.0231)	0.8439* (0.0228)	0.8450* (0.0228)
Size Quartile 3		0.7345* (0.0218)	0.7171* (0.0213)	0.7218* (0.0214)
Size Quartile 4		0.4671* (0.0214)	0.3931* (0.0183)	0.3931* (0.0183)
Appalachian Region		-----	-----	-----
Interior Region		1.1239* (0.0511)	1.1875* (0.0541)	1.1796* (0.0537)
Western Region		0.7888* (0.0837)	0.6875* (0.0731)	0.6749* (0.0718)
Log Likelihood	-74,836.369	-74,649.520	-74,385.362	-74,220.528

* Denotes significance at .01





Figure 1. Kaplan-Meier Hazard for U.S. Coal Mines

Figure 2. Total Hours, Production, and Employment: 1972 to 1996








Figure 4. Failure Function by Productivity Quartile




Figure 5. Failure Function by Size Quartile

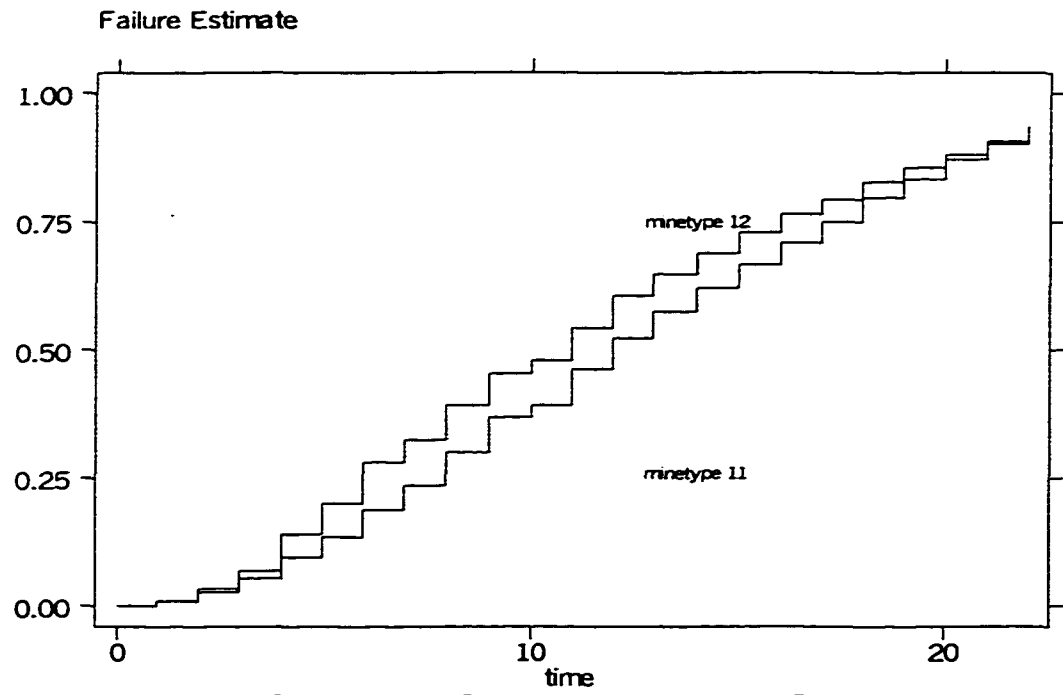


Figure 6. Failure Function by Mine Type




Figure 7. Failure Function by Region

4. Gross Employment Flows in U.S. Coal Mining

Abstract: This paper examines the patterns of job creation and destruction in the U.S. coal mining industry and compares those patterns to known regularities in U.S. manufacturing. Using a unique panel data set containing the statistical universe of coal mines from 1973 to 1996, this paper calculates rates of job creation, job destruction, job reallocation, and excess job reallocation. The main findings of this paper are as follows: i. Annual employment flows in coal mining are substantially higher than corresponding flows in manufacturing, ii. The high rate of job creation, destruction, and reallocation is attributable in part to mine openings and closing, iii. A significant amount of job destruction and job creation is attributable to temporary shutdowns and re-openings, and iv. Job destruction is counter-cyclical and more volatile than job creation—similar to the known patterns in manufacturing.

4.1 Introduction

Over the course of the last decade, economists have been increasingly interested in the job creation and job destruction process in economies. This fact is evidenced by the recently published third volume of the Handbook of Labor Economics that has a separate chapter devoted to the employment flows literature.⁴⁴ One key finding in this literature is that rates of job creation and job destruction (hence, job reallocation) are quite high in comparison to net changes in employment. For example, in U.S. manufacturing, net employment change between 1973 and 1992 averaged about -0.8% annually over the period—a relatively small net change. However, the underlying average annual job creation and job destruction rates are quite high—averaging about 9% and 10% , respectively. These are ten times the net employment change. Clearly, the job

⁴⁴ The chapter is entitled “Gross Job Flows” and is authored by Steven Davis and John Haltiwanger and appears in the Handbook of Labor Economics (1999), Volume 3B.

creation and job destruction rates show that job reallocation greatly exceeds the net change in a given period. These findings have implications in a number of strands of economic literature including labor economics, macroeconomics, and industrial economics.

First, with respect to labor economics, the literature on gross employment flows (job creation and job destruction) has highlighted the importance of demand-side fluctuations in employment reallocation. The amount of observed demand-side fluctuations has impacts for analyzing worker mobility, in examining dislocations of workers (e.g., unemployment), and for analyzing matching models of workers and firms.⁴⁵ Second, macroeconomists have shown great interest in analyzing job creation and job destruction series as a mechanism to understand business cycles better. The empirical data from U.S. manufacturing suggest that job creation and job destruction behave quite differently over the business cycle. Job destruction appears to be disproportionately intense during recessionary periods and also appears to be more volatile than job creation. Macroeconomic models with asymmetries in adjustment costs across hiring and firing or where there are significant fixed costs of altering employment levels have been developed to explain these differences in the nature of the series.⁴⁶ Finally, job creation and job destruction series, in conjunction with information

⁴⁵ For example, Hopenhayn and Rogerson (1993) examine interaction of government policy that reduces worker mobility and employment reallocation in a general equilibrium framework. Mortensen and Pissarides (1994) develop a model of unemployment and job creation and job destruction.

⁴⁶ Research in the macroeconomics on employment flows include Blanchard and Diamond (1990), Caballero and Engel (1993), Caballero, Engel, and Haltiwanger (1997), and Campbell and Kuttner (1997).

on business openings and closings, have been used by industrial economists to examine industrial dynamics and patterns of industrial evolution. These calculations have highlighted the importance of entry and exit, the growth of plants and firms, and the impact of reallocation of output and employment shares on productivity growth.⁴⁷ In short, the development of new data sets that allow researchers to examine the employment dynamics of individual plants and firms has created new empirical insights that have formed the basis for a wide number of both theoretical and empirical studies.

