UNIVERSITY OF OKLAHOMA

GRADUATE COLLEGE

ESSAYS IN FINANCIAL MARKETS

A DISSERTATION

SUBMITTED TO THE GRADUATE FACULTY

in partial fulfillment of the requirements for the

Degree of

DOCTOR OF PHILOSOPHY

By

WENBIN CAO
Norman, Oklahoma
2016

ESSAYS IN FINANCIAL MARKETS

A DISSERTATION APPROVED FOR THE
MICHAEL F. PRICE COLLEGE OF BUSINESS

BY

_____
Dr. Pradeep Yadav, Chair

_____
Dr. Scott Linn

_____
Dr. Tor-Erik Bakke

_____
Dr. Hamed Mahmudi

_____
Dr. Robert Dauffenbach

I dedicate this work to my wife, Xiaoman Duan; without her love, support, and understanding this would not have been possible. I dedicate this work to my parents, Xuefeng Cao and Lixia Wang, who are my biggest advocates and cheerleaders. I dedicate this work to our daughter, Julia Cao, who illuminates our lives.

Acknowledgements

I thank Pradeep Yadav, my advisor, for his encouragement, mentorship, and guidance throughout my doctoral program. I am grateful to Scott Linn, Vahap Uysal, Louis Ederington, Xiaoman Duan, and Scott Guernsey for the opportunity to collaborate on research over the past five years. I am grateful to Scott Linn, Tor-Erik Bakke, Hamed Mahmudi, and Robert Dauffenbach for their commitment and support as members of my dissertation committee.

Table of Contents

List of Tables

List of Figures

Abstract

This dissertation is a collection of three essays that analyze the impact of economic uncertainty on financial market activities and evaluate alternative quantitative models for economic uncertainty based on financial asset prices. Chapter 1 is motivated by the fact that major economic and political shocks, such as the Cuban missile crisis, the 9/11 terrorist attacks, and the 2008 financial crisis, trigger spikes in market-wide uncertainty. It develops a dynamic trading model to analyze the impacts of these uncertainty shocks on the behaviors of market liquidity and shows how such impacts differ from those caused by shocks to economic conditions. According to my model, an uncertainty shock triggers a temporary decline in market liquidity, because an uncertainty shock introduces ambiguity and learning can resolve this ambiguity. Meanwhile, incorporating the notion of time-varying uncertainty aversion, my model implies that a shock to economic condition generates a persistent decline in market liquidity, since learning does not affect uncertainty aversion. My VAR estimations using monthly US data for 1962--2013 lend support to my model implications. An uncertainty shock generates a rapid drop and rebound in overall stock market liquidity for around five months on average, while a shock to economic condition leads to a persistent decline in market liquidity for up to a year.

Chapter 2 examines the ability of five basis alternative option pricing models to price the early exercise premium (EEP) in American put prices: Black-Scholes model, Heston (1993) stochastic volatility model, and three jump pricing models – Merton (1976), Madan et al. (1998), and Carr and Wu (2003). After duly accounting for the market implied value of the Fleming and Whaley (1993) wild card option, we find that jump

models perform best in pricing observed EEP. Importantly, all models consistently and significantly underprice observed EEP, where this underpricing is more pronounced for short term in-the-money EEP. We argue and empirically demonstrate that trading costs in the option market generate a significant EEP by incentivizing and rewarding early exercise of American options that would alternatively have been "sold" in the market.

Chapter 3 examines the frequency and character of price jumps in front month oil and natural gas futures prices, where prices are sampled every five seconds over the period 2006-2014. We find that an infinite activity jump diffusion process describes crude oil and natural gas futures returns combined with a process involving much larger but less frequent jumps. We further find that jumps account for respectively 36 and 41 percent of the realized variances of the crude oil and the natural gas returns.

Chapter 1:

Learning Under Ambiguity and Market Liquidity

I. Introduction

Market-wide uncertainty fluctuates over time and appears to increase dramatically after major economic and political shocks, such as the Cuban missile crisis, the 9/11 terrorist attacks, and the 2008 financial crisis. These uncertainty shocks (second-moment shocks) adversely impact economic activities and are major concerns of policy makers.[1] These shocks by nature can make market participants question whether the economy will stay on the same track and further consider multiple possible trajectories of the future economy. Equivalently, these shocks introduce ambiguity or Knightian uncertainty in the economy and in the financial market. If market participants are ambiguity averse, as suggested by the Ellsberg Paradox (Ellsberg (1961)), these uncertainty shocks will cause drastic declines in market liquidity. This idea has sparked a strand of studies that attribute the disappearing market liquidity in the recent global financial crisis as originating from the ambiguity being introduced and market participants' ambiguity aversion.[2] However there is an important question remains unanswered: How to analyze the impact of these uncertainty shocks and that of those economic condition shocks (first-moment shocks) on the dynamics of market liquidity? Which shocks, uncertainty shocks or economic

[1] Bloom (2009) demonstrates the adverse impact of such uncertainty shocks on investments and unemployment and highlights a distinction in their impacts between these uncertainty shocks (second-moment shocks) and level shocks (first-moment shocks).
[2] Notable examples include Routledge and Zin (2009) and Easley and O'Hara (2010b), among others.

conditions shocks, have more persistent impacts on market liquidity? Moreover, what do we actually observe?

To address this question, I first focus on the informational nature of ambiguity and the role of learning in resolving ambiguity (Epstein and Schneider (2007, 2008); Klibanoff, Marinacci, and Mukerji (2009)). Ambiguity, by nature, results from agents' lack of relevant information to be confident in a specific probabilistic model for a random variable, such as asset payoffs. This is particularly prevalent after the onset of an uncertainty shock. In subsequent periods, more relevant information starts to accumulate. By updating their beliefs, agents become more confident in a specific model, and their perceptions of ambiguity gradually decline. Consequently, we should expect in general that the effect of ambiguity is relatively short-run due to learning.

Second, I consider the notion of "time-varying risk aversion" (Campbell and Cochrane (1999); Barberis, Huang, and Santos (2001)) to analyze the impact of declining economic conditions. This notion relates to the insight that market participants are more averse to uncertainty (risk and ambiguity) in economic downturns. Since learning only affects agents' perception of uncertainty but not their aversions to uncertainty, we should expect in general that the impact of a change in uncertainty aversion is more persistent. Note that although it is common to see that economic uncertainty can elevate in economic downturns, it is reasonable to examine both in separate channels, as noted in Bloom (2009) that shocks to economic conditions can leave economic uncertainty intact and vice versa.

Building upon these two foundations of thoughts, I consider a multi-period market setting in which the preference of a representative market maker over his ambiguous

terminal wealth is portrayed by the smooth ambiguity model (SAM) of Klibanoff, Marinacci, and Mukerji (2005). Although ambiguity is also modeled as a multiple-prior belief in SAM, the same as that in the maxmin preference of Gilboa and Schmeidler (1989), the SAM allows the explicit modeling of the degree of ambiguity aversion, which is infinite in the maxmin preference, according to Klibanoff, Marinacci, and Mukerji (2005). Within the general SAM framework, I employ the ambiguity robust Arrow-Pratt preference (ARAP) introduced by Maccheroni, Marinacci, and Ruffino (2013). This ARAP preference approximates the SAM in the same way as the classic Arrow-Pratt preference approximates the expected utility model and offers flexibilities in modeling ambiguity aversion as well as analytical tractability.

To model the information structure, I assume the representative market maker receives a noisy signal about the (ambiguous) future liquidation value of the asset whose liquidity is supplied by him in each period. In this case, the information structure evolves consistently with the insights in Epstein and Schneider (2007, 2008) and Klibanoff, Marinacci, and Mukerji (2009) that learning affects agents' beliefs about the ambiguous future value by affecting agents' beliefs about the validity of alternative models. In another word, agents initially consider the ambiguous future value as depicted by a range of alternative models and assign a prior weights to the models. After incorporating more relevant information, their beliefs will converge to a particular model (the "right" model), and this is the exact nature of how learning resolves ambiguity.[3]

Supported by these pillars, my model predicts distinct behaviors of market liquidity (price impact in specific) in subsequent periods after an uncertainty shock and

[3] Klibanoff, Marinacci, and Mukerji (2009) prove that this convergence will occur if the dimension of the parameter space is finite. The setting being considered here satisfy the convergence condition.

after an economic condition shock. To be compatible with the unified theoretical framework of market liquidity being analyzed in Vayanos and Wang (2012), I consider price impact, reflecting the impact of a one-unit trade on transaction price. According to my model, price impact will increase temporarily after a pure uncertainty shock that introduces ambiguity. When ambiguity increases initially, a one-unit trade will have a larger impact on transaction price, reflecting a higher required compensation for the market maker to supply liquidity. Due to learning, the market maker's perception of ambiguity declines gradually, translating into a lower required compensation hence lower price impact. In this sense, my model features a market liquidity cycle caused by uncertainty shocks. At the onset of an uncertainty shock, market liquidity dries up due to the spike in ambiguity. In subsequent periods, market liquidity increases and converges to a steady-state level, subject only to ongoing fluctuations (risk), until another uncertainty shock.

In contrast, price impact will increase persistently after a pure economic condition shock, which increases market makers' uncertainty aversion. When the market maker becomes more averse to uncertainty, he will require a higher compensation for supplying liquidity even facing the same level of risk or ambiguity. Because learning does not affect his uncertainty aversion, an increase in the market maker's uncertainty aversion will have a longer-term impact on price impact.

Although learning is the key to understanding the distinction in the behaviors of market liquidity in subsequent periods following an uncertainty shock and an economic condition shock, my model based on learning under ambiguity is free from a common concern faced by models based on learning under pure risk. Models with learning under

pure risk can indeed generate a convergence result (e.g., Vives (1995)). However, they feature a type of long-run equilibria with no uncertainty and perfect liquidity, which is too ideal to observe in reality. In my model, the convergence to the steady-state equilibrium can be viewed as a process of resolving the impact of model uncertainty. In the short run, market liquidity is affected by both the uncertainty about the "correct" model and the uncertainty within each possibly "correct" model. In the long run, market liquidity is only affect by the uncertainty within the correct model.

To empirically assess my model implications, I estimate a range of vector autoregressions (VAR) using detrended monthly US data from 1962--2013. The main response variable is overall stock market liquidity.[4] In particular, I construct the Amihud (2002) measure at a monthly frequency, using daily stock return and volume data from the CRSP. I employ the Amihud measure because my theoretical model features price impact, and Goyenko, Holden, and Trzcinka (2009) demonstrate that the Amihud measure is the best measure for price impact based on low-frequency data. I aggregate the constructed market liquidity series for individual stocks to get the overall market liquidity series (e.g., Naes, Skjeltorp, and Odegaard (2011)).

The impulse variables are classified into two categories: economic condition series and economic uncertainty shock series. I consider the S&P 500 index level and the Federal Funds Rates (FFR) as proxies for economic conditions. The economic uncertainty shock series is a 0/1 indicator series, constructed following Bloom (2009), who demonstrates that using this indicator series yields better identifications of the uncertainty

[4] I study overall stock market liquidity because individual stock liquidity is subject to idiosyncratic factors, which are not the focus of my theoretical model and the effects of which are diluted at the aggregate level.

impacts. For robustness purposes, I also consider alternative strategies to construct this indicator series.

My empirical findings lend support to my model predictions that overall stock market liquidity declines temporarily after an uncertainty shock and persistently after an economic condition shock. For example, during the entire sample period, an uncertainty shock triggers a rapid drop in market liquidity for up to three months and a subsequent rebound till month five. In contrast, overall stock market liquidity will exhibit a persistent decline for up to one year after a one-standard-deviation decrease in the detrended S&P 500 index level and even longer for a one-standard-deviation increase in the FFR. These results are robust to alternative estimation strategies.

The present study contributes to both the theoretical and empirical literature on market microstructure. My model contributes to the growing literature on ambiguity and market microstructure that is highlighted by Easley and O'Hara (2010b), who focus on market designs and show that exchanges can enhance market quality by reducing ambiguity. Ozsoylev and Werner (2011) analyze how adverse selection and ambiguity jointly affect liquidity. Easley and O'Hara (2010a) illustrate how ambiguity and heterogeneous beliefs can result in an equilibrium where no one is willing to trade.[5] Routeledge and Zin (2009) focus on the bid-ask spread to show how ambiguity can reduce liquidity, especially the liquidity of derivatives. My model differs from the above-mentioned models in two important ways. First, the above-mentioned models are all based on the *maxmin* preference introduced by Gilboa and Schmeidler (2009). Under the *maxmin* preference, decision makers have an infinite degree of ambiguity aversion,

[5] In addition, Easley and O'Hara (2010a) adopts a different approach to model decisions under uncertainty, the Bewley (2002) approach.

according to Klibanoff, Marinacci, and Mukerji (2005). Based on the ARAP preference, my model allows explicit identifications of both the level of ambiguity and the degree of ambiguity aversion. This identification feature allows one to treat the degree of ambiguity aversion similarly to that of risk aversion, in which case market participants can display time-varying risk aversions as well as time-varying ambiguity aversions. Second, my model highlights the role of learning in resolving ambiguity and explaining the distinct behaviors of market liquidity following uncertainty shocks and economic conditions shocks. In comparison, learning is not considered in the above-mentioned models, hence ambiguity is more persistent in the above-mentioned models.

My empirical findings generate new insights into the existing literature on determinants of overall stock market liquidity. Chordia, Roll, and Subrahmanyam (2000) show that market liquidity of individual stocks moves together with each other over time. Accordingly, they introduce the notion of "market liquidity." Hameed, Kang, and Viswanathan (2010) show that market liquidity decreases after large declines in the stock market level. Chung and Chuwonganant (2014) find that market liquidity decreases in stock market volatility. However, these two studies do not address the question about which, stock market condition or uncertainty, has a longer-run impact on stock market liquidity.

II. Learning Under Ambiguity and Market Liquidity

I examine a finite horizon model with final period $T + 1$. There are two types of assets: a risk-free asset, whose price is normalized to unity, and an ambiguous asset, whose final payoff D is a random variable and is realized at $T + 1$. Specifically, the

ambiguity about final payoff is modeled as a multiple-prior belief about its stochastic nature. Asset payoffs are defined on a probability space (Ω,F, P), and under P the payoff $D \sim N(\mu_A, \sigma_0^2)$. In this case, agents can determine with high precision the variance of the asset value, σ_0^2, but not the mean μ_A. They believe μ_A can be any real value; each possible value $\mu_{A,i}$ corresponds to a model $m(\mu_{A,i}) = N(\mu_{A,i}, \sigma_0^2)$, and collectively $M = \{m(\mu_{A,i}) | i \in I\}$ is the set of plausible models. Following the setting of the smooth ambiguity model of Klibanoff, Marinacci, and Mukerji (2005), I assume that agents assign a (subjective) prior probability measure Q for alternative models to be true. Since alternative models only differ in their mean parameters$\mu_{A,i}$, I can simply describe Q as$\mu_A \sim N(\mu_0, \sigma_A^2)$. Note that if there is a single prior or Q is a Dirac delta measure, my model degenerates to the standard learning setting. In my current setting, the level of ambiguity is measured byσ_A^2, the dispersion of agents' prior beliefs.

To highlight the impact of ambiguity on market liquidity, I study a simple trading model with two kinds of agents: "market makers" and "investors". Market makers are competitive; there exists a representative market maker whose preference over ambiguous final wealth W follows the ambiguity robust Arrow-Pratt (ARAP) preference introduced by Maccheroni, Marinacci, and Ruffino (2013):

$$V(W) = E_{m_B}[W] - \frac{\rho}{2} Var_{m_B}[W] - \frac{\theta}{2} Var_Q\left[E_P[W]\right]. \tag{1.1}$$

Two observations are worth discussing in equation (1.1). First, ρ and θ refer to the risk aversion parameter and ambiguity aversion parameter, respectively. Second, the probability measure m_B is a special element in the family of multiple priors M. According to Maccheroni, Marinacci, and Ruffino (2013), the probability measure $m_B$is

called the reduction of Q, the probability measure for priors, on the probability space for the asset payoff Ω. Specifically, m_B is defined

(1.2) $$m_B(A|\mu_0, \sigma_A^2) = \int_{\{\mu_{A,i}|i\in I\}} m\left(A\middle|\mu_{A,i}\right) dQ\left(m\left(\mu_{A,i}\middle|\mu_0, \sigma_A^2\right)\right), \forall A \in F.$$

With the above setting, the following lemma shows the representation of m_B.

Lemma 1.1. Let $M = \{m(\mu_{A,i}, \sigma_0^2)|i \in I\}$ be the set of alternative models, Q be the probability measure over$M$, and Q is represented by a normal distribution $N(\mu_0, \sigma_A^2)$. The measure m_B in equation (1.1) is represented with a normal distribution $N(\mu_0, \sigma_0^2 + \sigma_A^2)$.

The above lemma can be interpreted as collapsing model uncertainty onto the uncertainty about final payoff. If $\theta = 0$, the agent is ambiguity neutral and the ARAP preference degenerates to the classic Arrow-Pratt preference. The Ellsberg (1961) paradox suggests that agents display ambiguity aversion, hence θ should be positive in general.

The ARAP preference states that the certainty equivalent is the expected value under the collapsed model $E_{m_B}[W]$ minus a risk discount $\frac{\rho}{2}Var_{m_B}[W]$ and an ambiguity discount $\frac{\theta}{2}Var_Q\left[E_P[W]\right]$, where the ambiguity discount increases in the level of ambiguity about the mean. According to Maccheroni, Marinacci, and Ruffino (2013), the ARAP preference approximates the smooth ambiguity model of Klibanoff, Marinacci, and Mukerji (2005) in the same way as the standard Arrow-Pratt preference approximates the expected utility model. The main merits of the ARAP preference are its analytical tractability and its inheritance of the flexibility in explicitly modeling the level of ambiguity aversion. In general, the smooth ambiguity model nests the *maxmin* preference

model of Gilboa and Schmeidler (1989), shown by Klibanoff, Marinacci, and Mukerji (2005). The *maxmin* preference model is obtained in the smooth ambiguity model when the level of ambiguity aversion approaches infinity.