This paper will extend this research by creating a new data series on job creation and job destruction for the coal mining sector of the U.S. economy. The previous U.S. literature on job creation and job destruction has focused predominately on U.S. manufacturing industries. Research by Davis and Haltiwanger (1992), Davis, Haltiwanger, and Schuh (1996) and Dunne, Roberts, and Samuelson (1989) all examined employment flows in U.S. manufacturing. There is a small literature examining non-manufacturing sectors stemming from work done by the U.S. Small Business Administration and by Foster, Krizan, and Haltiwanger (1998). However, neither of these studies is able to develop annual time series on non-manufacturing industries, and in fact, all are constrained to examining gross employment flows over five year intervals. Hence, the time-

⁴⁷ Industrial organization papers that examine employment dynamics include Evans (1987), Dunne, Roberts, and Samuelson (1989) and Troske (1996). Papers examining reallocation and productivity growth include Bailey, Hulten, and Campbell (1992), Olley and Pakes (1996) and Aw, Chen, and Roberts (1997).

series data on non-manufacturing sectors in the U.S. is relatively scarce.⁴⁸ This new data on the U.S. coal mining sector has the strengths both of providing an annual time series of employment flows over a long period of time and of being able to study those changes outside of the manufacturing sector.

Moreover, coal mining has some interesting institutional features that may yield differences in the job creation and destruction series as compared to manufacturing. First, the life span of a coal mine is limited by coal reserves present at the mine. As reserves run out, mines (even very productive ones) are forced to close; see Merrell (1999). This suggests that mine openings and mine closings may play a larger role in employment reallocation in the mining sector as compared to manufacturing. Second, the mining sector is more heavily unionized than manufacturing and is more subject to strikes and lockouts. Both of these related factors may induce more volatility into the job creation and job destruction series. Finally, coal mining has undergone dramatic changes in mine technology, mine type (underground versus surface), and locations (Western coal versus Appalachian coal) that have yielded substantial improvements in labor productivity. These changes have resulted in a substantial net decline in employment in the industry even though production has risen markedly.

The main findings of this paper are fourfold. First, annual employment flows in coal mining are substantially greater than employment flows in manufacturing. In coal mining, job creation and job destruction exceed

⁴⁸ There are some studies of individual states that include non-manufacturing data. Davis and Haltiwanger (1999) summarize the various US and international studies. As yet, there are no U.S.

manufacturing levels on average by roughly 50% and 80%, respectively. Second, the high rate of job creation, job destruction, and employment reallocation is due to the fact that the proportion of employment gained through mine openings and the proportion of employment lost through mine closings is much greater in mining than in manufacturing. In mining, on average, two-thirds of annual job destruction is attributable to mine closings, while one-half of annual job creation is attributable to mine openings. Third, a significant fraction of job destruction and job creation due to mine closings and mine openings is temporary in nature. Unlike manufacturing plants, it is quite common for coal mines to close for a year (or more) and then re-open. Fourth, with respect to cyclicalities of job destruction and job creation, job destruction appears to be more volatile than job creation in coal mining—similar to the pattern in manufacturing as well. However, the prediction of certain models that recessions tend to have a “cleansing” effect is more difficult to ascertain.⁴⁹ This is because the energy shock induced recessions of the 1970s have a different impact on the coal mining sector than on the economy as a whole; coal actually *expanded* during these energy driven recessions. In the deep recession of 1982-1983, the job creation and job destruction series in coal mimic the corresponding series in manufacturing. However, overall, it is the case that job reallocation is counter-cyclical in coal mining.

studies that examine the U.S. coal mining industry.

⁴⁹ A particular line of modeling in the macroeconomics literature has developed macroeconomic models incorporating non-convex costs of adjustments. Papers in this line of literature (e.g.,

The remainder of the paper proceeds as follows. In the next section, we describe measurement issues and the data used in the study. In section 4.3, we provide an analysis of the time series and distribution of employment flows for coal mining. Section 4.4 provides some brief closing comments.

4.2 Measurement and Data Issues

The data used in the analysis come from the Mine Safety and Health Administration's (MSHA) data files on employment, hours and production for all coal mines in the United States. Unlike other data sources, the data on U.S. coal mining represent a complete sample of all mines in every year. Data are available quarterly on the number of workers in the mine, the number of hours worked, and the short tons of coal produced. In the analysis that follows, we will focus our attention on annual (rather than quarterly) job creation and destruction statistics for two reasons. First, the annual data are more comparable to other studies (see Davis and Haltiwanger (1998)) to which we make direct comparisons. Second, the annual data have fewer measurement error problems than the quarterly data.⁵⁰

Table 1 provides some descriptive statistics on the number of mines, the number of employees in coal mining, and the average mine size for 1975, 1980, 1985, 1990, and 1995. The data show that there has been a marked decline in

Caballero and Hammour (1994)) predict that firms use recessions as periods of restructuring, and hence, predict that employment reallocation in an economy will tend to be counter-cyclical.

⁵⁰ The quarterly coal mining data appear to have "fourth-to-first" quarter discontinuity in employment flows. That is, the measured employment flow going from the fourth quarter of period $t-1$ to the first quarter of period t generally is quite large. This may be due to improper data smoothing that would impute data to remaining quarters in a year when in fact the mine is closed. This will have a tendency to bunch up mine closures in the fourth-to-first quarterly period. We are

both the number of mines and in mine employment in the late 1980s and early 1990s. From 1975 to 1995, employment fell by over 50% for these mines. However, mine size as measured by average employment, has declined by only 10% over the same period. While not shown in Table 1, the decline of employment is not reflective of a decline in production. In fact over this period, coal production increased by 66% in these mines even in the face of the steep decline in employment and worker hours. The dramatic decline in workers accompanied by increasing production combine to show substantial growth in labor productivity since the middle of the 1980s. This growth in labor productivity has been extensively examined in Berndt and Ellerman (1997).