I further assume that at each date t, the market maker sets a market clearing price p_t after observing the exogenous aggregate trading order q_t. Hence, the market maker's final wealth is

(1.3) $$W_T = \sum_{t=0}^{T}(p_{t+1} - p_t)x_t = \sum_{t=0}^{T} -(p_{t+1} - p_t)q_t, with\ p_{T+1} = D,$$

where x_t is the demand of the market maker, and the second equality is based on the market clearing condition that$x_t + q_t = 0$. Note that this assumption of inelastic investor demand is made for notational simplicity only, and does not qualitatively affect the results I derive below. The key to my liquidity model is the idea that one can "invert prices from quantities". That is one can determine the equilibrium prices consistent with the representative market maker's optimization behavior. My focus is to determine how price impact depends on the level of ambiguity σ_A^2 as well as on the degrees of ambiguity aversion θ and risk aversion ρ. As a result, all that matters is that investors have demand curves that result in market makers choosing to hold, at the market prices, a position of $x_t = -q_t$ that we observe in the data. In another word, I abstract away the motivation behind each trade. Investors may trade for pure liquidity purposes or for superior information. Nevertheless, at the aggregate level, these idiosyncratic reasons play lessor roles, leaving economic condition and economic uncertainty as the dominating factors.

To better illustrate how ambiguity affects price impact, I consider a static model first. In this case, $T = 0$, and the market maker chooses the optimal demand x by

maximizing the ARAP preference of the ambiguous wealth at time 1. The equilibrium price is determined from the market clearing condition.

Proposition 1.1. The market maker chooses the trading quantity x to maximize his ARAP preference of terminal wealth:

(1.4) $$V^*(W) = \max_x E_{m_B}[(D-p)x] - \frac{\rho}{2} Var_{m_B}[(D-p)x] - \frac{\theta}{2} Var_Q\big[E_P[(D-p)x]\big].$$

His trading quantity x in accordance is

(1.5) $$x^* = \frac{\mu_0 - p}{\rho(\sigma_0^2 + \sigma_A^2) + \theta\sigma_A^2}.$$

The equilibrium price p is

(1.6) $$p = \mu_0 + \lambda q, where\ \lambda = \rho(\sigma_0^2 + \sigma_A^2) + \theta\sigma_A^2.$$

The parameter λ, which measures price impact or inverse market depth, is an important aspect of market illiquidity. λ shows the change in transaction price p in response to a trade quantity of q and a liquid market is defined with a low price impact. Based on equation (1.6), λ will increase if the level of risk or ambiguity increases, or if the degree of risk aversion or ambiguity aversion increases. These comparative statics reflect the fact that price impact depends on the market maker's required compensation for trading the uncertain asset. Interestingly, even if the market maker is ambiguity neutral $\theta = 0$, as long as there is ambiguity $\sigma_A^2 > 0$, a one unit trade will move the transaction price by $\rho\sigma_A^2$ more than the pure risk case. Generally, if the market maker is ambiguity averse, price impact will be further increased by the required compensation for ambiguity $\theta\sigma_A^2$. As a result, we can see in this simple case that an increase in either the level of ambiguity σ_A^2 or the degree of ambiguity aversion θ will increase price impact or decrease market liquidity.

This static model generates insights regarding how price impact depends on risk, ambiguity, risk aversion, and ambiguity aversion, however, it is silent on whether uncertainty and uncertainty aversion affect price impact in the same way or in different ways. This static model only implies correlations between price impact and its four determinants, and we are still not able to differential between the impacts of ambiguity and that of uncertainty aversion on price impact. Hence, we must examine a dynamic model.

Nevertheless the capability to separately identify the impacts of risk, risk attitude, ambiguity, and ambiguity attitude on price impact is the novel feature of my model. Although my result is not the first one that analyzes the impact of ambiguity on price impact, previous results in general are silent on the differential impacts of ambiguity and ambiguity attitude. For example, Easley and O'Hara (2010b) and Ozsoylev and Werner (2011) both examine the impact of ambiguity on price impact. However, ambiguity attitude is not explicitly featured in their results because of their choices of the *maxmin* preference. As discussed before that the *maxmin* preference is a special case of the smooth ambiguity model being utilized here, compared to my model, their models can be interpreted as a description of price impact in a scenario when investors are extremely ambiguity averse. Since ambiguity attitude, like risk attitude, features a specific dimension of agents' general attitudes towards uncertainty, it is reasonable to believe that individuals can display different attitudes toward ambiguity. Meanwhile, it is also reasonable to believe that individuals' ambiguity attitudes can change through time. In this sense, an explicit accounting for the ambiguity attitude can lead to richer predictions

and generate better understandings of the reality. More importantly, this explicit modeling of ambiguity aversion is valuable for calibration exercises.

Because differentiating the impacts of uncertainty shocks and that of economic condition shocks is not viable in a static model, I develop a dynamic model. The defining feature of my dynamic model is the market maker's learning under ambiguity. Intuitively, ambiguity can arise due to a lack of relevant information to pin down the exact model. This is possible after an uncertainty shock. However, as agents gather more information through time, ambiguity can decrease gradually as a result of updating, until some new shock introduces a high level of ambiguity again.

I assume agents receive a noisy signal $s_t = D + \epsilon_t$ at each trading date from $t = 1$ to $t = T$. Here ϵ_t is the noise of the signal and assumed to be i.i.d. $N(0, \sigma_s^2)$ and independent of D. These noisy signals can be interpreted as relevant public information about the ambiguous fundamental value of the asset. I do not consider the case that the variance of the signal noise σ_s^2 is ambiguous such as Epstein and Schneider (2008). Allowing σ_s^2 to be ambiguous only adds another dimension of ambiguity without qualitatively change any result under my current choice of the smooth ambiguity model.

As a benchmark, let's consider the standard Bayesian setting first. When the prior probability measure for alternative models Q_0 is a Dirac delta measure δ_{μ_0}, there will be a single prior model for D: $m(\mu_0) = N(\mu_0, \sigma_0^2)$. As a result, at any time t, the conditional distribution for D on a series of signals up to time t $\{s_i\}_{i=1}^t$ has the following standard form

$$D|\{s_i\}_{i=1}^t \sim N\left(\frac{\tau_s \bar{s_i} t + \tau_0 \mu_0}{\tau_0 + \tau_s t}, (\tau_0 + t\tau_s)^{-1}\right), \tag{1.7}$$

where $\overline{s_t}$ denotes the average of $\{s_i\}_{i=1}^{t}$, and $\tau_s = \sigma_s^{-2}$ and $\tau_0 = \sigma_0^{-2}$ are noted as the precisions of the signal s_i and of the asset value D.

In my interested setting, there is a multiple-prior belief for D, characterized by a set of alternative models ($M_0 = \{\mathrm{m}(\mu_{A,i})|\mathrm{i} \in \mathrm{I}\}$) and a normal prior ($Q_0$) for alternative models to be true, indexed by the ambiguous mean parameter μ_A. To model learning under this ambiguous setting, I follow the general structure introduced by Epstein and Schneider (2007, 2008) and Klibanoff, Marinacci, and Mukerji (2009), among others. In this case, signals $\{s_i\}_{i=1}^{t}$ affect the market maker's belief of D, M_t, by affecting his belief of alternative models, Q_t, which can be represented by a posterior distribution for the ambiguous mean parameter $\mu_A^{(t)}$

$$\mu_A^{(t)}|\{s_i\}_{i=1}^{t} \sim N\left(\frac{\tau_s \overline{s_t} t + \tau_A \mu_0}{\tau_A + \tau_s t}, (\tau_A + t\tau_s)^{-1}\right), \tag{1.8}$$

where $\tau_A = \sigma_A^{-2}$. In this case, the market maker will have a refined belief Q_t forM_t. Specifically, the market maker will find it more likely for the mean of $\mu_A^{(t)}|\{s_i\}_{i=1}^{t}$ to approach the average of all the signals $\overline{s_t}$ as time t increases, based on equation (8). Interestingly, in the limit, Q_t will converge to a Dirac delta measure $\delta_{\overline{s_t}}$, when ambiguity is completely learned away. This feature that ambiguity can be learned away is consistent with the insight of Epstein and Schneider (2007) and particularly the one of Klibanoff, Marinacci, and Mukerji (2009), who proved that as long as the parameter space is finite, learning will resolve all the ambiguity.

This dynamic behavior of the posterior Q_t for the ambiguous parameter μ_A will directly affect the ARAP preference based on the posterior M_t for the final payoff D. Recall in equation (1.1) there are three probability measures: m_B, the collapsed model,

Q_t, the measure for alternative models, and P_t, the measure for payoff D. In this case, if Q_t is updated as equation (1.8), the collapsed model m_B will have the following form

(1.9) $$m_B^{(t)}|\{s_i\}_{i=1}^t \sim N\left(\frac{\tau_s \bar{s_i} t+\tau_A \mu_0}{\tau_A+\tau_s t}, (\tau_A + t\tau_s)^{-1} + \sigma_0^2\right).$$

This updating of the collapsed model m_B directly affects $E_m^B[W]$ and $Var_m^B[W]$, the first two terms in the ARAP preference illustrated in equation (1.1). In addition, although under P_t the conditional mean of D is still μ_A, the updating changes the distributional belief for μ_A from the prior $Q_0 = N(\mu_0, \sigma_A^2)$ to the posterior Q_t defined by equation (1.8). Consequently, $Var_Q[E_P[W]]$, the third term in the ARAP preference, will be affected.

Equation (1.9) suggests that the market maker will view the conditional mean of D as approaching the average of all the signals $\bar{s_i}$ and the conditional variance of D as approaching the unconditional variance or risk of D, when t goes to infinity. Once again, this result reveals a key feature of a dynamic model that learning can gradually resolve ambiguity. Ambiguity can have important implications in the short run, but in the long run only risk matters because investors become more confident about the correct model as they learn more information.

After I illustrate how beliefs are updated, the dynamic trading model with learning can be directly derived. The general settings are similar to those in pure risk based models, such as the ones in Kyle (1985) and Vives (1995) among others. Dynamic models of this sort are solved via backward induction. To see how learning under ambiguity can differ from learning under risk, I start with the benchmark result under risk, the standard Bayesian setting.

Proposition 1.2. Suppose there is a single prior $N(\mu_0, \sigma_0^2)$ for the terminal payoff D, and the market maker receives a signal $s_t = D + \epsilon_t$ at each date t. Then his trading quantity at time t, x_t, is

$$x_t^* = \frac{\frac{\tau_s \bar{s_t} t + \tau_0 \mu_0}{\tau_0 + \tau_s t} - p_t}{\lambda_t}. \tag{1.10}$$

The equilibrium price p_t is

$$p_t = \frac{\tau_s \bar{s_t} t + \tau_0 \mu_0}{\tau_0 + \tau_s t} + \lambda_t q_t, where\ \lambda_t = \rho(\tau_0 + t\tau_s)^{-1}. \tag{1.11}$$

This proposition is similar to the result in Vives (1995) and can be proved in a similar fashion. The parameter of interest is still λ_t, the price impact parameter at time t. We can see that λ_t decreases in time, indicating that the market becomes more and more liquid through time as a result of the market maker's updated belief of D.

This proposition highlights a testable empirical prediction: a change in the asset risk level will cause a temporary increase in price impact. This is true because λ_t is dominated by the term $t\tau_s$, which increases through time. Consequently, the shock effect of τ_0 on λ_t dissipates as time goes by.

This proposition highlights the role of learning in distinguishing dynamic trading models from static trading models. In static models, one can only predict the correlation between market liquidity and a particular variable, such as the risk level. In dynamics models, one can further predict the time series behavior of market liquidity in relation to a change in the risk level.

One limitation of the Bayesian benchmark is that both risk and price impact will approach zero as trading horizon increases. This is true if we let the trading horizon approaches infinity. The idea that there is no risk and no illiquidity in the long run may

seem to be too ideal, especially if one tries to connect the model with actual data. The following proposition will show that this concern is completely mitigated in my model of learning under ambiguity.

Proposition 1.3. Suppose there is a multiple-prior belief $M_0 = \{m(\mu_{A,i}) | i \in I\}$ for the terminal payoff D and a normal prior Q_0 for alternative models to be true, and the market maker receives a signal $s_t = D + \epsilon_t$ at each date t. Then his trading quantity at time t, x_t, is

$$x_t^* = \frac{\frac{\tau_s \bar{s_i} t + \tau_A \mu_0}{\tau_A + \tau_s t} - p_t}{\lambda_t}. \tag{1.12}$$

The equilibrium price p_t is

$$p_t = \frac{\tau_s \bar{s_i} t + \tau_A \mu_0}{\tau_A + \tau_s t} + \lambda_t q_t, where\ \lambda_t = \rho((\tau_0 + t\tau_s)^{-1} + \sigma_0^2) + \theta(\tau_A + t\tau_s)^{-1}. \tag{1.13}$$

This proposition can be proved in a similar way to Proposition 1.2. The distinction between learning under ambiguity and the standard Bayesian learning can be observed by comparing equations (1.11) and (1.13). In equation (1.13), λ_t converges to $\rho\sigma_0^2$, when the trading horizon approaches infinity and to 0 in equation (1.11). This result indicates that in a long horizon ambiguity will be completely learned away and market liquidity is only affected by risk. In another word, ambiguity only matters in a short horizon, while risk matters both in the short and long horizons.

To better illustrate the empirical implications of my model, I first conduct a set of numerical examples. Specifically, I generate a time series of λ_t based on equation (1.13) for a range of ambiguity levels, measured by σ_A^2. Here I let σ_A^2 take values of 2, 4, and 6. In all three cases, the common parameters are $\sigma_0^2 = 2$, $\sigma_s^2 = 1$, $\rho = 2$, and $\theta = 8$. Since there is no consensus for ρ and θ, I select those two in reference to Ju and Miao (2012).

Figure 1.1 shows how ambiguity and learning jointly determine the dynamics of price impact. If we consider $\sigma_A^2 = 4$ as a benchmark, we can see that an increase in ambiguity ($\sigma_A^2 = 6$) will lead to a sharp increase in price impact in the first few periods. However, the difference between the price impact curve for $\sigma_A^2 = 4$ and the one for $\sigma_A^2 = 6$ decreases as t increases. For a relatively long horizon, the difference is almost zero. The same is true for the opposite case $\sigma_A^2 = 2$. Collectively, these results demonstrate that ambiguity affects market liquidity primarily over short horizons, as in long horizons ambiguity is resolved by learning, and market liquidity converges to the benchmark under pure risk.

[Place Figure 1.1 about here]

As mentioned in the introduction that it is widely acknowledged that an uncertainty shock introduces ambiguity, my model implies that market liquidity declines temporarily after an uncertainty shock. This implication can be directly related to Figure 1 and interpreted as the effect of learning. In the next empirical section, I will formally test this implication.

In addition, my model implies that observed stock return variance can be decomposed into a predictable component, driven by risk, and an unpredictable component, driven by ambiguity, introduced by unpredictable uncertainty shocks. While the former implication requires formal empirical estimation, this implication is intuitive. For example, the ARCH/GARCH literature indicates that stock return variance can be estimated with relatively high accuracy, and abnormal high-variance regimes are short-lived. This observation is consistent with the implication discussed here in that otherwise we would have observed long-lived high variance regimes.

To analyze how changes in economic conditions affect price impact dynamics, I conduct a set of numerical examples based on various degrees of uncertainty aversion (risk aversion and ambiguity aversion). Specifically, I generate a time series of λ_t based on equation (1.13) for a range of risk aversion coefficients ρ and a range of ambiguity aversion coefficients θ. First, to see the effect of risk aversion, I let ρ take values of 2, 3, and 4, holding $\theta = 8$. Second, to see the effect of ambiguity aversion, I let θ take values of 7, 9, and 11, holding $\rho = 2$. In both cases, the fixed parameters are $\sigma_0^2 = 2$, $\sigma_s^2 = 1$, and $\sigma_A^2 = 4$. In this case, I am examining the effects of risk aversion and ambiguity aversion on market liquidity dynamics separately.

Figure 1.2 indicates that an increase in uncertainty aversion causes market liquidity to decline persistently. In the upper panel, I show the dynamics of price impact for different degrees of risk aversion; in the lower panel, I plot the same for different degrees of ambiguity aversion. First, focusing on the upper panel, we can see that an increase in risk aversion (from $\rho = 3$ to $\rho = 4$) leads to a sharp and persistent increase in price impact, and the same is true for the opposite case $\rho = 2$. This finding reiterates the idea that learning only changes the perception of uncertainty but not the attitudes toward it. Second, looking at the lower panel, we can see that an increase in ambiguity aversion (from $\theta = 9$ to $\theta = 10$) leads to a significant increase in price impact for the first few periods nevertheless this increases decays through time. This decaying effect is driven primarily by the diminishing level of ambiguity, because it is the product of ambiguity and ambiguity aversion that affects price impact.

[Place Figure1. 2 about here]

This set of results suggest that a shock to economic condition will lead to a persistent decline in market liquidity. This is true based on the notion of time-varying risk aversion. Market participants become more averse to uncertainty in economic downturns, and this elevated uncertainty aversion leads to a persistent decline in market liquidity.

III. Empirical Findings

The empirical design is motivated by Bloom (2009). One focus here is on the impacts of uncertainty shocks on stock market liquidity. Given the infrequent nature of these shocks, I select a sample period as long as possible. In this sense, to construct a monthly stock market liquidity series, I obtain daily stock price, holding period return, and trading volume data from CRSP for 1962--2013. I select common stocks (share code of 11) traded on the NYSE to make the sample as homogeneous as possible. Next, I construct a monthly stock market volatility series using the VXO index for post-1986 periods and the actual stock market volatility for the time before 1986.[6] In addition, I obtain the Federal Funds Rate (FFR) time series from the Federal Reserve Board of Governors (H15).

[Place Figure 1.3 about here]

To define volatility shocks, I follow the method introduced in Bloom (2009). Specifically, this monthly indicator series takes the value of 1 for each of the 18 shock events in Table 1, and 0 otherwise. These 18 shocks are identified when stock market volatility in a particular month exceeds the threshold of 1.65 standard deviation above the Hodrick-Prescott detrended (lambda = 129, 600) mean. In this way, this indicator series ensures that identification comes only from these large, and arguably exogenous,

[6] The VXO index is not available before 1986.

volatility shocks rather than from the smaller ongoing fluctuations. As a result, it is reasonable to believe that this indicator series captures periods when ambiguity matters most. Note that some shocks span multiple months, leaving the exact timing as a choice. I follow the scheme in Bloom (2009) to allocate each event from two alternative approaches. The primary approach is to label the month when the volatility obtains its local maximum while the alternative is to select the first month in which the volatility moves above the 2-standard-deviation threshold. Further, as suggested in Bloom (2009), shock events differ in their nature. To address this issue, I categorize events into terror, war, oil, or economic shocks. Since it is possible that economic shocks may not be completely exogenous in relation to financial market activities, in the robustness section, I construct an alternative volatility shock indicator series using only the arguably most exogenous terms of terror, war, and oil shocks.

Table 1.1 shows the 18 identified shock events, their dates under each timing scheme, and their classifications. The first 17 shocks are consistent with those in Bloom (2009) and the 18th occurs in September 2011 at the onset of the sovereign debt crisis, after the point where the Bloom sample ends. At first glance, one may wonder whether stock market levels would decrease during these volatility shocks since each appears to be adverse. Fortunately, such a concern is ameliorated. Bloom (2009) shows that the log detrended stock-market level has a correlation of -0.192 with the main 0/1 volatility shock indicator, a correlation of -0.136 with the 0/1 oil, terror, and war shock indicator, and a -0.340 correlation with the log detrended volatility index itself. This feature is not surprising. For example, events including the Cuban missile crisis raise volatility, leaving stock market levels intact, and other events including hurricane Katrina affect stock

market levels without raising volatility. Hence, focusing on these large volatility shock events would better suit the purpose of examining the effect of ambiguity.

[Place Table 1.1 about here]

I focus on market liquidity for an important reason. First, individual stock liquidity is subject to idiosyncratic factors, such as adverse selection (Kyle (1985)). Since I am interested in economy-wide factors, studying market liquidity is the best choice.

To be consistent with my model, I construct the Amihud (2002) measure at a monthly frequency using daily stock return and volume data to measure price impact. In the literature, there are two broad types of price impact measures based on the frequency of the data being used. The first class of measures is constructed using intraday data, e.g., the TAQ database. The second class is based on data with lower frequencies, like the CRSP data. The choice of a specific measure depends on the scope of the study. Normally, if one wants to study a longer time horizon, one would construct low frequency data based measures since the TAQ database starts from the 1990s while the CRPS data can be traced back to 1960s. Since my focus is on the effect of volatility shocks on market liquidity, I believe choosing a time horizon as long as possible can lead to the best inference possible, especially considering the rare nature of these volatility shocks. The choice of the Amihud measure is based on Goyenko, Holden, and Trzcinka (2009). They show that the Amihud measure best captures price impact through a comprehensive study. Since the Amihud measure is constructed for individual stocks, to create a monthly time series for the market wide price impact measure, I follow previous literature (e.g., Naes, Skjeltorp, and Odegaard (2011)) by using the equally weighted average of the Amihud measure across all stocks.

Figure 1.4 plots the time series of the market liquidity measure with 18 reference lines added indicating volatility shock events. Unlike the volatility series in Figure 1.3, market liquidity exhibits both a downward trend and a decreased variability. Note that the Amihud measure captures the price impact dimension and lower price impact is equivalent to higher liquidity. This observation is intuitive and can be interpreted as the consequence of the advancements in information technologies as well as improvements in market quality.

[Place Figure 1.4 about here]

I estimate a range of VARs using the constructed monthly time series from June 1962 to December 2013. The baseline estimation ordering is log(S&P 500 stock market index), the stock market volatility shock indicator series, Federal Funds Rates (FFR), and the market liquidity series. This ordering is based on the discussion in Bloom (2009) that shocks instantaneously affect stock markets (levels and volatility), then interest rates, and market characteristics. In the robustness session, I also examine a battery of alternative orderings. All variables are detrended with the Hodrick-Prescott(HP) filter (lambda=129,600) in the baseline estimation.

Table 1.2 provides descriptive statistics for the VAR system (Stock and Watson (2001)). Panel A presents the P-values from bivariate Granger causality tests. The null of these tests is that the cause variable does not Granger cause the effect variable. Both volatility and FFRs Granger cause market liquidity, indicating that innovations in the volatility and FFRs contain information about future market liquidity movements. Panel B shows the percentage decomposition of a variable's forecast variance into contributions of other variables based on forecast horizons. All three determinants of market liquidity

explain substantial proportions of the total forecast variance in innovations of liquidity. For example, at a 12 month forecast horizon, around 40 percent of the forecast variance of liquidity is explained by the three determinants. Collectively, estimates in Table 1.2 point out a set of strong interdependent relations between market liquidity and the three determinants.

[Place Table 1.2 about here]

Figure 1.5 plots the impulse response function of market liquidity to shocks in volatility for the whole sample period. We can see that market liquidity displays a sharp decline of around 1% within 2 months and a subsequent recovery to the pre-shock level at around month 5. The 2-standard-error bands (dashed lines) are plotted to show the significance of the result. This result is qualitatively consistent with the finding in Chordia, Sarkar, and Subrahmanyam (2005) and Chung and Chuwonganant (2014) that uncertainty adversely affects market liquidity. In addition, this result reveals a new property of the uncertainty impact on liquidity: such an impact is transitory.

[Place Figure 1.5 about here]

Next, I examine whether the ambiguity effect would differ in different time periods. A key motivation for this exercise is that the spread of information is much faster in more recent periods due to the advancements in information technology. In Figure 1.6, I plot the impulse response functions of market liquidity to shocks in volatility for periods 1962--1989 and 1990--2013 separately. We can see sharp differences in the magnitudes of impact and the speeds of rebound from the upper and lower figures. In the first period, a volatility shock generates a decline in liquidity up to 3% and the shock impact diminishes at around month 5. In contrast, a volatility shock only generates a decline in

liquidity up to 0.4%, and the speed of rebound is much faster: the shock impact disappears at around month 3.

[Place Figure 1.6 about here]

Figure 1.7 plots the impulse response function of market liquidity to changes in the stock price level and the FFR for the whole sample period. Compared with Figure 1.5, Figure 1.7 suggests fundamentally different responses of market liquidity. For example, a one-standard-deviation decrease in the stock price level is associated with a persistent drop in market liquidity up to 2% for around a year. This finding is in line with Hameed, Kang, and Viswanathan (2010), who demonstrate that stock market declines have a persistent adverse impact on market liquidity. In addition, a one-standard-deviation increase in the FFR is accompanied by a much more persistent drop in market liquidity up to 2.5%, and this effect lasts for over a year.

[Place Figure 1.7 about here]

Next, I examine whether the impact of uncertainty aversion would change in different time periods. In Figure 1.8, I plot the impulse response functions of market liquidity to shocks in the stock price level and the FFR for periods 1962--1989 and 1990--2013 separately. We can see the degree of persistence in the impact of uncertainty aversion is less in the latter period. In the former period, changes to either the stock price level or the FFR lead to a persistent response in liquidity for around a year, while in the second period such response persists for only six months. This difference suggests that market liquidity has improved significantly through time.

[Place Figure 1.8 about here]

I establish the robustness of my empirical findings through a battery of robustness analyses. The first set of robustness analyses aims to demonstrate that my findings are robust to alternative methods of constructing the shock indicators series. In addition, since the impulse-response analysis depends on VAR orderings, my second set of robustness analyses show that my results do not depend on a particular way of ordering.

In Figure 1.9, I show that my VAR result is robust to alternative measures of volatility shocks. In the first figure of the upper panel, I use the actual volatility series instead of the indicator series. In the next figure of the upper panel, I scale each 0/1 indicator by the actual detrended volatility to reflect the idea that shocks' effects may depend on shock magnitudes. In the first figure of the lower panel, I identify shocks using the first month instead of the month with the local maximum volatility. In the second figure of the lower panel, I exclude economic shocks to mitigate the concern that economic shocks may be interrelated to financial market activities. Findings from the three graphs strengthen the conclusion that an uncertainty shock triggers a temporary decline in market liquidity.

[Place Figure 1.9 about here]

In Figure 1.10, the impulse response functions of market liquidity to volatility shocks are presented for three alternative orderings of the orthogonal VAR. In the first figure of the upper panel, the VAR ordering is (stock market volatility shocks indicator, log(stock market return), Federal Funds Rates, market liquidity). In the second figure of the upper panel, the ordering is (stock market volatility shocks indicator, market liquidity). In the picture of the lower panel, the ordering is (market liquidity, FFR, log(stock market return), stock market volatility shocks indicator). It is evident that the

response of market liquidity to volatility shocks is irrelevant to actual ordering of the VAR. Collectively, these findings lend further support to my conclusion.

[Place Figure 1.10 about here]

IV. Conclusions

Major economic and political shocks, such as the Cuban missile crisis, the 9/11 terrorist attacks, and the 2008 financial crisis trigger spikes in uncertainty. If investors are ambiguity averse, these uncertainty shocks will cause drastic market reactions, driving the dynamics of overall stock market liquidity. The present paper develops a dynamic trading model with learning under ambiguity to analyze the distinction in the impacts of these uncertainty shocks from that of economic condition shocks on the price impact aspect of market liquidity. First, I consider the idea that an uncertainty shock introduces ambiguity and learning can resolve ambiguity, as suggested by Epstein and Schneider (2007, 2008), and Klibanoff, Marinacci, and Mukerji (2009). Second, to analyze the impact of economic condition shocks, I incorporate the notion of time-varying risk aversion, introduced by Campbell and Cochrane (1999) and Barberis, Huang, and Santos (2001). My model implies that an increase in ambiguity leads to a temporary spike in price impact, while an increase in either risk aversion or ambiguity aversion results in a persistent increase in price impact. This distinction in the behaviors of price impact highlights the role of learning in resolving ambiguity that is introduced by uncertainty shocks.

My empirical analyses lend support to my model implications. Market illiquidity, measured by price impact and constructed monthly using the Amihud (2002) measure, declines on average up to 1.5% and rebounds within the next five months after an

uncertainty shock. In contrast, market liquidity declines persistently for up to a year after a shock to economic conditions.

Chapter 2:

The Early Exercise Premium in American Put Prices[7]

I. Introduction

Most exchanged-traded options have the American exercise style, and this early exercise feature contributes significantly to the value of the options.[8] Since the seminal study of Bakshi et al. (1997), researchers have devoted substantial efforts on empirically testing alternative option pricing models' performances to explain European option prices.[9] At this point, however, we still know little about how well these models can match the magnitude of and the variation in empirically observed EEP. In this paper, we focus on put EEP and exploit a unique fact that the S&P 100 index trades both American options (OEX) and European options (XEO), allowing direct measurements of EEP, to address two important questions. First, how well can alternative option pricing models explain the cross section of put EEP? Second, do alternative models systematically overprice or underprice put EEP? If so, how can we reconcile that?

In the principle of parsimony, we compare the Black and Scholes (1973) model and its four extensions for pricing American puts and EEP. The extensions include the Merton (1976) jump diffusion model, the variance gamma model of Madan et al. (1998), the finite moment log stable model of Carr and Wu (2003), and the stochastic volatility model of Heston (1993). Accordingly, they increase the dimension of the Black and

[7] This chapter is based on collaborative work with Louis Ederington and Pradeep Yadav. The computing for this project was done at the Oklahoma University Supercomputing Center for Education & Research (OSCER).
[8] Examples include options on equity, currency, index futures, etc.
[9] Earlier studies in this direction are surveyed by Bates (2003).

Scholes model by three, two, one and four, and they can serve as the foundation for more sophisticated models.[10] For the sake of brevity, henceforth we refer to these models as the BS, MJ, VG, FMLS, and SV models, respectively.

Studying these models' performances in pricing put EEP can generate new insights about the implied early exercise boundaries of these models. Classic theories including Carr et al. (1992) based on arbitrage-free arguments show analytically that put EEP is the present value of the interests earned on the strike price net of the foregone dividends, when the underlying price is in the exercise region. This exercise region is defined by a model specific early exercise boundary, which is an optimal threshold for the underlying price to fall from above such that an early exercise is triggered. This early exercise boundary is the key to price EEP and to distinguish alternative models. It is not analytically tractable, and neither is its connection with the shape of the implied volatility surface. Hence, in contrast to previous empirical studies that focus on European option prices and assess alternative option pricing models based on their abilities to match the implied volatility surface, our study offers new insights about the relevance of alternative models based on the early exercise boundary.

We empirically measure market EEP by pairing quotes of an OEX option and an XEO option on the S&P 100 index with identical strike and maturity, if the two quotes

[10] For example, one can combine the Merton jump with the SV to get the SVJ model of Bates (1996) or further randomize the jump activity rate to arrive at the model studied in Bates (2000), and possibly add a jump in the volatility as well to get the SVJJ model studied in Duffie et al. (2000) and Pan (2002). Furthermore, one can introduce a particular dependence structure in the occurrences of jumps in the return equation and in the variance equation, as in the SVCJ model studied in Eraker et al. (2003) and Eraker (2004). In addition, one can use the variance gamma process or the α-stable process to model jumps instead of the compound Poisson process, and such generalizations are featured in Huang and Wu (2004) and Li et al. (2008). Alternatively, one can divert from the above-mentioned affine generalizations and resort to the time changed Levy process models, e.g., subordinating the variance gamma model by a Feller "clock" as in Carr et al. (2003) and Carr and Wu (2004).

are posted within 60 seconds. Our sample consists of high frequency quotes data on S&P 100 index options for the calendar period 2005.05—2007.12. We choose this period because short-term interest rates have been close to zero for a long time after the financial crisis in 2008, while the dividend yield on the S&P 100 index has been stable (around 2 percent). This makes the early exercise of OEX puts unattractive in the long period after 2008, because the interests earned can hardly exceed the forgone dividends. In contrast, the short-term interest rates are around 4 percent in our sample period, much larger than dividend yields. In this case, the early exercise of OEX puts can be optimal in general.

When implementing each model to determine its predicted EEP, we first back out the set of risk neutral parameters for each day using out-of-the-money (OTM) European option prices by minimizing the root mean squared pricing errors (RMSE). Next, we use their weekly averages as inputs to the Least Squares Monte Carlo (LSM) American option-pricing algorithm developed in Longstaff and Schwartz (2001). We choose the LSM algorithm based on its wide applicability for all types of models and for the convenience it offers in adjusting for discrete cash dividends (Harvey and Whaley (1992a, 1992b)) relative to alternative methods, including the spectral methods developed in Fang and Oosterlee (2009) and Ruijter and Oosterlee (2012) and the asymptotic expansion method in Medvedev and Scaillet (2010).11 Our approach of estimating the

[11] The LSM algorithm falls within the general framework of recursively solving an optimal stopping problem by comparing the continuation value and the immediate stopping proceeds at each step going backwards. The LSM algorithm relies on simulated sample paths to approximate the continuation value through a least squares exercise, while spectral methods rely on an orthogonal basis to approximate the transition density of the underlying, and hence the continuation value at each time. Aside from such a recursive method of solving the optimal stopping problem, finite differencing the PDE or partial integro-differential equation (PIDE) implied by the variational inequality subject to a set of initial and boundary conditions is also popular in the literature. We acknowledge the better convergence of the solution generated by finite difference methods, and we opt not to use them simply due to the fact that one may need to make ad hoc adjustments for each model to meet the stability conditions.

parameters each day follows the approach used widely in the literature for the purpose of pricing the cross section of European options – for example, Bakshi et al (1997), Carr et al. (2002), Carr and Wu (2003), and Carr et al. (2007). We estimate the parameters daily even though model parameters are theoretically constant overtime. The reason is highlighted by Carr et al. (2007): the parameters estimated each day from models that adequately fit the daily cross section of option prices vary significantly from day to day, and restricting the parameters to be constant across different days delivers sub-optimal pricing performance.[12]

When assessing the performances of alternative models in predicting EEP, we rank them based on their prediction RMSE, in accordance with Christoffersen and Jacobs (2004). They show that the choice of the loss function in the prediction step must be consistent with that in the estimation step, because otherwise the statistical distributions for pricing errors in the two steps are inconsistent, causing the inferences invalid. Since we minimize the RMSE when estimating model parameters, to make valid inferences, we must compare alternative models' predicting powers using the same metric.

It is also important to account for the value of a "wild card option" embedded in OEX contracts, when one is constructing the RMSE.[13] We determine the value of this wild card option based on the Fleming and Whaley (1994) approach, which requires the knowledge of the volatility parameter of the underlying during the wild card window. This parameter is not directly observable, because the stock market is closed during the

[12] Since our goal is to examine the pricing performance, it is necessary for us to estimate these parameters daily. This approach could not have been used if our purpose had been to form dynamic arbitrage portfolios or to estimate factor risk premia, like the variance or jump risk premia.

[13] Holders of the American style OEX contract can continue to exercise the option after 3:00 pm (when the stock market closes) till 3:20 pm, based on the closing cash index level at 3:00 pm. This flexibility has been termed as the wild card option, e.g., by Fleming and Whaley (1994). Also note that the wild card option does not exist for equity options.

wild card window. To circumvent this problem, we back out the wild card period volatility parameter from call EEP, because OEX calls should not be exercised early in our sample period for the high interest rates, leaving the wild card option the only sensible reason for early exercises.

We find that jump models are superior at matching the cross-sectional variation and magnitudes of market EEP. Since EEP depend directly on the model implied early exercise boundaries, this finding suggests that investors consider jumps as the first-order concern when valuing American options' flexibility to exercise early for earning interests purposes, as in classic theories (Carr et al. (1992)). This finding also generates new insights about jump models. Conventional wisdom suggests that jump models are essential to explain why market overprices OTM European puts relative to the BS model, while our results in contrast demonstrate that jump models are essential to explain why market overprices ITM put EEP relative to the BS model. In another word, we lend support to jump models from a completely different angle.