In this paper, we will utilize the data on mine-level employment to construct measures of annual gross employment flows. Our approach will follow that of Davis and Haltiwanger (1992) and Davis, Haltiwanger, and Schuh (1996) and define job creation (POS) and job destruction (NEG) accordingly as follows:

$$(1) \quad POS_{it} = \frac{TE_{it} - TE_{it-1}}{.5 * (TE_{it} + TE_{it-1})}$$

$$(2) \quad NEG_{it} = \left| \frac{TE_{it} - TE_{it-1}}{.5 * (TE_{it} + TE_{it-1})} \right|$$

where TE_{it} represents total employment at mine i in quarter one in period t and TE_{it-1} represents total employment in mine i in quarter one in period $t-1$.⁵¹ TE includes total employment at the mine but does not include coal processing

currently investigating the source of the problem. We do not believe that this will impact our first quarter to first quarter measures.

operations or white-collar workers. The POS_{it} variable is constructed for all plants with expanding employment between $t-1$ and t , while NEG_{it} is constructed for all plants with contracting employment between $t-1$ and t . The denominator used here is the same as Davis and Haltiwanger (1992) and represents the average employment at a mine in periods t and $t-1$. To be sure, this formulation constrains the job creation and job destruction rates to the interval $[0,2]$ where an opening or closing mine has a job creation or destruction rate of two, respectively. Finally, both POS_{it} and NEG_{it} are used to construct aggregate measures of job creation and job destruction. POS_t is the employment weighted average of job creation in period t , and NEG_t is the employment weighted average of job destruction in period t . These are main variables used throughout the analysis that follows.

In addition to the POS_t and NEG_t aggregate employment flow variables, we also construct the following three additional measures of aggregate employment change:

$$(3) \quad NET_t = POS_t - NEG_t$$

$$(4) \quad REAL_t = POS_t + NEG_t$$

$$(5) \quad EXCESS_t = REAL_t - |NET_t|$$

NET_t is net employment growth rate in coal mining, $REAL_t$ represents total reallocation of employment across mines between $t-1$ and t , and $EXCESS_t$ measures excess employment reallocation that exceeds the amount necessary to accomplish the net change. In this case, $EXCESS$ is merely $REAL$ minus the

⁵¹ Throughout the paper, a first quarter to first quarter comparison is used. This is identical to the

absolute value of NET. To be sure, it used to measure the amount of employment reallocation that occurs in excess of the amount needed to facilitate the net employment change in the industry.

4.3 Annual Employment Flows

This section presents the empirical analysis of gross employment flows in the U.S. coal mining sector. First, we present some aggregate facts about employment flows. Second, we present some facts about the importance of mine openings and closures. Finally, we discuss the role of mine heterogeneity in describing differences in the patterns of mine growth.

4.3.1 Aggregate Trends

The job creation, job destruction, and reallocation series are presented in Table 2 for the period 1973 through 1996. The median annual job creation rate is 0.158, while for job destruction the median rate is 0.186. These rates are considerably higher rates of job creation and job destruction than those observed in manufacturing. Figure 1 plots the job creation and destruction series for coal mining. Examining the job creation series, we see that job creation was relatively high in the 1970s, a period of rapidly rising energy prices, and yet job creation was relatively low in the 1980s and 1990s, a period of substantial employment decline in the coal mining sector. In general, the opposite patterns occur for the job destruction rates. It is generally lower in the 1970s and rises thereafter.

construction used by Davis, Haltiwanger, and Schuh (1996).

A couple of specific points are worth noting. First, the large spike in job destruction in 1977 accompanied by the large spike in job creation in 1978 is due to an industry strike that occurred in the 1977-1978 period. The Energy Information Agency (1995) reports that the longest major coal miners' strike since 1960 occurred in the 1977-1978 period and lasted 111 days. The impact on our data is to reduce significantly employment in the first quarter of 1978. This creates a large job destruction rate for the 1977-1978 period and a corresponding large creation rate from 1978 to 1979. There are other strikes that do occur over the period, but none have the impact of the 1977-1978 strike.

Second, the level of job creation and job destruction is much higher in coal mining than what is found in other sectors of the economy. In U.S. manufacturing, job creation and destruction rates average around 0.09 and 0.10, respectively, for the 1970s and 1980s. Given that the coal data are a new data source, one naturally must wonder if these differences result from measurement problems. To gauge the quality of the measures presented in this paper, we identified a corroborating data source that provides gross employment flows data for mining industries. For the 1995 to 1996 period, the U.S. Census Bureau constructed estimates of job creation and destruction rates for the mining sector—including coal mining; also included are nonferrous metals mining, iron ore mining, and stone quarries. The Census Bureau reports that the mining sector had a job creation rate of 15.2% in the period 1995 to 1996 and a destruction rate of 24.1% over the same period. The data from the MSHA indicate that the coal

mining job creation rate was 15.1% over the period, and the job destruction rate was 24.2% over the same period. The data from the Census Bureau and MSHA are surprisingly similar. Given the fact that the MSHA data have not been previously used for this purpose, the closeness of the MSHA-based figures to the Census data increases our confidence in this new data source.

In addition to the job creation and job destruction, we are also interested in net employment change and excess job reallocation. Figure 2 presents information on job creation, job destruction, net change, and excess reallocation normalized around the series means. This graph allows a better comparison of the relative fluctuations of all of these series over time. The most dramatic episodes are the during strike period (1977-1978) which affects the net change and overall reallocation series and during the 1982-1983 recession where net change in industry employment exceeds -20%. Note, in the strike period, excess reallocation is relatively unaffected since the strike will affect both the net and reallocation terms similarly. The key point here is that impact of recessions on the job creation and job destruction series is somewhat asymmetric. Job destruction appears more cyclically sensitive than job creation. This is particularly true in the large (non-energy induced) recession of 1982-1983. Job destruction increases markedly, while job creation falls modestly. If one computes the coefficient of variation for each series (omitting the strike years), one finds that the coefficient of variation for job destruction series is 37.4% higher than the coefficient of variation for job creation (32.87 versus 23.92). In

short, job destruction is more cyclically volatile than job creation—a finding consistent with the U.S. manufacturing evidence.⁵²

With respect to employment reallocation and excess employment reallocation, we see that excess reallocation is relatively stable over the 22 year period: around 25%. This rate is higher than the analogous rate in manufacturing which is roughly 18%. In terms of the cyclicity of reallocation, we find that reallocation is negatively correlated with industry growth. That is, employment reallocation is counter-cyclical. The Pearson correlation (again, excluding strike years) between NET and REAL is estimated at -0.579 (.006). This is also quite consistent with the patterns found in manufacturing where employment reallocation is found to be counter-cyclical, and this finding also agrees with the predictions of models of job creation and job destruction as offered by Caballero and Hammour (1994) and Campbell and Fisher (1996). These models predict asymmetries in adjustment of firm employment over the cycle based on asymmetries and non-convexities in adjustment costs. This will cause firms to concentrate their employment changes into discrete episodes. In terms of timing across the cycle, the opportunity cost (lost revenue) of restructuring (e.g., retooling a mine, changing technology, retraining workers) in an expansionary period is greater than the lost revenue in a contractionary period.⁵³ Hence, firms

⁵² Note, however, that the 1974-1975 recession is not apparent in the coal data. This is because this recession was energy related, and energy related industries felt little of the impact of the recession on output or employment.