We in addition document that all models underprice market EEP and this underpricing is more pronounced for short term ITM puts. We provide an explanation based on the unique feature of American options, the flexibility to exercise early, which generates a transaction cost saving option. When the half spread of an American option exceeds its time value, an investor can exercise the option early as a more profitable alternative to selling it. Following this idea, we augment the Longstaff-Schwartz algorithm to quantitatively determine the value of this transaction cost saving option for each option series in our sample. Based on its explanatory power and the empirical regularity that the highest proportion of short term ITM OEX puts exhibit larger half

spreads than time values, we find that this transaction cost saving option is the reason why market EEP are larger than the combination of model predicted EEP and wild card premia.

This study contributes to several strands of literature. First, our study contributes to the literature on model specification by uncovering new properties of jump models. Prior studies focus on alternative models' performances in pricing the cross section of European options. They find that jump models are essential for explaining the high implicit negative skewness, observed in the implied volatility smirk for short-term options, while the importance of modeling jumps seems to diminish for longer horizons. Notably, studies in this category include Bakshi et al. (1997), Pan (2002), Eraker (2004), Huang and Wu (2004), and, Carr and Wu (2007). Here we find that jump models are essential for matching the magnitudes of EEP regardless of the maturity dimension. This finding suggests that modeling jumps is the key to match the implied early exercise boundaries from put EEP.

Second, it contributes to the empirical literature on the EEP. Jorion and Stoughton (1989) perform an early study that focuses on the Black-Scholes model only and use daily currency option data that is potentially suffering from significant synchronization errors, which they duly acknowledge. Harvey and Whaley (1992a, 1992b), Fleming and Whaley (1994), and Broadie et al. (2000) focus on the relative contributions of dividends and stochastic volatility in the call EEP. However, the market value of EEP in these studies is not directly observable. None of these studies focus on American puts nor do they undertake a comprehensive empirical assessment of alternative models, as we do.

II. Option Pricing Theories

In this section, we outline the theoretical foundation for our analysis: the risk neutral pricing framework. First, we describe the European option-pricing problem. Second, we discuss the American option-pricing problem and the theoretical properties of the EEP. Third, we illustrate model specifications for the underlying dynamics.

Given a complete filtered probability space $(\Omega, \mathcal{F}, \mathbb{Q}, \{\mathcal{F}_t\})$, the price of a European call on a non-dividend paying asset under the risk neutral probability measure $\mathbb{Q}$ can be described as

(2.1) $$Call(t,T,K|S(t)) = e^{-r(T-t)}E^{\mathbb{Q}}[\max\{S(T)-K\}|\mathcal{F}_t],$$

where $S(t)$ stands for the stock price at time t, K is the strike price, r is the risk-free rate, and T refers to time to maturity. If the underlying asset pays dividends, we need to adjust the $S(t)$ with the sum of the present values of the cash dividends paid during the life of the option. As shown, e.g., by Bakshi et al. (1997) and Carr and Madan (1999), equation (2.1) can be rewritten as

(2.2) $$Call(t,T,K|S(t)) = S_t\Pi_1 - e^{-r(T-t)}K\Pi_2,$$

where $\Pi_1 = \overset{\mathbb{P}}{\Pr}\{S_T > K\}$, the physical probability of the call finishing in the money, and $\Pi_2 = \overset{\mathbb{Q}}{\Pr}\{S_T > K\}$, the risk-neutral probability of the call finishing in the money. Their expressions are given by the following:

(2.3) $$\Pi_1 = \frac{1}{2} + \frac{1}{\pi}\int_0^{\infty} Re\left[\frac{e^{-iulnK}\phi_T(u-i)}{iu\phi_T(-i)}\right] du,$$

and

(2.4) $$\Pi_2 = \frac{1}{2} + \frac{1}{\pi}\int_0^{\infty} Re\left[\frac{e^{-iulnK}\phi_T(u)}{iu}\right] du,$$

where $\phi_T(u) = E^{\mathbb{Q}}[e^{iulnS_T}]$, is the characteristic function of lnS_T.

Although we only explicitly describe the European call-pricing problem above, European put prices can be easily determined through put-call-parity:

(2.5) $$Put_t = Call_t - S_t + e^{-r(T-t)}K.$$

Same as above, if the underlying asset pays dividends, one needs to adjust for the present value of cash dividends paid during the life of the option.

The American option-pricing problem is more complicated than its European counterpart due to its path dependent nature. We begin with the simplest case being the pricing of American options on non-dividend paying assets, which is formally analyzed in Carr et al. (1992). We then move to a more general case where the underlying asset pays dividends and the underlying asset dynamics is driven by a discontinuous jump process, which is formally analyzed in Gukhal (2001).

In the case of a non-dividend paying asset, an American call option with a finite maturity should never be exercised prior to maturity; while an American put option can be optimally exercised early. Hence, the American put commands a premium over its European counterpart.

Mathematically, the problem of pricing a finite maturity American put option can be expressed as the following optimal stopping problem

(2.6) $$Put^{Am}(t,T,K|S(t)) = \sup_{t<\tau\leq T} e^{-r(\tau-t)}E[\max\{K - S(\tau)\}|\mathcal{F}_t]$$

where τ is a stopping time with respect to the filtration $\{\mathcal{F}_t\}$. The solution to (2.6) is attained at $\tau^* = \inf\{t|S_t \geq S^*(t)\}$, and $S^*(t)$ is called the early exercise boundary. The literal meaning of (6) is that the price of an American put is the discounted cash flow at a random time τ^* when the stock price first reaches the optimal early exercise boundary $S^*(t)$. To elaborate, the holder of an American put on a non-dividend paying underlying

should exercise early only if the underlying's price at time t ($S(t)$) falls from above to a time dependent critical value $S^*(t)$. When $S(t)$ is above $S^*(t)$, the holder should continue to hold the American option, hence the region $\{S(t)|S(t) > S^*(t)\}$ is called the continuation region. In the continuation region, the value of an American option equals that of a European option. On the contrary, when $S(t)$ hits the $S^*(t)$ for the first time, the holder should exercise the option and as a result receives payoff $K - S^*(t)$, where K is the strike price. In this way, the region $\{S(t)|S(t) \leq S^*(t)\}$ is called the exercise region, and specifically, the curve $\{S(t)|S(t) = S^*(t)\}$ is called the early exercise boundary.

Even in the simplest case where $S(t)$ is modeled as a geometric Brownian motion, equation (2.6) has no analytical solution. Carr et al. (1992) provide an alternative characterization for the solution of (2.6) as

(2.7) $$Put^{Am}(t,T,K|S(t)) = Put^{Eu}(t,T,K|S(t)) + EEP(t,T,K|S(t)).$$

The above expression says that the price of an American put equals the sum of the price of its European counterpart and an EEP. In addition, they illustrate an approach to construct a portfolio that replicates the cash flow from holding an American put. To do so, an investor should go long in the put in the continuation region, and exercise the put when the underlying hits the early exercise boundary and enters the exercise region receiving the payoff of $K - S^*(t)$. When the underlying's price is in the exercise region, the investor should invest K in risk-free bonds and short one unit of the underlying. As long as the underlying's price stays in the exercise region, the investor earns interest. If the underlying's price trajectory is continuous as modeled by a diffusion process, the underlying's price transition from the exercise region to the continuation region is

smooth. The investor can buy the option at the early exercise boundary at price $K - S^*(t)$, and spend $S^*(t)$ to close the short position in the underlying. With this strategy, the investor's payoff at maturity is the same as that of a European put, but the investor collects interest along the way, that would otherwise be impossible while holding a European put. Thus, the EEP is the interest earned when the underlying's price is in the exercise region.

Alternatively, if the underlying asset pays dividends, then American calls can be exercised early and command premia over their European counterparts. In this way, the call EEP is the dividends received in the exercise region net of the interest expense paid on the strike, while the put EEP in this case is the interest earned on the strike net of the foregone dividend in the exercise region. These statements follow a similar argument as that in the above dividend free case.

Gukhal (2001) generalizes equation (2.7) by studying the case when the trajectory of the underlying dynamics is discontinuous and modeled as a jump diffusion process. He shows that the put EEP in this case equals the present value of interest earned on the strike price in the exercise region minus the rebalancing cost due to jumps from the exercise region into the continuation region. The main distinction is that when the underlying price jumps upward from the exercise region into the continuation region in a discontinuous way, the investor has to pay a price higher than the exercise payoff to buy back the American put. Recall our previous discussion of the continuous case that, to replicate the cash flow of an American put, an investor needs to buy back the option when the underlying price hits the early exercise boundary from below at the price of the option payoff. But if the underlying price jumps upward directly into the exercise region,

bypassing the early exercise boundary, the investor will not be able to buy the option at the price of the option payoff since the American option price is greater than its payoff in the continuation region. On the other hand, when the underlying price jumps downward from the continuation region into the exercise region, the investor does not need to pay the additional rebalancing cost, since the price of an American put equals the exercise payoff in the exercise region.

To summarize, the put EEP is the present value of interest earned on the strike net of foregone dividends in the exercise region when the underlying trajectory is continuous. Additional rebalancing costs due to jumps from the exercise region into the continuation region need to be subtracted when the underlying trajectory is discontinuous. Similarly, the call EEP is the present value of dividends received net of interest paid on the strike in the exercise region, and rebalancing costs due to jumps need to be adjusted. Nevertheless, this statement does not necessarily imply that the put EEP or call EEP under a continuous model of the underlying should be greater than that under a jump discontinuous model. The optimal exercise decision as well as the replicating strategy depends on the first passage probability of a specific model and the associated early exercise boundary. There is no analytically tractable way to compare the first passage probability and the early exercise boundary of different models.

Motivated by the above theoretical insights, we begin with the Black and Scholes (1973) model as a benchmark where the underlying is modeled as a geometric Brownian motion. The risk neutral dynamics of the BS model reads

(2.8) $$dS_t = rS_t dt + \sigma S_t dW_t,$$

where W_t is a standard Brownian motion, r is the risk free rate, σ is the volatility parameter.

However, extensive empirical evidence has emerged that undermines the assumption of lognormal returns in the Black-Scholes model. The first contradicting feature is the "volatility smirk". For a given option maturity, the Black-Scholes implied volatilities for out-of-the-money (OTM) options are much higher than those of call options that are equally OTM. Second, as documented by Carr and Wu (2003), the implied volatility smirk does not flatten out but steepens slightly as maturity increases, for up to two years.

Widely acknowledged, the implied volatility smirk indicates that the risk-neutral distribution of the underlying index return is negatively skewed and fat-tailed. Specifically, extreme downward movements in the underlying return occur with larger probabilities and magnitudes than otherwise predicted by a normal distribution. Previous studies either introduce jumps in the underlying return or allow the instantaneous variance to be stochastic to generate the implied volatility smirk.

On the other hand, the maturity pattern, that the smirk persists as maturity increases, is inconsistent with the central limit theorem in which we expect the smirk to flatten out through the maturity horizon. To capture the persistent smirk pattern, researchers either resort to models with both jumps in return and stochastic volatility or jump models with infinite variation. In the former case, the jump component can generate short-term high negative skewness in the return distribution, and the stochastic volatility component slows down the convergence of the return distribution to normal for increased

maturities. In the latter case, the required condition for the central limit theorem to hold is violated, allowing the smirk arising from jumps to persist as maturity increases.

Generalizations of the Black-Scholes models all aim to incorporate at least one aspect of the abovementioned empirical regularities. Given the complexity of the American option pricing problem and the gap in the literature, we opt to proceed from a set of the most parsimonious generalizations. These include the jump-diffusion model first introduced by Merton (1976), the variance gamma model of Madan and Seneta (1990), Madan et al. (1998), the finite moment log stable model of Carr and Wu (2003), and the stochastic volatility model of Heston (1993).

The Merton jump diffusion model (MJ) is a four-parameter model with finite jump intensity and is constructed by introducing a compound Poisson process in the BS model to generate infrequent jumps. The risk neutral dynamics of the underlying level follows

(2.9) $$dS_t = \left(r - \lambda\mu_J\right)S_t dt + \sigma S_t dW_t + J_t S_t dN_t,$$

where N_t is a Poisson counting process with intensity λ, and J_t is the random jump size, distributed as $\ln(1+J_t) \sim N(\ln\left(1+\mu_J\right) - \frac{\sigma_J^2}{2}, \sigma_J^2)$ ($N(,)$ refers to the normal distribution). Note that if λ is restricted to zero (meaning no jumps will occur), the MJ model will degenerate to the BS model. Based on the model construction, we can see that jumps are discrete events, and in a finite time interval ($[0, T]$), the model is expected to jump a finite number of times, λT. Furthermore, the mean jump size μ_J controls the skewness of the return distribution, a negative μ_J indicating a negatively skewed return distribution.

The variance gamma model (VG) is a three-parameter pure jump model with infinite jump intensity and finite variation in its sample path. It is constructed by substituting the Brownian motion component in the BS model with a variance gamma process, a process constructed by subordinating the Brownian motion with a gamma process. Following the original notation in Madan et al. (1998), we write the VG model under the risk neutral measure as

(2.10) $$S(t) = S(0)\exp(rt + X(t;\sigma,\nu,\theta) + \omega t),$$

where $X(t;\sigma,\nu,\theta)$ denotes the VG process, and $\omega = \frac{1}{\nu}\ln\left(1 - \theta\nu - \frac{\sigma^2\nu}{2}\right)$ is the martingale correction term. The VG process can be obtained by evaluating a Brownian motion with drift at a random time given by a gamma process. Specifically, if

(2.11) $$b(t;\theta,\sigma) = \theta t + \sigma W(t),$$

then the VG process $X(t;\sigma,\nu,\theta)$ is defined as

(2.12) $$X(t;\sigma,\nu,\theta) = b(\gamma(t;1,\nu);\theta,\sigma),$$

where $\gamma(t;1,\nu)$ refers to the gamma process with unit mean rate, and variance rate ν. The VG model is a three-parameter model: (1) σ the volatility of the Brownian motion, (2) ν the variance rate of the gamma time change and (3) θ the drift in the Brownian motion with drift. The VG process provides additional control over the shape of the distribution beyond the volatility or standard deviation. One not only controls the skewness via parameter θ but also the kurtosis with parameter ν. For example a negative θ indicates a negatively skewed distribution, and a positive ν leads to a fat-tailed distribution. Different from the MJ model, the VG model exhibits infinitely many small jumps for the Levy measure of the VG process, and is singular at zero jump size.

However, the sample paths of the VG process display finite variation (Huang and Wu (2004)).

The finite moment log stable model (FMLS) is a two-parameter pure jump model with infinite jump intensity and infinite return variation but finite moments for the underlying level. Following the original notation in Carr and Wu (2003), we write the risk-neutral dynamics of the FMLS model as

(2.13) $$\frac{dS_t}{S_t} = rdt + \sigma dL_t^{\alpha,-1},\ t \in [o,\mathcal{T}],\ \alpha \in (1,2),\ \sigma > 0,$$

where the increment $dL^{\alpha,-1}$ has an α-stable distribution with zero drift, dispersion of $dt^{1/\alpha}$, and a skew parameter -1: $L_\alpha(0,\ dt^{\frac{1}{\alpha}},-1)$, and $\mathcal{T}$ is some arbitrary distant horizon. The FMLS model has two parameters, namely σ and α. The one of particular interest is the parameter α, which governs the thickness of the tail. A special case is for $\alpha = 2$, in which case the FMLS model degenerates to the BS model; hence in general α is required to be less than 2. Under such parameter specification, Carr and Wu (2003) show that the left tail of the increment is fat (decaying based on a power law) and the right tail is thin (decaying exponentially). In this case, both the jump intensity and the return variance become infinite. The infinite return variance renders the relevance of the central limit theorem (CLT), since for the CLT to hold a random variable should have finite variance. Finally, although the log return exhibits infinite variance, such infinite variation stems only from the left tail not the right tail, as a result, the stock price level remains finite, hence the option price.

The Heston (1993) stochastic volatility model (SV) is a five-parameter two-dimensional model with continuous sample paths that allows the instantaneous return

volatility (standard deviation) of the BS model to come from a Feller process. The SV model takes the following form

(2.14) $$dS_t = rS_t + \sqrt{V_t}S_t dW_t^{(S)},$$

where V_t is the instantaneous variance and is governed by the following Feller process

(2.15) $$dV_t = \kappa(\theta - V_t)dt + \sigma_v\sqrt{V_t}dW_t^V,$$

where κ is the mean reverting speed parameter, θ refers to the long-run variance, σ_v is the volatility of the variance process, and the two standard Brownian motions have a correlation coefficient ρ. Furthermore, for the variance process to remain positive, the Feller condition has to be satisfied: $2\kappa\theta - \sigma_v^2 > 0$. The SV model can generate a slow decaying volatility smirk at the maturity dimension. However, it suffers from its limitation to generate the extremely high level of implicit skewness normally observed in index option data. For example, the short-term implied volatility curve of the SV model is "too flat" compared to that of either jump models or to real data.

III. Data and Market EEP

We obtain intraday trades and quotes for American and European puts and calls on the S&P 100 index (symbols: OEX and XEO) across different strikes and maturities from Tick Data Corporation for period 2005.05.02—2007.12.31. This is the best period to study the EEP in OEX puts, because the extremely low interest rates in later periods diminish the incentive to exercise American puts early in general.[14] We collect the daily cash dividend series of the S&P 100 index from Bloomberg. We select daily interest rate

[14] The data is unavailable from Tick Data for period before 2005.05.

series of the Treasury bills for corresponding maturities from the Federal Reserve Board of Governors (H15) and linearly interpolate when necessary.

Several filters for the option sample are applied. We require that the bid price should be strictly positive, and the ask price should be no less than the bid price. We pair a European option and an American option with the same strike and time to maturity, if their quotes are posted within 60 seconds. In this way, we use the option quote midpoint as the fundamental value of the option, and the EEP is the difference between the quote midpoint of the American option and that of its European counterpart. We require the option maturity to be no less than 6 days and no greater than 180 days and the moneyness (K/S) to be in the range of [0.8, 1.2], since options otherwise tend to be too illiquid to contain valuable information (Huang and Wu (2004)). After we apply these filters, our sample consists of 91,083 puts and 88,424 calls.

To describe our sample, we show in Figure 2.1 the moneyness and maturity distributions and in Table 2.1 the summary statistics. We can see from Figure 2.1 that there are more OTM puts than ITM puts and more ITM calls than OTM calls. This character has been documented in the SPX market as well and is interpreted as an institutional feature (e.g., Huang and Wu (2004)). When dealing with option values, we normalize them by the S&P 100 index level and multiply them by 10^4. Two observations emerge from Table 2.1. First, market EEP are economically significant. For example, the average EEP/XEO for long term ITM puts is around 8.2 percent, which means that for an XEO put with market price of 100 dollars the OEX put with identical strike and maturity has a market price of roughly 108 dollars. Second, the absolute magnitudes of put EEP (EEP/S0) are significantly large and display substantial variations across moneyness and

maturity categories. For example, the full sample average EEP/S0 is 16 basis points with a standard deviation of 26 basis points. Looking at the subsamples, we can see that the category mean of EEP/S0 can be as large as 63 basis points.

[Insert Table 2.1 about here]

[Insert Figure 2.1 about here]

Before we evaluate alternative models' performances in pricing the EEP, it is informative to examine some empirical properties of market EEP. Table 2.2 and Figure 2.2 shows the comparative statics of market EEP. In Table 2.2, we regress market EEP on moneyness (K/S_0), moneyness-squared, maturity, interest rates, and dividend yields. In the regressions, we standardize all explanatory variables to aid interpretations. In Figure 2.2, we plot market EEP against moneyness and maturity.

The estimates reported in Table 2.2 the patterns shown in Figure 2.2 are consistent with the usual analysis of EEP, especially the analysis of the relations between EEP and interest rates/dividend yields. The estimates worth mentioning are the coefficients of interest rates and dividend yields, in addition to the coefficients of moneyness and maturity. For example, if the interest rate increases by one standard deviation, the average put EEP/S0 will increase by around 3 basis points, where the average is 16.4 basis points (Table 2.1). The same increase for dividend yield will decrease the average EEP/S0 by around 1 basis point. In addition, it is evident that market EEP display significant variation across moneyness and maturity. These findings point to the fact that put EEP in this period are informative about the theory of American option pricing hence model specification that are outlined previously.

[Insert Table 2.2 about here]

[Insert Figure 2.2 about here]

IV. Empirical Methodology

To price the EEP contained in the American option price, we first need to estimate the vector of model parameters, Θ. We adopt a well-established approach in the literature (Bakshi et al. (1997); Carr et al. (2002); Carr and Wu (2003); Carr et al. (2007)) by estimating each model on a daily basis, using the cross-section of European options on both the moneyness and the maturity dimensions. Specifically, at day t, we minimize the square root of the mean-squared pricing error (RMSE),

$$\Theta_t \equiv \arg\min_{\Theta_t} \sqrt{mse_t}. \tag{2.16}$$

The mean squared pricing error (MSE) is defined as

$$mse_t = \frac{1}{N_t} \sum_{i=1}^{n_{t,\tau}} \sum_{j=1}^{n_{t,k}} e_{ij}^2, \tag{2.17}$$

where $n_{t,\tau}$ and $n_{t,k}$ denote, respectively, the number of maturities and the number of moneyness levels per each maturity at date t, and e_{ij} represents the pricing error at maturity i and moneyness j.

The Pricing error matrix $e = \{e_{ij}\}$ is constructed as

$$e = \begin{cases} O(\Theta) - O(put), & if\ K < S, \\ O(\Theta) - O(call), & if\ K > S, \end{cases} \tag{2.18}$$

where $O(\Theta)$ denotes model predicted option prices as a function of the parameter vector Θ, and $O(put)$ and $O(call)$ are the market prices of puts and calls. Note that we are using OTM puts and calls to construct the pricing error. Such a construction is standard in the literature, e.g., Carr and Wu (2003), Huang and Wu (2004) and others. Specifically, as discussed in Carr and Wu (2003) and Huang and Wu (2004), ITM options are less

sensitive to model specifications for their positive intrinsic values, and their prices are much larger than their OTM counterparts; hence, using them to construct the pricing error tends to cause parameters to disproportionately favor the information contained in them. They also note that OTM options tend to be more liquid in general, which makes them more reliable when there is an inconsistency between their quotes and their ITM counterparts.

After we estimate the model, we move on to price American options and EEP using the least square Monte-Carlo (LSM) method, developed by Longstaff and Schwartz (2001). The finite maturity American option pricing problem can only be solved numerically even for the simplest BS model. Hence, the choice of a specific numerical method really depends on the actual problem at hand. The LSM method, like any other recursive method for dynamic programming problems, starts from discretizing the time domain, e.g., T exercise steps per year where T stands for the number of days to maturity.[15] The value of the option, $V(T, S(T))$, at the final step is the exercise payoff, $Exe(T, S(T), K)$ ($S(T) - K$ for calls and $K - S(T)$ for puts). Then moving backwards, at each step t between T and 1, the value of the option at step t, $V(t, S(t))$ is given by

(2.19) $V\big(t, S(t)\big) = \max\{e^{-r\Delta t} E^{\mathbb{Q}}[V(t+1, S(t+1)|S(t)], Exe(t, S(t), K)\}$,

where the first term in the maximum operator represents the continuation value and the second term represents the immediate exercise payoff. Thus, the value of the American option today is obtained by repeating this procedure until step 1. The LSM algorithm starts from simulating a $N \times T$ matrix of sample paths, where N is the number of paths,

[15] Obviously here we are valuing the American option approximately as a Bermuda option that allows for exercise each day during its life.

and T is the number of exercise steps. Then at each time step, t, we select those paths where the option with strike K is in the money in those paths $\{S(t)\}^{ITM} = \{S^i(t)|Exe(t,S(t),K) > 0\}$, and use all values in $\{S(t)\}^{ITM}$to approximate the time t conditional expectation of the option value at time $t+1$, the first term in the maximum operator of equation (18).[16] To implement the LSM algorithm, we generate 2×10^6 (10^6plus 10^6 antithetic) sample paths, adjust cash dividends along the paths, and use a 10-degree polynomial of $\{S^{ITM}(t)\}$ to approximate the conditional expectation.

There are two primary reasons to choose the LSM method over its alternatives. First, as documented, it is inappropriate to apply a constant dividend yield to the dividend series of the S&P 100 index, thus the method should allow adjusting discrete dividends along the exercise steps without causing numerical instabilities. Second, the LSM method is easy and convenient to implement, only requiring one to simulate the underlying dynamics. Although if one considers all available numerical methods for pricing American options, the LSM method is not among the most accurate methods due to its slow convergence ($O\left(\frac{1}{\sqrt{N}}\right)$) where N refers to the number of simulated paths). However, since this problem is shared by all Monte-Carlo simulation based methods, and we are applying the same method to different models, it should not cause any systematic bias in the pricing result.

When implementing the LSM method to price the cross-section of options in a given day, for both calls and puts, and for both sample periods, we opt to use the weekly average of the daily estimated parameters, instead of directly using the set of parameters

[16] For more information about the mechanics of the LSM algorithm, one can refer to Longstaff and Schwartz (2001).

estimated from European options in the day. An insight from Carr et al. (2007) is that, in general, option-pricing models are not able to achieve both quality cross-sectional pricing performance and time series consistency. That is, if one implements this daily estimation routine by minimizing the daily root mean squared pricing error, then there will be inevitable variations (sometimes large) in the daily parameter series. On the other hand, if one restricts the set of structural parameters to be the same over the entire sample period, then the pricing performance may not be adequate. This observation suggests a trade-off between pricing performance and time series consistency in estimating model parameters. As a result, we use the weekly average of the parameters estimated daily as a compromise to this trade-off. In addition, using the weekly average may partially immunize us from the possibility that model parameters may experience structural breaks or display regime-switching behavior, a possible scenario that should not be ignored.

For the OEX contract, there is another institutional specific determinant of its value, in addition to the flexibility to exercise early, the defining feature of American options. The CBOE grants holders of the OEX contract a "wild card option", allowing them to exercise the option until 3:20 pm (i.e., up to twenty minutes after 3:00 pm when the stock market closes) based on the 3:00 pm level of the S&P 100 index. Consequently, this institutional feature should be accounted for when evaluating alternative models' performances in pricing the EEP.

This wild card option can be another reason to exercise an OEX option early and adds value to the option. Despite the fact that the actual S&P 100 index level remains constant during this twenty minutes window, investors can still find it optimal to exercise the OEX contract based on their perception of the fundamental value of the index during

this time. For example, a holder of an OEX put may exercise the option at 3:20 pm, if she believes the fundamental value of the index has experienced a significant increase during this time window. In this case, the immediate exercise proceeds become the difference between the strike price and the 3:00 pm close price: here the expected discounted continuation value is the conditional expectation of the discounted optimal exercise proceeds in the future conditional on the 3:20 pm unobserved (increased) fundamental value. An exercise decision would be made, if the exercise payoff exceeds the expected discounted continuation value. Thus, defined as the price difference between an OEX contract and an XEO contract with identical strike and maturity, the market EEP observed here also contains the fair value of this wild card option.

We implement the Fleming and Whaley (1994) method to determine the value of this wild card option in the OEX contract. The basic idea of the Fleming and Whaley algorithm is to first set up a binomial tree, where each step corresponds to the end of the wild card period, i.e., 3:20 pm. Second, at each node j of step t, value the wild card option and add it to the expected discounted future option value. Third, conditional on the underlying value at the node, assuming a log normal distribution for the underlying value, one can value the wild card option at note j of step t as the conditional expectation of the option "beginning ITM": i.e., the conditional expectation of the underlying being in the money at 3:00 pm when knowing its value at 3:20 pm.[17] In this case, the value of the American option at node j of step t is the maximum of the proceeds from immediate exercising and the continuation value plus the wild card option value.

[17] Note that the Fleming and Whaley algorithm is limited to the case where the underlying follows a log-normal distribution as the algorithm is based on a binomial tree, which implicitly assumes that the underlying is modeled as a geometric Brownian motion.

A practical difficulty of implementing the Fleming and Whaley (1994) method lies in choosing the right value of the standard deviation parameter of the log normal distribution for the underlying value during the wild card period. A natural response is to use the same standard deviation value implied from the cross section of European option prices that is used to determine the BS model EEP. However, as pointed out by Fleming and Whaley (1994), the realized standard deviation of the wild card period return may be higher than the realized standard deviation calculated using the 24 hour return window,[18] and clearly the European option implied standard deviation corresponds to a 24 hour window. One may scale the European option implied standard deviation by the ratio of the standard deviation of the realized return of the S&P 500 futures during the wild card period to that of the realized return of the S&P 500 futures over a 24 hour window. However, this still uses historical estimates to calculate an ex ante estimate of future volatility over the wild card period, and exposes one to a difficult problem, to convert the ratio based on the S&P 500 futures to the right one for the S&P 100 index.

To circumvent this problem, we invert the wild card period standard deviation parameter daily from the EEP in the cross section of OEX calls. This approach is justified by the fact that dividend yields are much smaller than interest rates in our sample period, and it is in general not optimal to exercise OEX calls early for interests/dividends purposes. Hence, to the best of our knowledge, the early exercises of OEX calls are predominantly motivated by the wild card option.[19] Operationally, we perform a daily RMSE minimization routine, similar to the one characterized by equations (16) to (18).

[18] In Fleming and Whaley (1994), they use the prices of the S&P 500 futures to reach this conclusion.
[19] For robustness purposes, we in addition exclude short-term ITM OEX calls for a market friction determinant of American option value, as inspired by Battalio et al. (2014).

The only difference here is that the pricing error is constructed by subtracting the Fleming and Whaley wild card option value with the OEX call EEP, where the former is a function of the implicit wild card period standard deviation.

V. Model Performance Analysis

Table 2.3 shows the estimation results. In Panel A, we report the results from estimating the five selected models, the Black-Scholes (BS), the Merton jump (MJ), the variance gamma (VG), the finite moment log stable model (FMLS), and the Heston stochastic volatility (SV). We implement the procedure mentioned by equations (2.16)—(2.18), where we choose the daily optimal set of parameters, $\Theta_{t,j}^{*}$, by minimizing the square root of daily mean squared pricing errors, $MSE_{t,j}$, for each model. In this way, the direct output from this estimation is a series of optimal parameters and a series of daily root mean squared pricing errors for each model and each day. In Panel B, we report the estimation results for the wild card period volatility parameter, based on the procedure outlined in Section IV. This exercise gives us a daily series of the volatility parameter for the risk neutral dynamics of the S&P 100 index during the wild card period.

[Insert Table 2.3 about here]

Estimated parameters in Table 2.3 are comparable with previous studies. For model parameters, we compare our estimates with previous studies using the S&P 500 index option (SPX). For example, in the BS model, the constant volatility parameter, σ_w, has a mean of around 13%, a value close to the time series average of the VXO index for the same period. In the MJ model, the mean jump intensity parameter, λ, is close to 0.66, meaning that on average investors expect 0.66 jump per year. The mean jump size

parameter μ_J is negative, indicating that jumps are on average downward. In addition, a negative μ_J also indicates a negatively skewed risk-neutral underlying distribution, which is consistent with the risk-neutral distributional properties of the S&P 500 index. Next, in the VG model, the ν parameter that controls for the kurtosis of the risk neutral return distribution is positive, indicating excess kurtosis. The θ parameter, controlling for the skewness of the risk neutral return distribution, is negative, suggesting that the return distribution is negatively skewed. Next, in the FMLS model, the parameter governing the thickness of the tail, α, has an average estimate of 1.65, a value larger than one, similar to the estimate of 1.56 reported in Carr and Wu (2003). This observation suggests that the left tail of the S&P 100 index risk-neutral distribution is slightly thinner than that of the S&P 500 index, which is no surprise since the S&P 100 index is composed of large-cap stocks. In the SV model, the average mean reverting coefficient, κ, is around 7.9, close to the value found in previous studies. The correlation parameter ρ has a value close to -0.5. This correlation parameter is negative in most studies that are based on the S&P 500 index options, meaning that volatility and return are negatively related (the leverage effect). This correlation parameter controls the skewness of the underlying distribution: where the distribution becomes more negatively skewed if ρ gets closer to -1. Lastly, for the wild card period volatility parameter σ_{wc}, our mean estimate is 0.131, larger than the mean estimate of 0.125 for the BS volatility parameter σ_W. This finding is consistent with the one documented in Fleming and Whaley (1994) that the wild card period volatility is higher than the volatility over a 24-hour window.

In terms of the in-sample fit, measured by a model's mean RMSE, we observe that jump models perform significantly better than diffusion models. In terms of the in-

sample fit, it should not be surprising to see all generalizations of the BS model perform better. However, it is surprising to see the best performance from the FMLS model in that the FMLS model only has two parameters. Based on the principle of parsimony, we argue that the FMLS model should be the best one in this regard. Such a distinguishing performance of the FMLS model is consistent with its distributional properties. On one hand, it features a slow decaying left tail, which captures the high implicit negative skewness observed in the index return risk-neutral distribution. Moreover, its infinite return variation property frees it from the rule of the central limit theorem, allowing the negative skewness to persist through the maturity dimension. Combining these two features, one can say that the extremely parsimonious FMLS model shares properties of both the other two jump models and the SV (if not others), and should perhaps be the default building block for more sophisticated models.

Having estimated all the models, we proceed to compare alternative models' performances in pricing the cross section of put EEP. In this exercise, choosing the correct pricing error metric is an important issue. According to Christoffersen and Jacobs (2004), one should be consistent in the in-sample estimation exercise and the out-of-sample prediction exercise, because otherwise the inference based on the out-of-sample prediction would be biased.[20] In this spirit, we rank alternative models based on the RMSE metric, which was the same metric in our in-sample estimation exercise.

Table 2.4 shows our out-of-sample prediction analysis results. Based on the estimated parameters for each model for the underlying process, we obtain the model predicted American option prices from the LSM algorithm, outlined in Section IV, and

[20] For example, one should not use the RMSE in the in-sample estimation exercise but use a different metric such as the mean pricing error (MPE) in the out-of-sample prediction exercise.

we also obtain the model predicted European option prices from the Monte-Carlo simulation exercise, to be consistent.[21] Hence, the difference is the model predicted EEP. Note that the model predicted EEP cannot be compared with the market EEP directly, and the same applies to model predicted American put prices, because the market EEP also contains the wild card premium. Consequently, the pricing error for EEP in this case is defined as

$$errModel = \frac{EEP^{Model} + EEP^{Wild\ Card} - EEP^{Market}}{S0} \times 1000, \tag{2.20}$$

where we normalize the error by the underlying level $S0$ and scale it by 10^4 like before. We report the RMSE summary for the overall sample in Panel A to rank the overall performances of alternative models. In addition, we also report conditional summaries in Panels B and C based on moneyness and maturity. The main reason is that since the early exercise decision is mostly relevant for ITM options, especially long term ITM options, it is crucial to examine how alternative models perform in pricing the EEP of these options. In addition, since alternative models are originally introduced to correct the moneyness-maturity associated biases in the cross section European option prices, this group based comparison allows one to better relate our findings to the existing literature.

[Insert Table 2.4 about here]

Results in Table 2.4 indicate that jump models, especially the FMLS model, are superior at pricing the cross section of EEP. For example, the RMSE of the FMLS model is around 9 basis points; based on a paired t-test, it is significantly better than that, around

[21] In the LSM algorithm, one starts with simulating a large number of sample paths for the underlying. The American option price is obtained by applying the Longstaff and Schwartz (2001) method on the simulated sample paths, while the European option price is obtained by computing the expected discounted terminal payoff using the same simulated sample paths.