⁵³ See Davis and Haltiwanger (1999) for a more detailed discussion of these issues.

will have a tendency to concentrate their employment restructuring episodes in downturns.

4.3.2 Mine Openings and Closings

Figure 3 provides a breakdown of the employment flows by source of employment flow: mine openings, mine closings, continuing mine expansion, and continuing mine contraction. One should be careful to note that we treat both a de novo mine and a reopened mine as an opening, and a mine shutdown is considered a closing even if it is temporary. Basically, an opening is any mine that goes from zero to positive employment between t and $t+1$, and a closing is any mine that goes from positive employment to zero employment between t and $t+1$. Figure 3 shows that the opening and closings series are quite important in terms of job creation and job destruction. About two-thirds of job destruction is attributable to mine closure while about one-half of job creation is attributable to mine openings. This importance of closing mines is particularly pronounced in the late 1980s and 1990s. The job destruction due to mine closings series depicted in Figure 3 is substantially higher than the other three series in the late 1980s and 1990s. This pattern of extensive job creation and job destruction due to openings and closings is starkly different than what is found in manufacturing. In manufacturing, job destruction due to closures accounts for about 23% of overall job destruction, and job creation due to openings only accounts for 16% of total job creation. Hence, openings and closings play a much more important role

in reallocating employment in the mining sector than would be found in manufacturing.

The important role that openings and closings play is further highlighted in Figure 4a and Figure 4b. These figures present data on the micro-distribution of job creation and job destruction across mines. The data presented are in rate form and are constrained to the $[0,2]$ interval. Recall, a rate equal to two indicates an opening or closing in the relevant panel. Looking first at the job creation panel (Figure 4a), it is clear that mining is characterized by relatively small employment adjustments of continuing mines (the large bars at the left of the graph) and a large number of openings (the large spike at the rightmost column of the graph). In fact, over 35% of all employment expansions in coal mining are mine openings. Figure 4b shows that closings make up an even higher percentage of mine contraction episodes. A little over 50% of all contraction episodes result from mine closures.

A key difference between manufacturing and mining is that a significant fraction of mine closings and openings are temporary in nature. Table 3 provides some basic information regarding the nature of mine reopenings and temporary closings. The table provides data on temporary closings and reopenings for five years: 1975, 1980, 1985, 1990, and 1995. Column 2 reports the percent of closing mines in a year that will re-open in the future. Data for 1995 will be censored quite heavily, so it is not presented. The data show that a substantial fraction of closing mines will re-open. For example, in 1985, 28.4% of the mines that close

that year will reopen between 1986 and 1996 (the last year of the data). A similar pattern is observed for mine reopenings. Column 3 presents data on the percent of opening mines that were in previous operation in the sample. Again, because our data only go back to 1973, data for 1975 will be censored heavily and as before, will not be presented. In 1985, slightly over 40% of opening mines had been operating in the 1973-1983 period. On an employment weighted basis, temporary closings account for approximately 20% to 30% of job destruction attributable to plant closings. In the case of mine openings, re-openings account for 40% to 50% of job creation due to mine openings. Again, it is difficult to be overly precise here because of censoring in these data series. However, the key point is that temporary mine closings and mine reopenings are an important feature of the job creation and job destruction process in the coal mining sector.

4.3.3 Mine Heterogeneity and Employment Growth

As Figure 4a and 4b indicate, there are considerable cross-sectional differences in mine gross employment flows. Previous studies by Davis and Haltiwanger (1992) and Davis, Haltiwanger, and Schuh (1996) find that most plant-level changes in employment are due to idiosyncratic shocks. That is, neither plant characteristics nor industry or economy-wide shocks can explain a substantial component of the differences in plant or firm growth. To examine this issue in the case of coal mines, we perform a simple decomposition of the variation in growth attributable to mine specific time invariant factors (fixed effects), common time shocks (time effects), and idiosyncratic effects (error).

The fixed effects will pick up mine characteristics that are time invariant including location and mine type. The time effects will pick up aggregate industry shocks. Finally, we examine two measures of mine employment growth—a measure of the growth calculated as in expression (1) above and an absolute value of growth. In the latter case, we are trying to account for changes in employment at the mine level but not the direction of the change. Finally, it is important to recognize that this approach is entirely non-structural in nature; we simply estimate a fixed effects model to examine the amount variation explained by the various components.⁵⁴

The results are of the decomposition are presented in Table 4. The decomposition shows that common time shocks explain little of the mine-level differences in employment growth. Common time effects (in Column 2) can explain only 2.4% of the overall variation in mine-level employment flows. Similarly, when also controlling for fixed mine effects, only 13% of the variation in mine-level employment flows can be explained by fixed mine effects and time effects. In fact, a joint statistical test of the significance of the mine fixed effects cannot reject the null hypothesis of no joint effect; the F-value is 0.2. Alternatively, when one examines the absolute value of growth (Column 3), mine fixed effects along with time effects can explain about 38% of the variation across mines, though as expected common shocks explain little of the variation in absolute employment flows. The fact that mine effects are more important in

⁵⁴ In this case, we estimate a basic effects model of the kind $y_{it} = a_i + b_t + e_{it}$ where the a_i 's represents the mine fixed effects, b_t represents the common shocks, and e_{it} represents

explaining differences in the absolute employment flow suggests that some mines have a tendency to have high levels of job creation and destruction, while other mines may have a tendency to have relatively low rates of job creation and job destruction. Such differences may be due to differences in the costs of hiring and laying off workers or may be due to differences in the nature of demand shocks across the various locations. However, it is the case that some mines have persistently higher rates of employment flows than other mines.

4.4 Conclusion

This paper documents the patterns of employment flows in the coal mining sector. The contributions of the paper are threefold. First, the paper develops a new data series on employment flows for the coal mining sector. Previous studies have focused predominately on the manufacturing sector or on particular states. Coal mining has features that generate substantially different patterns than manufacturing. In particular, mine openings and closings play a dominant role in job creation and job destruction; this is not the case in manufacturing. Second, temporary closings and mine reopenings are important sources of employment flows. This suggests that either local shocks that mines receive are either large or that costs of temporary shutdown and re-startup are relatively low. Finally, while there are important differences between mining and manufacturing, there are also some striking similarities. The level of job creation and job destruction greatly exceeds the net changes, as is true in manufacturing.

idiosyncratic effects.

In addition, employment reallocation is counter-cyclical, as it is in manufacturing—also consistent with the manufacturing evidence.

Table 1. Summary Statistics on Number of Mines, Total Employment, and Mine Size

Year	Number of Mines	Total Employment in Active Mines	Average Employment per Mine	Median Employment per Mine
1975	2,355	155,023	65.8	21
1980	2,788	184,038	66.0	20
1985	2,336	138,658	59.4	19
1990	1,969	109,134	55.4	20
1995	1,323	77,868	58.9	24

Source: Author's tabulation from MSHA mine-level data.

Table 2. Employment Flows in Coal Mining: 1973-1996

Year	Job Creation Rate	Job Destruction Rate	Net Employment Growth	Employment Reallocation Rate	Excess Reallocation Rate
1974	0.201	0.125	0.076	0.326	0.250
1975	0.249	0.080	0.170	0.329	0.159
1976	0.224	0.125	0.100	0.349	0.249
1977	0.194	0.123	0.070	0.317	0.247
1978	0.136	0.412	-0.276	0.549	0.273
1979	0.519	0.133	0.386	0.652	0.266
1980	0.193	0.247	-0.054	0.439	0.385
1981	0.159	0.142	0.017	0.301	0.285
1982	0.168	0.176	-0.007	0.344	0.337
1983	0.103	0.352	-0.250	0.455	0.206
1984	0.189	0.105	0.084	0.294	0.210
1985	0.134	0.231	-0.097	0.364	0.267
1986	0.149	0.214	-0.066	0.363	0.298
1987	0.122	0.202	-0.080	0.324	0.244
1988	0.139	0.177	-0.038	0.315	0.277
1989	0.134	0.180	-0.046	0.314	0.269
1990	0.158	0.151	0.007	0.309	0.302
1991	0.127	0.186	-0.058	0.313	0.255
1992	0.126	0.210	-0.084	0.337	0.253
1993	0.121	0.244	-0.123	0.365	0.242
1994	0.196	0.186	0.010	0.383	0.372
1995	0.160	0.218	-0.059	0.378	0.319
1996	0.128	0.242	-0.114	0.370	0.256

Table 3. Temporary Mine Closings and Mine Re-openings

Year	Percent of Closing Mines that will Re-Open	Percent of Opening Mines that were Previously in Operation
1975	34.0	N/A
1980	27.6	37.1
1985	28.4	41.4
1990	19.6	37.4
1995	N/A	45.1

Table 4. Across Mine Variation in Employment Flows

Effect	Mine-Level Employment Growth: Percent of Explained Variation	Absolute Value of Mine-Level Employment Growth: Percent of Explained Variation
Time Effects	2.4%	0.1%
Time and Mine Fixed Effects	13.1%	38.2%

Note: The analysis of variance includes 64,547 observations representing 11,691 individual mines.

Figure 1. Job Creation and Job Destruction in Coal Mining: 1973-1996




Figure 2. Employment Reallocation




Figure 3. Employment Flows – Continuing Mines, Openings and Closings
1973 to 1996




Figure 4a: Distribution of Job Creation at Mines





Figure 4b: Distribution of Job Destruction at Mines



5. Conclusion

This dissertation examines microeconomic adjustment dynamics in the U.S. coal mining industry. Studying these sorts of adjustment dynamics can be instructive to those who are interested in understanding the engines that drive industry dynamics and the process of creative destruction. Specifically, this dissertation studies three instruments available to managers (or owners) of coal mines to adjust to changing economic environments.

First, as a mechanism to correct for declining productivity performance, I examine the use of acquisitions. The idea is that the creative destruction process can manifest itself through turnover in the control of an establishment. In this case, I develop a simple stochastic dynamic programming model that describes the problem of matching owners to coal mines. By observing productivity over time, owners learn about their relative productivity (a signal as to the quality of the match) and then decide whether to continue as owners or not. This model makes two empirically testable hypotheses that, if true, would corroborate the notion that acquisitions are corrective forces in the coal mining industry. These two predictions are as follows: i. Acquired mines ought to exhibit lower productivity prior to having been acquired than mines that were not acquired, and ii. Extant acquired mines ought to exhibit productivity gains after having been acquired.

In reduced form regressions, I find that both hypotheses are supported by the data. That is, I estimate that in the pre-acquisition period, mines targeted for

acquisition are between 5% and 12% less productive than mines that were not targets. Additionally, I find that surviving acquired mines have significantly faster productivity growth than the non-acquired mines in the three periods immediately after the acquisition event periods. One should be careful when thinking about this last result as the estimates likely suffer from sample selection bias since the data used to estimate the growth equations are conditioned on survival. Failure probit models suggest that this could be true since I find that having been acquired is a significantly positive determinant of the likelihood of failing. However, the faster post-acquisition productivity growth is highest in the year immediate following the acquisition—a period where sample selection bias would be least severe.

Second, I study the determinants of coal mine failure—recalling that one of the mechanisms through which the creative destruction process works is establishment turnover. Relying on a theoretical structure that incorporates market demand conditions, productivity, and resource depletion as state variables, I present the hypothesis that demand conditions and productivity ought to lower hazard rates while resource depletion ought to raise the hazard rates. The basis for the first assertion comes from standard models of firm dynamics that argue that high realizations of prices and productivity raise current and expected profitability and hence lower the failure probability. However, the second assertion, that resource depletion would raise the hazard, stems from simply

noting that because coal is non-renewable, every unit produced must lower expected future profitability, *ceteris paribus*.

Taken to the data, the empirical problem is to take the age of a coal mine and disentangle the learning effects from the resource depletion effects. In Cox proportional hazard models, controlling for market demand conditions and mine heterogeneity, I argue that including both a measure of cross-sectional productivity differences and a measure of a mine's age will disentangle these two effects. The cross-sectional productivity controls would control the learning effects, and the age will capture the depletion effects—as long as both are included in the same specification. I estimate that prices have a very slight tendency to raise the hazard; the magnitude, though statistically significant, is so slight that it is unclear that the percent change in oil prices have much material effect on the baseline hazard. I further find that productivity seems to lower the hazard—consistent with most models of industry learning. Finally, I find that the age of the mine significantly raises the baseline hazard. Altogether, these results support the theoretical structure developed in chapter three.

Finally, in an entirely empirical chapter, I study the employment adjustment dynamics in coal mining and compare the rates of job creation, job destruction, and job reallocation to known empirical regularities in manufacturing industries. Four salient features emerge from that data. First, the rates of job creation and job destruction (hence, reallocation) are substantially higher in coal mining than in manufacturing. Second, the large rates of job creation and job

destruction are attributable in part to mine openings and closings. Third, a large portion of the gross employment flows is attributable to the temporary nature of a large portion of mine openings and closings. Finally, similar to the evidence from manufacturing studies, job destruction is counter-cyclical with respect to the business cycle and also is more volatile over time than the job creation rate. All of these findings illustrate how the creative destruction process manifests itself in the employment adjustment mechanism.