10 basis points, of the second best model, the MJ model. The RMSE of the two diffusion models are significantly larger than that of the worst performing jump model, the VG model. Next, based on our group based comparisons, we can see that jump models continue dominating diffusion models. For example, looking at the three ITM categories in Panel B, where early exercises are most relevant, the FMLS model significantly outperform all other models. Once again, according to Christoffersen and Jacobs (2004), our inference regarding model performances based on this RMSE metric is statistically consistent.

In addition, the MPE summary reported in Panel C suggests that jump models are essential to explain the actual magnitudes of market EEP. For example, the average EEP for long term ITM puts is 64 basis points (Table 2.1), while the average MPE is 8 basis points for our best performing model, the FMLS model. Similarly, the average EEP for intermediate term and short term ITM puts are 38 and 18 basis points, and the average MPE are 7 and 6 basis points for the FMLS model. In sharp contrast, the average MPE for the BS model is several times larger than that of the FMLS mode. For example, the average MPE for long term ITM EEP is 29 basis points for the BS model.

The results reported in Table 2.4 reveal new insights about jump models. Existing empirical studies based on the cross section of European option prices as highlighted by Carr and Wu (2003) conclude that jumps are the reason why the prices of OTM puts, especially those of short term OTM puts, are much more expensive than that predicted by the BS model. The rationale of such a conclusion follows the distributional properties of jump models. In contrast, our findings lead to the conclusion that the actual magnitudes of market EEP, especially ITM market EEP, are much larger than that predicted by the

BS model, and jump models are essential for closing the gap between market EEP and that the EEP predicted by the BS model. Since EEP are outcomes of investors' optimal timing behaviors, our findings also reveal the fact that investors consider jumps in the underlying returns as a first-order concern when making optimal exercise decisions.

The results reported in Table 2.4 also reveal a puzzle that market EEP are overpriced, and this overpricing is more pronounced for short term ITM puts. In another word, market EEP are much larger than the combination of model predicted EEP and the wild card premia, even in the case of our best performing model, the FMLS model. Nevertheless, the degree of this overpricing is not uniform. For example, the average of the unexplained market EEP amounts to 33 percent for short term ITM EEP, while that decreases to 13 percent for long term ITM EEP. Hence, in the section to follow, we formally examine the driver for this overpriced market EEP.

VI. Does Transaction Cost in the Option Market Explain the Overpricing of Market EEP?

In the previous section, we uncover a new property of jump models: jump models are essential to explain why ITM EEP are more expensive than that predicted by the BS model. Nevertheless we find that market EEP are larger than the combination of model predicted EEP and the wild card premia, and this overpricing is more pronounced for short term ITM EEP. The objective of this section is to uncover the underlying reason for this overpriced market EEP. Since the theoretical models being considered previously all feature perfect market, here we relax that assumption and consider the role of transaction cost in the option market.

Our investigation is motivated by the idea that the presence of transactions in the market for American options can make early exercise a more profitable alternative to selling an option, provided its time value being less than the half spread. In this case, when driving by liquidity needs, an investor can exercise the option other than selling it at the bid price to save transaction cost that equals the half spread minus the time value. In this regard, the flexibility to exercise an American option early can be viewed as an option that saves transaction cost, and this flexibility should add value to American options. Hence, market EEP should contain a component arising from this "transaction cost saving option".

In Figure 2.3, we show and describe the prevalence of this transaction cost saving option in the market for OEX puts. We can see that this transaction cost saving option is present in many cases and most relevant to shorter-term ITM OEX puts. For instance, we can see it in the left figure that many cases for ITM OEX puts are present, where the time value is less than the half spread. Meanwhile, the right figure shows that this happens for OEX puts with relatively short maturity. Collectively, these two observations are consistent with our finding that the overpricing of market EEP is more pronounced for short term ITM puts and provide foundations for our hypothesis that this transaction cost saving option is the reason why market EEP are larger than the sum of model predicted EEP and the wild card premia.

To demonstrate that this transaction cost saving option is indeed the reason why market EEP are overpriced, we augment the Longstaff-Schwartz algorithm to quantitatively determine the value of this transaction cost saving option embedded in each OEX put in our sample. The main complication here is to distinguish between optimal

exercise and liquidity driven selling and exercise. In our computation, we acknowledge that investors will optimally exercise OEX puts whenever possible; when selling an OEX put prior to a model implied optimal exercise time, they will compare the time value of the put with its half spread. Further note that once investors optimally exercise the put, this transaction cost saving option is no longer relevant, since the option position is closed by the optimal early exercise. To implement this idea, we utilize the optimal exercise matrix $M_{N\times T}$, an intermediate product of the Longstaff-Schwartz algorithm. Discussed in detail in Section IV, this optimal exercise matrix shows the time for optimal early exercise within each sample path. Based on this optimal early exercise matrix, we model liquidity driven selling with a probability p, which is the expected probability that an option is exercised each day.[22] Hence, the value of this transaction cost saving option in sample path i is the expected present value of the saved transaction cost on each day going forward until the optimal exercise time as indicated by the i^{th} row of $M_{N\times T}$, where the expectation is under the measure induced by this probability p. Consequently, the value of this transaction cost saving option is the cross sample path average of these determined values.

Before we compare the remaining pricing errors with our predicted values of the transaction cost saving option based on our augmented Longstaff-Schwartz algorithm, it is informative to examine the properties of such values. For this purpose, we investigate how this transaction cost saving option value depends on model specification, moneyness, maturity, the bid-ask spread, and the liquidation probability p.

[22] Note that if $p = 0$, rational investor will never exercise OEX puts before the model determined optimal exercise time.

In Table 2.6, we provide the results of our numerical experiments. For model parameters, we use their mean estimates that are summarized in Table 2.3 to ensure all models are matched to the same set of data. In all four panels, we choose the following benchmark environmental parameters: S0=100, K = 105, T = 60, spread = 0.9 (bid = option value*spread), and p = 0.25. In each panel, we alter one particular environmental parameter while setting others to the benchmark values. For example, in Panel A, we let K vary from 90 to 110 and fix other parameters. In all four panels, we choose r=0.04 and assume no dividends, where the latter choice is to ensure that our estimates are only affected by transaction costs.[23] When reporting the results, we include both the mean estimates and their standard errors across all sample paths in our augmented Longstaff-Schwartz algorithm, scale them by the underlying level (S0=100), and multiply them by 10^4. In this case, these estimates have the same scale (basis points) as our pricing errors.

Results from Table 2.6 reveal several insights regarding the transaction cost saving option value in American puts. First of all, this option value is larger for ITM puts. For example, in Panel A, this option value is around 25 basis points for K/S0=1.1 while almost 0 for K/S0=0.9, across all five models. Besides, all three jump models predict larger values for OTM puts, consistent with the finding in Table 4 that all three jump models predict larger intermediate term OTM put EEP. Since this option value is more relevant to ITM American puts, we focus on ITM puts in the next three panels. Secondly, this option value displays an inverted U shape with respect to maturity. For example, in the BS model, this option value attains its maximum of 27 basis points in T=60 days, yet

[23] In our empirical analysis later, we use the actual dividend series like that in our other pricing analyses, nevertheless it is clear that the presence of non-zero dividends will decrease the value of this transaction cost saving option.

this optimal maturity differs across models. Thirdly, from Panel C, we can see that this option value declines monotonically with the spread. This observation is consistent with the idea that this transaction cost option value will become larger if the transaction cost is higher. From Panel D, we can see that this option value increases monotonically with the probability of liquidation p. This observation indicates that a higher need to sell the option will make this transaction cost saving option more valuable. In addition, the results from all four panels show that no model dominates others in pricing this transaction cost saving option. For example, the BS model generates the highest values in Panel C but the three jump models generate the highest values in Panel D. This observation implies that this transaction cost saving option value not only depends on model implied optimal exercising rule but also on model predicted option values.

Having established several properties for the transaction cost saving option value, we now proceed to determine the actual value of this option utilizing our sample information. The complication of this analysis is the unobservable expected liquidation probability p. To find the best possible proxy for it, we use the realized selling intensity ($\widehat{p_{t,K,T}}$) of an option series on a given day defined as $\widehat{p_{t,K,T}} = \frac{sell_{t,K,T}}{open\ interest_{t,K,T}}$ in analogous to the "expectation hypothesis". Note that the quantity $sell_{t,K,T}$ is the number of contracts sold, identified with the Lee and Ready (1991) algorithm, and the subscript t, K, T refer to the option series with strike K maturity T on trading date t. Further note that not all options are traded in a given day nor have positive open interest, making this probability undefined in many cases. Instead of assigning each option series a probability each day, we construct a moneyness-maturity group based probability on a daily basis, since theoretically options in the same moneyness-maturity group should have the same

probability of liquidation. To be consistent with our previous analyses, we use the same nine moneyness-maturity groups as that in Table 2.4.

[Insert Table 2.6 about here]

In Table 2.6, we report the power of this transaction cost saving option in explaining the remaining pricing errors that are summarized in Table 2.4. Guided by the same principle, we compare alternative models based on the RMSE metric, where we construct the pricing error for each option series at each day by subtracting from the market EEP the model predicted EEP, the wild card option value, and the estimated transaction cost saving option value, and we scale this difference by the underlying level and multiply it by 10^4. We report the mean and its standard error for the RMSE time series. Next, to assess their actual magnitudes, we also construct a MPE metric for these pricing errors. In this case, we can examine how well these determined transaction cost saving option values can match the magnitudes of the remaining pricing errors. In addition, we regress the remaining pricing errors excluding our model determined transaction cost saving option values on transaction cost saving option values. Based on this exercise, we can further evaluate how well these transaction cost saving option values can match the cross-sectional variation in the remaining pricing errors.

The estimates in Table 2.6 provide evidence that the presence of this transaction cost saving option in OEX options is the reason why market EEP are overpriced. First, results in Panel A lead to the consistent conclusion that the FMLS model is the best model for the risk-neutral underlying process. This statement is supported by the fact that the minimum RMSE, 5 basis points, appears in the column for the FMLS model. In addition, recall that the RMSE for the FMLS model was 9 basis points in Table 2.4, where we did

not account for the transaction cost saving option values. This comparison indicates that accounting for the transaction cost saving option values results in significantly improved overall fit of market EEP. Second, results in Panel B suggest that the model determined transaction cost saving option values can match the magnitudes of the remaining pricing errors well. For example, the MPE of the FMLS model, the best model, is 0.9 basis points, which amounts to 5 percent of the mean market EEP. In this sense, it suggests that with a 5 percent confidence level market EEP can be completely explained by the sum of model determined EEP, wild card premia, and transaction cost saving option values. Next, results in Panel C show that the model determined transaction cost saving option values can match the cross-sectional variation in the remaining pricing errors (Market EEP-Model EEP-Wild Card Premia). The coefficients of the transaction cost saving option value are positively significant across all models. Note that none of the coefficients are close to one, which should be in the ideal case. Nevertheless this observation is not surprising because the presence of measurement errors in the explanatory variables will bias the coefficients towards zero, where such measurement errors originate from our choice of the realized probability of liquidation instead of the expected probability of liquidation.

VII. Robustness

In the previous sections, we reach two conclusions. First, based on the RMSE metric, we conclude that jump models, especially the FMLS model, perform best in pricing the cross section of market put EEP. Second, after incorporating the model predicted transaction cost saving option values, we find the overpricing of market EEP is

resolved based on the results for the FMLS model. To demonstrate that our conclusions are not driven by our chosen empirical methodology, we resort to another test, the orthogonality test, that is popular in the in the empirical European option pricing literature (e.g., Madan et al. (1998); Bates (2000)).

The underlying idea of this orthogonality test follows the fact that generalized option pricing models, such as jump models, stochastic volatility models, or their affine combinations, are introduced to correct the biases associated with the BS model in matching the implied volatility surface (IVS). Hence, if a generalized model is successful in matching the IVS, then the remaining pricing errors should exhibit little dependence on moneyness and maturity. We acknowledge that the cross section of EEP is driven primarily by the early exercise boundaries and the associated first-passage probabilities, whose relation with the IVS is not clear. Yet we still believe this test can be informative about the performances of alternative models.

In Table 2.7, we report our results for the orthogonality test. The dependent variable is the remaining pricing error, constructed by subtracting model EEP, wild card premia, and the transaction cost saving option values from market EEP. The independent variables are moneyness, moneyness-squared, maturity, interest rates, and dividend yields, where moneyness-squared is included to account for any nonlinearity associated with moneyness. To better interpret the regression estimates, we standardize all the explanatory variables.

[Insert Table 2.7 about here]

The results in Table 2.7 reinforces our conclusion that jump models, especially the FMLS model, perform best in pricing the cross section of put EEP. Recall in Table

2.2 that the R-squared is 0.737 in the regression of market put EEP on this set of explanatory variables, however the R-squared here decreases sharply for each model. In particular, the R-squared for the FMLS model is only 0.078. This result indicates that the FMLS model is able to explain the majority of the cross-sectional variations in market EEP, leaving the explanatory power of option characteristics small.

VIII. Conclusions

The S&P 100 index has both European (XEO) and American style (OEX) options traded on it, which makes the market early exercise premium directly observable. We focus on the EEP in American put prices for period 2005.05—2007.12, when short-term interest rates are much larger than the dividend yields on the S&P 100 index such that the early exercise for American puts can be optimal. We use high quality transactions data to provide, for the first time, a comprehensive empirical study on the relative merits of alternative generalizations of the BS model for improving the pricing quality of EEP. Specifically, we study four parsimonious generalizations of the BS model, namely the Merton jump diffusion model, the variance gamma model, the finite moment log stable model, and the Heston stochastic volatility model, motivated by the distinct sample path properties.

With our carefully designed tests and our acknowledgement of the wild card option embedded in OEX options, we generate unique insights regarding market EEP as well as alternative models, when all models are estimated using European option prices. We find that jump models are superior at matching the cross-sectional variation and magnitudes of market EEP. Since EEP depend directly on the model implied early exercise boundaries, this finding suggests that investors consider jumps as the first-order

concern when valuing American options' flexibility to exercise early for earning interests purposes, as in classic theories (Carr et al. (1992)). This finding also generates new insights about jump models. Conventional wisdom suggests that jump models are essential to explain why market overprices OTM European puts relative to the BS model, while our results in contrast demonstrate that jump models are essential to explain why market overprices ITM put EEP relative to the BS model.

We in addition document that all models underprice market EEP and this underpricing is more pronounced for short term ITM puts. We provide an explanation based on the unique feature of American options, the flexibility to exercise early, which generates a transaction cost saving option. When the half spread of an American option exceeds its time value, an investor can exercise the option early as a more profitable alternative to selling it. Following this idea, we augment the Longstaff-Schwartz algorithm to quantitatively determine the value of this transaction cost saving option for each option series in our sample. Based on its explanatory power and the empirical regularity that the highest proportion of short term ITM OEX puts exhibit larger half spreads than time values, we find that this transaction cost saving option is the reason why market EEP are larger than the combination of model predicted EEP and wild card premia.

Chapter 3:

Infinite versus Finite Jump Processes in Commodity Futures Returns: Crude Oil and Natural Gas[24]

I. Introduction

Commodity futures serve both as an important investment vehicle for gaining exposure to physical commodities as well as a crucial centerpiece of the risk management strategies employed by both investors and businesses seeking to hedge commodity price exposure. Crude oil and natural gas futures are widely used for risk management but in recent years have also come to be appreciated as pure investment vehicles, a good example being the advent of the United States Oil Fund that follows a strategy of holding long positions in the nearby NYMEX crude oil contract (Gorton et al. 2012). For instance average daily open interest in exchange traded oil futures contracts increased 227% between 2000 and 2015 (from 515 x 10^3 contracts in 2000 to 1,686 x 10^3 contracts in 2015 (U.S. EIA, https://www.eia.gov/finance/markets/financial_markets.cfm). However this actually understates the economic importance of the oil derivatives market as the vast majority of positions are in the form of the less transparent over-the-counter (OTC) forwards and swaps.

The behavior of oil and natural gas futures returns have in recent years exhibited volatility far in excess of financial futures, such as stock index futures and treasury futures.[25] Plots of the absolute values of log price changes shown in Figures 3.1 and 3.2

[24] This chapter is based on collaborative work with Scott Guernsey and Scott Linn.
[25] We follow general convention and for convenience often will refer to log price changes in futures prices as returns.

suggest jumps are potentially an important driving factor. This has not gone unnoticed by both policy makers, hedges and speculators as well as investors. For instance, Timothy Massad, chairman of the U.S. Commodity Futures Trading Commission has remarked that "hourly flash events" in recent times have occurred predominantly in the crude oil futures market and less so in other futures markets, including the markets for the E-mini futures and 30 Year Treasury futures.[26] Large price changes of this sort are important because they directly impact investors' and hedgers' marked-to-market trading positions. As Bakshi and Panayotov (2010) have noted, exposure to potential margin calls increases substantially when the values of asset holdings exhibit jumps.

Given the expansion of these markets over the last decade and the practical significance of jumps in oil and natural gas futures prices a clear characterization of those jumps will both help refine our understanding of the underlying processes as well as sow the seeds for analysis of the implications for derivative valuation as well as risk management practices. The existing pool of research emphasizes the behavior of stochastic convenience yields and unspanned stochastic volatility when modeling the term structure of oil and gas futures returns. Notable examples include Schwartz (1997) and Trolle and Schwartz (2009).

The aim of the present study is to formally test for the existence of jumps and to characterize the nature of jumps in oil and natural gas futures returns. This inquiry is of importance since it addresses a fundamental question about model specification. In

26 Source: "Remarks of Chairman Timothy Massad before the Conference on the Evolving Structure of the U.S. Treasury Market", October 21, 2015. Hourly flash events are defined as episodes in which the price of a contract moved at least 200 basis points within a trading hour-but returned to within 75 basis points of the original or starting price within the same hour.

addition, correctly specified models are crucial for investment and risk management purposes.