References

Chapter One

Acs, Z.J. and Audretsch, D.B. (1990) *Innovation and Small Firms*. Cambridge, MA: MIT Press.

Baily, M.N., Bartlesman, E.J., and Haltiwanger, J.C. (1996) "Downsizing and Productivity Growth: Myth or Reality?" *Small Business Economics*, v. 8, no. 4, pp. 259-278.

Bailey, M.N., Hulten, C., and Campbell, D. (1992) "Productivity Dynamics in U.S. Manufacturing Plants," *Brookings Papers on Economic Activity: Microeconomics*, pp. 187-267.

Baldwin, J.R. (1995) *The Dynamics of Industrial Competition*. Cambridge: Cambridge University Press.

Caves, R.E. (1998) "Industrial Organization and New Findings on the Turnover and Mobility of Firms," *Journal of Economic Literature*, v. 36, no. 4, pp. 1947-1982.

Dunne, T., Roberts, M.J., and Samuelson, L. (1989) "The Growth and Failure of U.S. Manufacturing Plants." *Quarterly Journal of Economics*, v. 104, no. 4, pp.671-698.

Ericson, R. and Pakes, A. (1995) "Markov-Perfect Industry Dynamics: A Framework for Empirical Work." *Review of Economic Studies*, v. 62, no. 1, pp. 53-82.

Griliches, Z. and Regev, H. (1992) "Productivity and Firm Turnover in Israeli Industry: 1979-1988," NBER Working Paper no. 4059.

Haltiwanger, J.C. (1997) "Measuring and Analyzing Aggregate Fluctuations: The Importance of Building from Microeconomic Evidence," *Federal Reserve Bank of St. Louis Review*, May/June, pp. 55-78.

Hopenhayn, H.A. (1992) "Entry, Exit, and Firm Dynamics in Long Run Equilibrium." *Econometrica*, v. 60, no. 5, pp. 1127-1150.

Jovanovic, B. (1979) "Job Matching and the Theory of Turnover." *Journal of Political Economy*, v. 87, pp. 972-990.

Jovanovic, B. (1982) "Selection and the Evolution of Industry." *Econometrica*, v. 50, no. 3, pp. 649-670.

Lichtenberg, F. R. and Siegel, D. (1987) "Productivity and Changes in Ownership of Manufacturing Plants." *Brookings Papers on Economic Activity: 1987*, pp. 643-673.

Maksimovic, V. and Phillips, G. (1999) "The Market for Corporate Assets: Who Engages in Mergers and Asset Sales and Are There Efficiency Gains?," Mimeo, Robert H. Smith College of Business, University of Maryland.

Olley, G. S. and Pakes, A. (1996) "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica*, v. 64, n. 6, pp. 1263-1297.

Roberts, M.J. and Tybout, J.R. (1996) *Industrial Evolution in Developing Countries: Micro Patterns of Turnover, Productivity, and Market Structure*. In M.J. Roberts and J.R. Tybout, eds., New York: Oxford University Press.

Schumpeter, J.A. (1942) *Capitalism, Socialism, and Democracy*. New York: Harper and Brothers.

Chapter Two

Bailey, M.N., Hulten, C., and Campbell, D. (1992) "Productivity Dynamics in U.S. Manufacturing Plants," *Brookings Papers on Economic Activity: Microeconomics*, pp. 187-267.

Baumol, W.J. (1967) *Business Behavior, Value, and Growth*. New York: Harcourt, Brace, and World.

Berndt, E.R. and Ellerman, A.D. (1996) "Investment, Productivity, and Capacity in U.S. Coal Mining: Micro vs. Macro Perspectives." Working Paper, MIT.

Brandenburger, A. and Polak, B. (1996) "When Managers Cover Their Posteriors: Making Decisions the Market Wants to See." *RAND Journal of Economics*, v. 27, pp. 523-541.

Brown, C. and Medoff, J.L. (1987) "The Impact of Firm Acquisitions on Labor." NBER Working Paper.

Butler, J.S. and Robert Moffitt. (1982) "A Computationally Efficient Quadrature Procedure for the One-Factor Multinomial Probit Model." *Econometrica*, v. 50, no. 3 pp. 761-764.

Dixit, A.K. and Pindyck, R.S. (1994) *Investment Under Uncertainty*. Princeton University Press: Princeton, NJ.

Dunne, T., Roberts, M.J., and Samuelson, L. (1989) "The Growth and Failure of U.S. Manufacturing Plants." *Quarterly Journal of Economics*, v. 104, no. 4, pp.671-698.

Ellerman, A.D., Stoker, T.M., and Berndt, E.R. (1998) "Sources of Productivity Growth in the American Coal Industry," Working Paper, MIT.

Energy Information Administration. (1993) *The Changing Structure of the U.S. Coal Industry: An Update*. Washington, D.C.: DOE/EIA-0513(93).

Jarrell, G.A., Brickley, J.A., and Netter, J.M. (1988) "The Market for Corporate Control: The Empirical Evidence Since 1980," *Journal of Economic Perspectives*, v. 2, pp. 49-68.

Jensen, J. B., and McGuckin, R.H. (1997) "Firm Performance and Evolution: Empirical Regularities in the U.S. Microdata," *Industrial and Corporate Change*, v. 6, no. 1, pp. 25-47.

Jensen, M.C. (1988) "Takeovers: Their Causes and Consequences." *Journal of Economic Perspectives*, v. 2, pp. 21-48.

Jensen, M.C. and Ruback, R.S. (1983) "The Market for Corporate Control: The Scientific Evidence." *Journal of Financial Economics*, v.11, pp. 5-50.

Jorgenson, D.W., Gollop, F., and Fraumeni, B. (1987) *Productivity and U.S. Economic Growth*. Cambridge: Harvard University Press.

Joskow, P. (1987) "Contract Duration and Relationship-Specific Investments: Empirical Evidence from Coal Markets." *American Economic Review*, v. 77, n. 1, pp. 168-85.

Jovanovic, B. (1979) "Job Matching and the Theory of Turnover." *Journal of Political Economy*, v. 87, pp. 972-990.

Lichtenberg, F. R. and Siegel, D. (1987) "Productivity and Changes in Ownership of Manufacturing Plants." *Brookings Papers on Economic Activity: 1987*, pp. 643-673.