We utilize recent developments and tests that are 'model free' to address four specific questions regarding the behavior of oil and natural gas futures prices.[27] (1) Whether jumps are present; (2) whether infinitely many small jumps in contrast to few but large jumps are present; (3) whether a smooth Brownian motion component is present in the general semi-martingale process; and (4) the contribution of jumps to total return variation. We take an agnostic view about the true nature of the return dynamics, yet we do require that the returns follow a general Ito semi-martingale process that may consist of a drift term, Brownian motion, and jump terms. We make only mild assumptions about the drift and the instantaneous volatility of the Brownian motion terms, assuming only that both can be described by general stochastic processes.[28] This allows us to focus our attention on the jump term, which can be decomposed into the sum of finitely many large jumps and a large number (infinitely) many small jumps. Large jumps are rare and are commonly associated with economy-wide news or events (Barro (2006)). In contrast, small jumps are generally attributed to market microstructure factors impacting prices at the intraday level. The merit of the model free method is its freedom from imposing a restrictive parametric form, allowing a more flexible and robust diagnosis of the data generating process.

The data we examine include front month contract prices recorded every 5-seconds for the period 2006-2014 for the crude oil and natural gas futures traded on

[27] See Ait-Sahalia and Jacod (2012) for a systematic overview of such model free tests, and Ait-Sahalia and Jacod (2009, 2010, 2011) for a formal treatment.
[28] In this sense, the structure may include mean reversion as well as smoothly time varying volatility.

NYMEX. Theoretically, jumps take place instantaneously, and in an ideal setting of continuous sampling it would be possible to exactly identify them. In practice, because of discrete sampling, we must make inferences based upon statistical tests. Since the asymptotic limits of the selected tests are based on a sampling interval that approaches zero, we work with prices observed at a high frequency (5 second interval here). In addition, guided by the insight and results in Bajgrowicz et al. (2015), we compute a statistic for each test based on the entire sample instead of each trading day to avoid the issue of multiple testing.[29] However, we do also present results after disaggregating the sample into calendar years.

We find that an infinite activity jump diffusion process describes crude oil and natural gas futures returns. We reject the null hypothesis that jumps are not present in either return series and further reject the null hypothesis that jumps in either return series only exist in large finite form. However, we also find that we cannot reject the hypothesis that Brownian motion is present in both return series, despite the presence of infinitely many small jumps. We end by exploring the contributions to overall variability of jumps in the two series. In the crude oil return series, jumps in total account for around 36 percent of the total return variation and large jumps contribute nearly 3 times more than small jumps. Jumps collectively account for around 41 percent of the total return variation in the natural gas return series, while large jumps also contribute almost three times the amount of small jumps.[30]

[29] Bajgrowicz et al. (2015) show that constructing the test statistic on a daily basis may lead to biased inferences regarding the presence of jumps.

[30] In contrast, Huang and Tauchen (2005) find the contribution of jumps to be around 7 percent based on 5 minute S&P 500 index series.

While several authors have explored the presence of jumps in oil and natural gas prices observed at the daily, weekly and monthly frequencies (Askari and Krichene, (2008); Lee, Hu and Chiou (2010); Gronwald (2012); Nomikos and Andriosopoulos (2012); Trolle (2014)) those studies do not address the question of whether infinite jump activity is present. Our study contributes to this literature but asking and answering the question of whether infinite jump activity is present and the extent to which infinite as well as finite jump activity drive overall price change variability.

Closest in spirit to our study is a study that examines daily data by Askari and Krichene (2008). Those authors estimate a set of parametric jump diffusion models that feature both finite jump activity and infinite jump activity and find the parameters of the jump component to be significant. Our paper is fundamentally different from these empirical studies, because we take a model free approach implemented on high frequency data. Conclusions based on model free approaches are more robust than those based on parametric models, and the examination of high frequency data allows us to cleanly distinguish between infinite versus finite jump process activity, which is hindered by data measured over longer time horizons.

II. Data

We study high frequency price data for crude oil and natural gas front month futures contract prices sourced from Tick Data, Inc. The sampling frequency is 5 seconds and spans 2006-2014. We study log price changes observed during trading days,

excluding overnight returns.[31] There are 8,932,703 data points for each return series. Table 3.1 provides annual and full sample descriptive statistics for the final sample of log-returns for both series.

[Insert Table 3.1 about here]

Table 3.1 emphasizes those measures that would be most reflective of the influence of jumps, the standard deviation, skewness and kurtosis. The full crude oil sample standard deviation is 0.3652.[32] The distribution is positively skewed (15.9) and kurtosis equals 10.6130 for the 2006--2014 period, clearly indicating that the sample is non-normal. The mean 5-second log-return for the natural gas sample equals -0.1102 with a daily full sample mean estimate of -436.5108. The minimum 5-second log-return for the same period is -0.2 while the maximum is 0.24. The sample standard deviation equals 0.5101. The natural gas data are positively skewed (24.1) and fat-tailed (15.8580) non-normal distribution. Generally the natural gas data exhibit greater volatility.

Figures 3.1 and 3.2 present a time-series plots of the absolute values of log price changes for each of the commodities' front month contract prices. The natural gas series (Figure 3.2) exhibits marginally higher volatility. Moreover, both series exhibit numerous sharp increases in volatility throughout the sample period. The crude oil log-return series exhibits its most pronounced turbulence during 2008-2009, however spikes in volatility occur throughout the sample period. We also see high turbulence in the natural gas log-returns during 2008-2009 but again spikes in volatility occur throughout. A histogram of the log-returns is shown in Figure 3.3. Note that the tails of the distributions are not

[31] In the infrequent event that either the crude oil or natural gas price series does not trade at a given 5-second interval we linearly interpolate when computing missing value to ensure an evenly-spaced sample.
[32] Standard deviation, skewness and kurtosis reported in Table 1 are the 5 second sample estimates times 10^3, while the minimum and maximum values are the sample values x 10^6

indicative of a normal distribution, which is consistent with our conclusions from the descriptive statistics. Our full samples are positively skewed and display overwhelming evidence of fatter tails. Hence, the preliminary diagnostics for both the crude oil and natural gas samples are consistent with the presence of jump activity.

[Insert Figure 3.1 about here]

[Insert Figure 3.2 about here]

[Insert Figure 3.3 about here]

III. Modeling Oil and Natural Gas Futures Price Changes

We do not make any explicit assumption about the parametric family of models for the log price series. We require only a modest structural assumption: Let $X_t = lnS_t$ denote the log price series and assume X_t follows a general Ito semi-martingale described by the following

(3.1) $$X_t = X_0 + \int_0^t b_s ds + \int_0^t \sigma_s dW_s + Jumps$$

(3.2) $$Jumps = \int_0^t \int_{\{|X| \leq \epsilon\}} x(\mu - \nu)(ds, dx) + \int_0^t \int_{\{|X| > \epsilon\}} x\mu(ds, dx)$$

where W_t denotes a standard Brownian motion term, μ is the jump measure of X_t, and its predictable compensator is the Levy measure ν.[33]

The first term in equation (3.2) represents innovations from "small" jumps, those with size less than a fixed arbitrary cutoff level ϵ, and the second term corresponds to "large" jumps, innovations with size larger than ϵ. For instance, large jumps could arise from rare economic events (Barro (2006)) at a frequency of daily or greater; small jumps

[33] For a formal treatment of the assumptions and conditions underlying the process, please refer to Ait-Sahalia and Jacod (2009).

may be driven by market microstructure effects occurring intraday. Mathematically, a semi-martingale process will always generate a finite number of large jumps on a finite time interval ($[0, T]$), as pointed out by Ait-Sahalia and Jacod (2011). Depending on the jump size behavior near the origin, we can classify a semi-martingale process as exhibiting finite or infinite jump activity. Notable examples of finite activity jump models include the models developed in Merton (1976) and Kou (2002). On the other hand, infinite activity jump models include the variance gamma (VG) model of Madan (1990) and Madan et al. (1998), the normal inverse Gaussian (NIG) of Barndorff-Nielson (1997), the CGMY model of Carr et al. (2002), and the finite moment log stable (FMLS) model of Carr and Wu (2003).

The baseline specification described in equation (3.1) accommodates both jumps as well as other sources of variability. In (3.1) both the instantaneous drift (b_t) and volatility (σ_t) can follow very general processes. For instance, the drift term can reflect a stochastic interest rate or convenience yield, as in the model of Schwartz (1997), and the volatility term can admit unspanned stochastic volatility, as in Trolle and Schwartz (2009). Generally, as pointed out by Ait-Sahalia and Jacod (2009), the volatility (σ_t) can follow an Ito semimartingale, which can exhibit jumps.

The distinction between jumps and continuous innovations comes from a discontinuity in the sample path of the variable of interest. Large changes can however arise when the sample path is continuous. Consider for example a constant volatility model, in which case the probability of extreme movements in the sample path is tiny regardless of the specification of the drift component. When the instantaneous volatility is however allowed to be stochastic, the probability of extreme movements is increased

significantly, nevertheless if this was the only driving force the sample path would still be continuous. In contrast jumps introduce discontinuities. The identification of jumps is therefore equivalent to identifying discontinuous movements in the sample path of a stochastic process.

IV. Do Oil and Natural Gas Returns Exhibit Jumps?

Based on the general framework outlined in Section III, consider a series $\{X_i\}_{i=1}^{N}$ within a time window $[0,T]$ at a discrete sampling frequency of Δ_n and $N = [T/\Delta_n]$, denotes the integer part of the ratio. We construct the increment $\Delta_i^N X = X_{i\Delta_n} - X_{(i-1)\Delta_n}$ as the empirical approximation of its theoretical limit $\Delta_s X = X_s - X_{s-}$. If X displays a discontinuous movement at t_0 in $[0,T]$, then $\Delta_{t_0} X$ will be nonzero. In contrast, $\Delta_{t_0} X$ would be zero if X is continuous in $[0,T]$. Empirically, since Δ_n cannot go to the limiting case, one may observe movements that appear to be jumps but are not. Any statistical test must have power to distinguish jumps from continuous path changes. At the same time the test should allow for a general specification of the continuous drivers of the process. The asymptotic model-free jump test developed in Ait-Sahalia and Jacod (2009) is our test statistic choice.

The essence of the jump test is to examine the ratio of power variations of the returns at two different sampling frequencies, which can cause different values of the ratio depending on whether jumps are present. Here the power variation of return measures the variability of return and is defined generally as

(3.3) $$B(p,\alpha,\Delta_n) = \sum_{i=1}^{N} \left|\Delta_i^N X\right|^p 1_{\{|\Delta_i^n X| \le \alpha \sigma_c \Delta_n^{\bar{\omega}}\}},$$

where p is the power, σ_c indicates the standard deviation of the Brownian component of X, Δ_n is the sampling interval (e.g., 5 seconds), $\varpi \in (0, \frac{1}{2})$, and α is a scaling parameter for the cutoff level.[34] As shown by Ait-Sahalia and Jacod (2009), $B(p, \infty, \Delta_n)$ displays different asymptotic limits (as Δ_n goes to zero) for $p > 2$, $p = 2$, and $p < 2$, and for cases where X displays jumps or not. Specifically, if X displays jumps, the asymptotic limit of $B(p, \infty, \Delta_n)$ will not depend on the sampling frequency Δ_n for $p > 2$. In contrast, if X is a continuous process, the asymptotic limit of $B(p, \infty, \Delta_n)$ will depend on the sampling frequency Δ_n for $p > 2$.

This sharp contrast of the behavior of $B(p, \infty, \Delta_n)$ leads to the construction of the following test statistic S_J, defined as

$$S_J = \frac{B(p, \infty, k\Delta_n)}{B(p, \infty, \Delta_n)}, \tag{3.4}$$

where $k > 1$ is a scaling coefficient for sampling frequency. The test statistic S_J will converge asymptotically ($\Delta_n \to 0$) to 1 under the null hypothesis that a jump process is present or to $k^{\frac{p}{2}-1}$ under the null of no jump process (Ait-Sahalia and Jacod, 2009). The intuition for the property of S_J follows the behavior of $B(p, \infty, k\Delta_n)$. When X jumps, the asymptotic limit of $B(p, \infty, \Delta_n)$ does not depend on Δ_n, meaning $B(p, \infty, \Delta_n)$ and $B(p, \infty, k\Delta_n)$ will have the same limit, thus S_J goes to 1 in the limit. When X is continuous, the asymptotic limit of $B(p, \infty, \Delta_n)$ depends on Δ_n, meaning $B(p, \infty, \Delta_n)$ and $B(p, \infty, k\Delta_n)$ will not have the same limit, thus S_J goes to $k^{\frac{p}{2}-1}$ in the limit. For instance, if $k = 2$ and $p = 4$, S_J will converge to 2 under the null of no jumps. Hence,

[34] We introduce the freedom for truncation here albeit we are not truncating, i.e., $\alpha = +\infty$. We introduce the truncated power variation to maintain consistency with subsequent analyses, where we will examine cases for different α values.

if the computed S_J is less than $C_{95}^{Normal} = 2 - z_{95}\sqrt{V_J^C}$, where Z_{95} represents the 95 percent normal quantile (1.67), and V_J^C is the asymptotic variance of the test statistic under the null of no jump process, we can reject the null of no jump process at the 95 percent level of confidence. Note that the choice of k and p here is arbitrary, as long as they satisfy the condition that $k > 1$ and $p > 3$. When testing under the null that a jump process is present, one can choose any pair of k and p, and with a given choice one will have the corresponding test statistics S_J and a critical value for the test statistics in order to reject the null. Ait-Sahalia and Jacod (2009) recommend the choice of $p = 4$ and $k = 2$ and validate the recommendation with simulation analysis.

When testing hypothesis based on the asymptotic test derived from the central limit theorem, one needs an estimate of the variance (V_J^C) of the test statistics under the null of no jump. Ait-Sahalia and Jacod (2009) show that the variance is given by

$$V_J^C = \frac{\Delta_n M(p,k) A(2p,\alpha,\Delta_n)_t}{A(p,\alpha,\Delta_n)_t^2}, \tag{3.5}$$

where $A(p,\alpha,\Delta_n)$ is defined as

$$A(p,\alpha,\Delta_n) = \frac{\Delta_n^{1-\frac{p}{2}}}{m_p} B(p,\alpha,\Delta_n). \tag{3.6}$$

Note that in (3.5) and (3.6), both the two scaling coefficients, $M(p,k)$ and m_p, and the variance V_J^C depend on α, the parameter that controls the cutoff level. According to Ait-Sahalia and Jacod (2009), the actual value of α is irrelevant asymptotically; however, in finite sample studies one can choose α to be larger than 3, meaning to sum p^{th} powers of return increments with magnitudes more than 4 standard deviations away. In our analysis,

we determine the variance and the associated rejection region for a range of αs, for robustness purpose.

In Table 3.2 we present the test results for the presence of jumps. The table reports the test statistics S_J and the 95 percent critical value, C_{95}^{Normal}, in order to determine a rejection decision for the null of no jumps based on a range of cutoff levels $\alpha\sigma_c$, where σ_c denotes the estimated standard deviation of the continuous component.[35] We choose $p = 4$ and $k = 2$ as suggested by Ait-Sahalia and Jacod (2009). Such values correspond to comparing the 4^{th} power variations at $\Delta_n = 5$ seconds and $2\Delta_n = 10$ seconds. In Panel A, we report the results for the crude oil sample, and Panel B reports the results for the natural gas sample.

[Insert Table 3.2 about here]

The results presented in Table 3.2 indicate that we can reject the null of no jumps in both the crude oil and natural gas samples. In Panel A, the S_J values should equal 2 under the null of no jumps and 1 under the alternative hypothesis of jumps. The asymptotic variances for S_J are computed based on a range of α values from 4 to 10. The rejection thresholds are increasing in α yet the empirical S_J are still less than the minimum of the thresholds at $\alpha = 10$. This suggests that even under the extreme case (increments with size larger than 10 times the standard deviation of the continuous component) we still reject the null of no jumps. The result is similar in Panel B for the natural gas sample. Taken together, the null hypothesis of no jumps is rejected for both crude oil and natural gas futures return series.

[35] Here we follow Ait-Sahalia and Jacod (2009) and construct a truncated quadratic variation estimator for the σ_c parameter.

V. Are There Infinitely Many Jumps in Oil and Natural Gas Returns?

Having rejected the null of no jumps, we next turn to the question of whether the series are characterized by finite or infinite jump activity. The test statistic S_F is calculated for this purpose and is defined as

(3.7) $$S_F = \frac{B(p,\alpha,k\Delta_n)}{B(p,\alpha,\Delta_n)},$$

where $B(p,\alpha,\Delta_n)$ is given by equation (3.3) and $p > 2$ (see Ait-Sahalia and Jacod (2011)). S_F is the ratio of truncated p^{th} power variations computed at different sampling frequencies, $k\Delta_n$ and Δ_n.

The test statistic $S_F$converges asymptotically ($\Delta_n \to 0$) to 1 under the null of infinite activity and to $k^{\frac{p}{2}-1}$ under the null of finite activity (Ait-Sahalia and Jacod (2011)). For instance, if $k = 2$ and $p = 4$, then S_F will converge to 2 under the null of finite activity. Hence, if the computed S_F is less than $C_{95}^{Normal} = 2 - z_{95}\sqrt{V_F}$, where Z_{95} represents the 95 percent normal quantile (1.67), and V_F is the asymptotic variance of the test statistic under the null of finite activity, then we can reject the null of finite activity at a 95 percent level of confidence.

The intuition for the behavior of S_F is as follows. Under the null of finite jump activity, meaning the number of jumps is finite within a finite time window, the truncation eliminates these jumps at some point along the asymptotic limit. The truncated power variation behaves close to the continuous part of the semi-martingale in this case. Thus in the limit the test statistics S_F under the null of finite activity converges to the same limit as S_J under the null of no jump. Under the alternative hypothesis of infinite jump activity, the truncation would never be able to fully eliminate jumps. With $p > 2$ the

power variation is dominated by jump increments, and for a sample size large enough the sampling frequency would not affect the value of the power variation. Thus S_F converges to 1 under the alternative hypothesis of infinite activity.