Lichtenberg, F. R. and Siegel, D. (1990) "The Effects of Leveraged Buyouts on Productivity and Related Aspects of Firm Behavior." *Journal of Financial Economics*, v. 27, pp. 165-194.

Lichtenberg, F. R. and Siegel, D. (1992a) "Productivity and Changes in Ownership of Manufacturing Plants." In F.R. Lichtenberg, ed., *Corporate Takeovers and Productivity*. Cambridge: MIT Press.

Lichtenberg, F. R. and Siegel, D. (1992b) "Leveraged Buyouts." In F.R. Lichtenberg, ed., *Corporate Takeovers and Productivity*. Cambridge: MIT Press.

Maksimovic, V. and Phillips, G. (1999) "The Market for Corporate Assets: Who Engages in Mergers and Asset Sales and Are There Efficiency Gains?," Mimeo, Robert H. Smith College of Business, University of Maryland.

Manne, H.G. (1965) "Mergers and the Market for Corporate Control." *Journal of Political Economy*, v. 73, pp. 110-120.

Matsusaka, J.G. (1993a) "Target Profits and Managerial Discipline During the Conglomerate Merger Wave." *Journal of Industrial Economics*, v. 41, pp. 179-189.

Matsusaka, J.G. (1993b) "Takeover Motives During the Conglomerate Merger Wave." *RAND Journal of Economics*, v. 24, pp. 357-379.

McGuckin, R.H. and Nguyen, S.V. (1995) "On Productivity and Plant Ownership Change: Evidence from the Longitudinal Research Database." *RAND Journal of Economics*, v. 26, pp. 257-276.

McGuckin, R.H. and Nguyen, S.V., and Reznick, A. (1995) "The Impact of Ownership Change on Employment, Wages, and Labor Productivity in U.S. Manufacturing 1977-87," Forthcoming in *Labor Statistics Measurement Issue*, eds. John Haltiwanger, Marilyn Manser, and Robert Topel. Chicago: NBER.

Meade, J.E. (1968) "Is the New Industrial State Inevitable?" *Economic Journal*, v. 78, pp. 372-392.

Morck, R. Schleifer, A., and Vishny, R. (1990) "Do Managerial Objectives Drive Bad Acquisitions?" *Journal of Finance*, v. 45, n. 1, pp. 31-48.

Mueller, D.C. (1969) "A Theory of Conglomerate Mergers." *Quarterly Journal of Economics*, v. 83, pp. 643-659.

Mueller, D.C. (1993) "Mergers: Theory and Evidence." In G. Mussanti, ed., *Mergers, Markets, and Public Policy*. Dordrecht: Kluwer.

Olley, G. S. and Pakes, A. (1996) "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica*, v. 64, n. 6, pp. 1263-1297.

Pakes, A. and Ericson, R. (1998) "Empirical Implications of Alternative Models of Firm Dynamics." *Journal of Economic Theory*, v. 79, pp. 1-45.

Ravenscraft, D.J. and Scherer, F.M. (1989) "The Profitability of Mergers." *International Journal of Industrial Organization*, v. 7, pp. 101-116.

Ritter, J.A. and Taylor, L.J. (1999) "Are Employees Stakeholders? Corporate Finance meets the Agency Problem." In Carrie Leana and Denise Rousseau, eds., *Relational Wealth*, New York: Oxford University Press, forthcoming.

Roll, R. (1986) "The Hubris Hypothesis of Corporate Takeovers." *Journal of Business*, v. 59, pp. 197-216.

Schleifer, A. and Vishny, R.W. (1989) "Management Entrenchment: The Case of Manager-Specific Investments." *Journal of Financial Economics*, v. 25, pp. 123-139.

Stigler, G.J. (1950) "Monopoly and Oligopoly Behavior." *American Economic Review*, v. 40, pp. 32-34.

Schleifer, A. and Summers, L. (1987) "Hostile Takeovers as Breaches of Trust." NBER Working Paper.

Chapter Three

Allison, P.D. (1984) *Event History Analysis: Regression for Longitudinal Event Data*. London: Sage Publications, Series/Number 07-046.

Bailey, M.N., Hulten, C. and Campbell, D. (1992) "Productivity Dynamics in U.S. Manufacturing Plants." *Brookings Papers on Economic Activity: Microeconomics*. pp. 187-267.

Bernard, A. and Jensen, J.B. (1997) "Why Some Firms Export," Mimeo, H. John Heinz III, School of Public Policy and Management, Carnegie Mellon University.

Berndt, E.R. and Ellerman, A.D. (1998) "Investment, Productivity, and Capacity in U.S. Coal Mining: Micro vs. Macro Perspectives." Working Paper, MIT.

Deily, M.E. (1991) "Exit Strategies and Plant-Closing Decisions: The Case of Steel." *RAND Journal of Economics*, v. 22, no. 2, pp. 250-263.

Dunne, T., Roberts, M.J., and Samuelson, L. (1989) "The Growth and Failure of U.S. Manufacturing Plants." *Quarterly Journal of Economics*, v. 104, no. 4, pp. 671-698.

Energy Information Administration (1995). *Coal Data: A Reference*, DOE/EIA-0064(93), Distribution Category UC-98.

Ericson, R. and Pakes, A. (1995) "Markov-Perfect Industry Dynamics: A Framework for Empirical Work." *Review of Economic Studies*, v. 62, no. 1, pp. 53-82.

Heckman, J. (1981a) "Statistical Models for Discrete Panel Data." In C. Manski and D. McFadden, eds., *The Structural Analysis of Discrete Data*, Cambridge, MA: MIT Press. pp. 114-178.

Heckman, J. (1981b) "The Incidental Parameters Problem and the Problem of Initial Conditions in Estimating Discrete Time-Discrete Data Stochastic Processes." In C. Manski and D. McFadden, eds., *The Structural Analysis of Discrete Data*, Cambridge, MA: MIT Press. pp. 179-195.

Heckman J. and Singer, B. (1984a) "Econometric Duration Analysis." *Journal of Econometrics*, v. 24, no. 1-2, pp. 63-132.

Heckman, J. and Singer, B. (1984b) "A Method for Minimizing the Impact of Distributional Assumptions in Econometric Models of Duration Data." *Econometrica*, v. 52, no. 2, pp. 271-320.

Heckman, J. and Singer, B. (1984c) "The Identifiability of the Proportional Hazard Model." *Review of Economic Studies*, v. 51, no. 2, pp. 231-241.

Heckman, J. and Walker, J.R. (198?)

Hopenhayn, H.A. (1992) "Entry, Exit, and Firm Dynamics in Long Run Equilibrium." *Econometrica*, v. 60, no. 5, pp. 1127-1150.

Jensen, J.B. and McGuckin, R.H. (1997) "Firm Performance and Evolution: Empirical Regularities in U.S. Microdata." *Industrial and Corporate Change*, v. 6, no. 1, pp. 25-47.

Joskow, P.L. (1987) "Contract Duration and Relationship-Specific Investments: Empirical Evidence from Coal Markets." *American Economic Review*, v. 77, no. 1, pp. 168-185.

Jovanovic, B. (1982) "Selection and the Evolution of Industry." *Econometrica*, v. 50, no. 3, pp. 649-670.

Kalbfleisch, J.D. and Prentice, R.L. (1980) *The Statistical Analysis of Failure Time Data*. New York: John Wiley and Sons, Inc.

Krautkraemer, J.A. (1998) "Nonrenewable Resource Scarcity." *Journal of Economic Literature*, v. 36, pp. 2065-2107.

Lambson, V.E. (1992) "Competitive Profits in the Long Run." *Review of Economic Studies*, v. 59, no. 1, pp. 125-142.

Lancaster, T. (1990) *The Econometric Analysis of Transitions Data*, New York, NY: Cambridge University Press.

Merrell, D.R. (1999) "Productivity and Acquisitions in U.S. Coal Mining." Mimeo, H. John Heinz III School of Public Policy and Management, Carnegie Mellon University.

Meyer, B. (1990,) "Unemployment Insurance and Unemployment Spells" *Econometrica*, v. 58, no. 4, pp. 757-782.

Pakes, A. and Ericson, R. (1998) "Empirical Implications of Alternative Models of Firm Dynamics." *Journal of Economic Theory*, v. 79, no. 1, pp. 1-45.

Pakes, A. and McGuire, P. (1994) "Computing Markov-Perfect Nash Equilibria: Numerical Implications of a Dynamic Differentiated Product Model." *RAND Journal of Economics*, v. 25, no. 4, pp. 555-589.

Pierce, W.S. (1996) *Economics of the Energy Industries, Second Edition*. Westport, Connecticut: Praeger Publishers.

Pindyck, R.S. (1980) "Uncertainty and Exhaustible Resource Markets." *Journal of Political Economy*, v. 88, no. 6, pp. 1203-1225.

Roberts, M.J. and Tybout, J.R. (1997) "The Decision to Export in Colombia: An Empirical Model of Entry with Sunk Costs." *American Economic Review*, v. 87, no. 4, pp. 545-564

Trussell, J. and Richards, T. (1985) "Correcting for Unmeasured Heterogeneity in Hazard Models Using the Heckman-Singer Procedure." In N. Tuma, ed., *Sociological Methodology*, Jossey-Bass Publishers, pp. 242-269.

Yamaguchi, K. (1991) *Event History Analysis*. London: Sage Publications Inc.

Chapter Four

Aw, B-Y., Chen, T-J., and Roberts, M.J. (1997) "Firm-Level Evidence on Productivity Differentials, Turnover, and Exports in Taiwanese Manufacturing," mimeo.

Baily, M.N., Hulten, C., and Campbell, D. (1992) "Productivity Dynamics in Manufacturing Plants," *Brookings Papers on Economic Activity, Microeconomics*, 1992, 187-267.

Berndt, E.R. and Ellerman, A.D. (1998) "Investment, Productivity, and Capacity in U.S. Coal Mining: Micro vs. Macro Perspectives." Working Paper, MIT.

Blanchard, O. and Diamond, P. (1990) "The Cyclical Behavior of Gross Flows of Workers in the U.S.," *Brookings Papers on Economic Activity* (1990:2), 85-155.

Caballero, R. and Engel, E. (1993) "Microeconomic Adjustment Hazards and Aggregate Dynamics," *Quarterly Journal of Economics*, 1993.

Caballero, R., Engel, E., and Haltiwanger, J.C. (1997) "Aggregate Employment Dynamics: Building from Microeconomic Evidence," *American Economic Review*, 87(1), March, 115-137.

Caballero, R. and Hammour, M. (1994) "The Cleansing Effect of Recessions," *American Economic Review*, 84, 1350-1368.

Caballero, R. and Hammour, M. (1997) "The Macroeconomics of Specificity," Working Paper.

Campbell, J.R. and Fisher, J. (1997) "Understanding Aggregate Job Flows," *Economic Perspectives*, 21(5), Federal Reserve Bank of Chicago (September/October) 19-37.

Campbell, J.R. and Kuttner, K. (1996) "Macroeconomic Effects of Employment Reallocation," *Carnegie-Rochester Conference Series on Public Policy*, 44(0), June, 87-116.

Davis, S.J. and Haltiwanger, J.C. (1992) "Gross Job Creation, Gross Job Destruction, and Employment Reallocation," *Quarterly Journal of Economics*, 107(3), 819-63.

Davis, S.J., Haltiwanger, J.C., and Schuh, S. (1996) *Job Creation and Destruction*, Cambridge, MA: MIT Press.

Dunne, T., Roberts, M.J., and Samuelson, L. (1989a) "The Growth and Failure of U.S. Manufacturing Plants," *Quarterly Journal of Economics*, 104(4), 671-98.

Dunne, T., Roberts, M.J., and Samuelson, L. (1989b) "Plant turnover and Gross Employment Flows in the U.S. Manufacturing Sector," *Journal of Labor Economics*, 7(1), 48-71.

Energy Information Administration (1995). *Coal Data: A Reference*, DOE/EIA-0064(93), Distribution Category UC-98.

Evans, D. (1987a) "Test of Alternative Theories of Firm Growth." *Journal of Political Economy*, 95(4), 657-74.

Foster, L., Haltiwanger, J.C., and Krizan, C.J. (1998) "Aggregate Productivity Growth: Lessons from Microeconomic Evidence," Working Paper, February.

Hopenhayn, H. and Rogerson, R. (1993) "Job Turnover and Policy Evaluation: A General Equilibrium Analysis," *Journal of Political Economy*, 101(5), 915-38.

Merrell, D.R. (1999) "Implications of Resource Exhaustion on Exit Patterns in U.S. Coal Mining," Mimeo, H. John Heinz III School of Public Policy and Management, Carnegie Mellon University.

Mortensen, D.T. and Pissarides, C. (1994) "Job Creation and Job Destruction in the Theory of Unemployment," *Review of Economic Studies*, 61, 397-415.

Olley, S. and Pakes, A. (1996) "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*, 64(6), November, 1263-1297.

Troske, K. (1996) "The Dynamic Adjustment Process of Firm Entry and Exit in Manufacturing and Finance, Insurance and Real Estate," *Journal of Law and Economics*. 39(2), October, 705-735.