The test of the null of finite activity is an asymptotic test derived from the central limit theorem. The variance (V_F) of the test statistics under the null is given by

$$V_F = \frac{N(p,k)B(2p,\alpha,\triangle_n)}{\left(B(p,\alpha,\triangle_n)\right)^2}. \tag{3.8}$$

Table 3.3 shows the results of tests of the null of finite jump activity for both the crude oil (Panel A) and natural gas (Panel B) price series. We compute the test statistics S_F for $p = 4$ and $k = 2$, as recommended by Ait-Sahalia and Jacod (2011). For robustness, we also consider a range of α, from 4 to 10.

[Insert Table 3.3 about here]

The results reported in Table 3.3 indicate that we can reject the null of finite activity in both the crude oil and natural gas samples. In Panel A, S_F should equal 2 under the null of finite activity and 1 under the alternative hypothesis of infinite activity. The rejection thresholds are increasing in α, yet the empirical S_F are still less than the rejection threshold in all cases. This observation suggests that even under the extreme case where increments with sizes larger than 10 standard deviations of the continuous component are considered as coming from large jumps, we still reject the null of finite jump activity. The result is similar in Panel B for the natural gas sample. Taken together, the results indicate that both return series exhibit infinite jump activity.

VI. Is Brownian Motion Present?

Thus far we have shown that the crude oil and natural gas returns series both exhibit infinite jump activity. One natural question remains: Should we take it for granted that a Brownian motion component is also present in the return series? For example, Carr et al. (2002) in an examination of a set of equity and equity index options conclude that pure infinite jump processes are adequate to model risk neutral stock index dynamics. We examine the question of whether a Brownian motion component is present in our sample data, given that the dynamics exhibit infinitely many small jumps. Note that the term "Brownian motion" here is not restricted to the standard Brownian motion. It also admits the case of stochastic volatility, where the instantaneous variance of the Brownian motion term is stochastic.

The test statistics S_W is defined as

$$S_W = \frac{B(p,\alpha,\triangle_n)}{B(p,\alpha,k\triangle_n)}, \tag{3.9}$$

where $p \in (0,2)$, according to Ait-Sahalia and Jacod (2010), who prove that under the null when Brownian motion is present, S_W converges asymptotically ($\triangle_n \to 0$) to $k^{1-p/2}$ for some $p \in (1,2)$ and $k \geq 2$. For instance, if $k = 2$ and $p = 1.5$, then S_W will converge to $2^{0.25} = 1.189$ under the null that Brownian motion is present. Hence, if the computed S_F is less than $C_{95}^{Normal} = 1.189 - z_{95}\sqrt{V_W}$, where Z_{95} represents the 95 percent normal quantile (1.67), and V_W is the asymptotic variance of the test statistic under the null that Brownian motion is present, we can reject the null at a 95 percent level of confidence.

To clarify the test, suppose that X contains a Brownian motion component, then the test statistics S_w converges to $k^{\frac{p}{2}-1}$. This happens because lower powers magnify

continuous increments; then for $p \in (0,2)$, $B(p,\alpha,\Delta_n)$ is dominated by continuous increments instead of jumps. For the sample size to be large enough, a property of Brownian motion suggests that $B(p,\alpha,\Delta_n)$ will depend on the sampling interval Δ_n, and $B(p,\alpha,\Delta_n)$ will not be the same when computed under Δ_n and $k\Delta_n$. Hence, as shown by Ait-Sahalia and Jacod (2010), S_w converges to $k^{\frac{p}{2}-1}$.

To reject the null that Brownian motion is present, we also rely on the asymptotic test derived from the central limit theorem. Similarly, the variance (V_W) of the test statistics under the null is given by

$$V_W = \frac{N_W(p,k)B(2p,\alpha,\Delta_n)}{\left(B(p,\alpha,\Delta_n)\right)^2}. \tag{3.10}$$

Table 3.4 shows the results for testing the null that Brownian motion is present for both the crude oil (Panel A) and natural gas (Panel B) data. Note the key distinction in calculating S_W is that we are using $p = 1.5 < 2$ instead of 4 when calculating S_J and S_F.

[Insert Table 3.4 about here]

The results reported in Table 3.4 indicate that we cannot reject the null that a Brownian component is present in both the crude oil and natural gas samples. In Panel A, the S_W should equal $2^{0.25}$ (≈ 1.19) under the null. The asymptotic variance for S_W is computed based on a range of α from 4 to 10, and the implied 95 percent rejection thresholds are reported. The test statistic S_w is larger than the 95 percent rejection thresholds regardless of the α value. Thus, we cannot reject the null of the presence of a Brownian component regardless of how we construct the test statistic. The result is similar in Panel B for the natural gas sample. Collectively, we cannot reject the null that a Brownian component is present in both return series.

VII. The Contribution of Jumps to Total Return Variations

To what extent do jumps contribute to total return variation is a natural question that emerges from the conclusions that both series exhibit infinite jump activity as well as Brownian motion. We define total return variation as the quadratic variation (QV) or the untruncated second order variation $B(2, \infty, \Delta_n)$. As discussed before, jumps and continuous increments contribute equally to $B(p, \infty, \Delta_n)$, when $p = 2$. In this sense, the QV is the best metric to examine the relative contributions of jumps and continuous increments.

The QV can be decomposed as

$$\begin{cases} \frac{B(2,\alpha,\Delta_n)}{B(2,\infty,\Delta_n)} & \%\ QV\ due\ to\ the\ continuous\ component \\ 1 - \frac{B(2,\alpha,\Delta_n)}{B(2,\infty,\Delta_n)} & \%\ QV\ due\ to\ the\ jump\ component \end{cases} \tag{3.11}$$

Note that the numerator is the truncated second order power variation $B(2, \alpha, \Delta_n)$, which is affected by the cutoff level α. Choosing α is equivalent to counting the increments with magnitudes less than α standard deviations of the continuous component. Given the conditional normality of the continuous component, an α with a value no less than 3 indicates extreme movements. We report our results based on a range of αs.

One can further decompose the QV due to jumps into the following

$$\begin{cases} \frac{U(2,\epsilon,\Delta_n)}{B(2,\infty,\Delta_n)} & \%\ QV\ due\ to\ large\ jumps \\ 1 - \frac{B(2,\alpha,\Delta_n)+U(2,\epsilon,\Delta_n)}{B(2,\infty,\Delta_n)} & \%\ QV\ due\ to\ small\ jumps \end{cases} \tag{3.12}$$

where $U(2, \epsilon, \Delta_n) = B(2, \infty, \Delta_n) - B(2, \epsilon, \Delta_n)$, meaning truncation from below, defined for an arbitrary ϵ. Note that the upper truncated power variation $U(2, \epsilon, \Delta_n)$ can be considered as the dual of the lower truncated power variation $B(2, \epsilon, \Delta_n)$, where the former focuses on large increments and the latter focus on small increments. We choose

$3 < \alpha < \epsilon$, indicating that we define increments with magnitudes within $(\alpha\sigma_c, \epsilon\sigma_c]$ as coming from small jumps and increments with magnitudes within$(\epsilon\sigma_c, \infty)$ as coming from large jumps, where σ_c is for the standard deviation of the continuous component. We report the results based on a range of ϵ.

In Table 3.5 we present the results of decomposing the return QV into contributions from large jumps, small jumps, and the Brownian component. Figures 3.4, 3.5 and 3.6 present plots of the relative contributions by year. The results are reported for a range of cutoff levels α at the first level to separate the QV due to jumps from those due to the Brownian component. After we decompose the QV into the relative contributions from jumps and the Brownian component, we further decompose the jump contribution into contributions from large jumps and small jumps by selecting three different values of ϵ, $1.5 \times \alpha$, $3 \times \alpha$, and $4.5 \times \alpha$. We find the results are not sensitive to the choice of ϵ.

[Insert Table 3.5 about here]

The results reported in Table 3.5 reveal that jumps contribute significantly to the QV of both the crude oil and natural gas return series. In panel A, the crude oil sample, the contribution of the jump component ranges from 38.71% to 34.12% for α ranging from 4 to 10 with an average of 36.10%. In panel B, the natural gas sample, the contribution of the jump component ranges from 44.05% to 38.23% for α ranging from 4 to 10 with an average of 40.62%. Note that the contribution of the jump component is decreasing in α, as expected, since selecting a larger value of α is equivalent to treating variations with larger magnitudes as coming from the continuous component. Recall that we are conservative in our analysis here in terms of the large α values being considered.

Since the continuous component is conditionally normal, increments larger than 3 standard deviations ($\alpha > 3$) are considered rare.

In addition, the results indicate that the contribution of large jumps exceeds that of small jumps in both the crude oil and natural gas return series. In Panel A, the crude oil sample, the average (in ϵ) contribution of large jumps ranges from 32.67% to 21.87% for α ranging from 4 to 10 with an average of 26.94%, while the average contribution (in ϵ) of small jumps ranges from 6.04% to 12.25% with an average of 9.16%. In Panel B, the natural gas sample, the average (in ϵ) contribution of large jumps ranges from 36.51% to 24.70% for α ranging from 4 to 10 with an average of 30.12%, while the average contribution (in ϵ) of small jumps ranges from 7.54% to 13.54% with an average of 10.50%. Collectively, these results highlight that the crude oil and natural gas returns series exhibit large jumps and small jumps, and both types of jumps make significant contributions to total variability.

Figure 3.4 shows a plot of the overall jump contribution to total return variation during the 2006-2014 period in Figure 3.4. We see that in 2006 the total realized variance attributable to jumps for crude oil was 24.11% whereas for natural gas it was 63.25%. There was a dramatic increase in the total jump component contribution for crude oil to 50.69% in 2009, a 110.24% increase from 2006. The natural gas total variation attributable to jumps declines from 2006 to 2009 by 13.39% to 54.78%. The maximum 5-second natural gas log-return is obtained in 2006 whereas the minimum is calculated in 2014 where we estimate a total jump realized variance of 65.4%. We conclude that the extrema log-return estimates do in fact occur when the total realized variance attributable to jumps is the highest for each respective energy futures' series. Another interesting

observation that points to the validity of our estimates is that the largest skewness estimates for the log-return distributions correspond to the years with the largest contribution of jumps to total return variation. Table 3.1 reports that when the crude oil obtains its min/max values and has an overall jump contribution of 50.69%, the skewness of the log-return distribution is estimated at 37.4. The natural gas log-return distribution has an estimated skewness of 40.2 in 2006 and -87.6 in 2014 which corresponds to the highest total jump contribution of 63.25% and 65.4%, respectively.

[Insert Figure 3.4 about here]

Figures 3.5 and 3.6 present plots by year of the finite activity jump contribution and infinite activity jump contribution to realized variance for each log-return series. In 2006 the large finite jumps contribute 11.7% for crude oil total return variation and 12.4% is attributed to infinite activity jumps. The natural gas finite jump component contributes 39.77% in 20006 whereas small infinite activity jumps supplies 23.48% of the realized variance. However, moving ahead to 2009 we observe that the finite jump component from crude oil increases by 266.84% to 42.92% while the infinite activity contribution decreases by 37.39% to 7.77%. During our sample period the extreme log-returns for crude oil are primarily captured by a dramatic increase in finite activity jumps in 2009. Similarly, the natural gas finite activity jump contribution increases from 2006 to 2009 by 15.21% to 45.82% whereas the infinite jump activity drops by 61.84% to 8.96%. We also observe for the natural gas log-returns that that the large jump component increase from 2009 to 2014 by 30.58% to 59.83% while the small infinite jumps contribution to realized variance declines by 37.83% to 5.57%. Hence, again the most extreme 5-second

natural gas log-returns occur when the finite jump activity is largest and the infinite activity is least pronounced.

[Insert Figure 3.5 about here]

[Insert Figure 3.6 about here]

VIII. Conclusions

Employing a model-free approach developed by Ait-Sahalia and Jacod (2009, 2010, 2011), we find that both high frequency oil and natural gas futures returns series exhibit jumps. Further we find that both series display infinite, small, jump activity, as well less frequent large jump activity, as well a continuous variation component. Approximately 36 percent of the total return variation of crude oil futures returns can be attributed to jumps and this proportion is 41 percent for natural gas. In addition infrequent large jumps contribute almost three times as much to crude oil futures return variability as do small and infinite jumps, while a similar relative contribution is also observed for the natural gas sample. Since the infrequent large jumps are normally associated with economy wide shocks while the small jumps are commonly interpreted as originating from market microstructure factors, the finding that large jumps contribute more to return variation than small jumps suggests that crude oil and natural gas returns are sensitive to economy wide shocks.

The finding that crude oil and natural gas returns exhibit jumps has important implications for derivatives pricing and risk management. It is widely acknowledged that incorporating jumps into the risk neutral dynamics of the underlying is necessary to explain the high level of skewness implicit in the cross-section of equity index option

prices (Carr and Wu (2003)). By implication, although we leave this for future research, jumps in the underlying futures prices for the prices of options on futures for oil and natural gas can have important risk management consequences. In addition, the demand for funding liquidity will be higher for speculators who trade crude oil or natural gas futures with a mark-to-market account, if the commodity returns jump (Bakshi and Panayotov (2010)).

References

Aït-Sahalia, Y., & Jacod, J. (2009). Testing for jumps in a discretely observed process. The Annals of Statistics, 37(1), 184-222.

Aït-Sahalia, Y., & Jacod, J. (2010). Is Brownian motion necessary to model high-frequency data?. The Annals of Statistics, 3093-3128.

Aït-Sahalia, Y., & Jacod, J. (2011). Testing whether jumps have finite or infinite activity. The Annals of Statistics, 39(3), 1689-1719.

Aït-Sahalia, Y., & Jacod, J. (2012). Analyzing the Spectrum of Asset Returns: Jump and Volatility Components in High Frequency Data. Journal of Economic Literature, 50(4), 1007-1050.

Amihud, Y. (2002). Illiquidity and stock returns: cross-section and time-series effects. Journal of financial markets, 5(1), 31-56.

Askari, H., & Krichene, N. (2008). Oil Price Dynamics (2002–2006). Energy Economics, 30(5), 2134-2153.

Bakshi, G., Cao, C., & Chen, Z. (1997). Empirical performance of alternative option pricing models. The Journal of Finance, 52(5), 2003-2049.

Bakshi, G., & Panayotov, G. (2010). First-passage probability, jump models, and intra-horizon risk. Journal of Financial Economics, 95(1), 20-40.

Barberis, N., Huang, M., & Santos, T. (2001). PROSPECT THEORY AND ASSET PRICES. Quarterly Journal of Economics, 116(1).

Barndorff-Nielsen, O. E. (1997). Processes of normal inverse Gaussian type. Finance and stochastics, 2(1), 41-68.

Barndorff-Nielsen, O. E., & Shephard, N. (2004). Power and bipower variation with stochastic volatility and jumps. Journal of financial econometrics, 2(1), 1-37.

Barndorff-Nielsen, O. E., & Shephard, N. (2006). Econometrics of testing for jumps in financial economics using bipower variation. Journal of financial Econometrics, 4(1), 1-30.

Barro, R. J. (2006). Rare disasters and asset markets in the twentieth century. The Quarterly Journal of Economics, 823-866.

Battalio, R. H., Figlewski, S., & Neal, R. (2014). Exercise Boundary Violations in American-Style Options: The Rule, not the Exception. Working Paper.

Bates, D. S. (1996). Jumps and stochastic volatility: Exchange rate processes implicit in Deutsche Mark options. Review of financial studies, 9(1), 69-107.

Bates, D. S. (2000). Post-'87 crash fears in the S&P 500 futures option market. Journal of Econometrics, 94(1-2), 181-238.

Bates, D. S. (2003). Empirical option pricing: A retrospection. Journal of Econometrics, 116(1), 387-404.

Bewley, T. F. (2002). Knightian decision theory. Part I. Decisions in economics and finance, 25(2), 79-110.

Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. The journal of political economy, 637-654.

Black, Fisher, (1976). Studies in stock price volatility changes, Proceedings of the 1976 Business Meeting of the Business and Economics Statistics Section, American Statistical Association, pp. 177–181.

Bloom, N. (2009). The impact of uncertainty shocks. Econometrica, 77(3), 623-685.

Brunnermeier, M. K., & Pedersen, L. H. (2009). Market liquidity and funding liquidity. Review of Financial studies, 22(6), 2201-2238.

Broadie, M., Detemple, J., Ghysels, E., & Torrés, O. (2000). American options with stochastic dividends and volatility: A nonparametric investigation. Journal of Econometrics, 94(1-2), 53-92.

Buraschi, A., & Jiltsov, A. (2006). Model uncertainty and option markets with heterogeneous beliefs. The Journal of Finance, 61(6), 2841-2897.

Campbell, J. Y., & Cochrane, J. (1999). Appendix for" By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior. Journal of Political Economy.

Carr, P., Geman, H., Madan, D. B., & Yor, M. (2002). The fine structure of asset returns: An empirical investigation. The Journal of Business, 75(2), 305-333.

Carr, P., Geman, H., Madan, D. B., & Yor, M. (2003). Stochastic volatility for Lévy processes. Mathematical Finance, 13(3), 345-382.

Carr, P., Geman, H., Madan, D. B., & Yor, M. (2007). SELF-DECOMPOSABILITY AND OPTION PRICING. Mathematical finance, 17(1), 31-57.

Carr, P., Jarrow, R., & Myneni, R. (1992). Alternative characterizations of American put options. Mathematical Finance, 2(2), 87-106.

Carr, P., & Madan, D. (1999). Option valuation using the fast Fourier transform. Journal of computational finance, 2(4), 61-73.

Carr, P., & Wu, L. (2003). The finite moment log stable process and option pricing. The Journal of Finance, 58(2), 753-778.

Carr, P., & Wu, L. (2004). Time-changed Lévy processes and option pricing. Journal of Financial Economics, 71(1), 113-141.

Carr, P., & Wu, L. (2007). Stochastic skew in currency options. Journal of Financial Economics, 86(1), 213-247.

Chen, H., Joslin, S., & Ni, S. (2014, October). Demand for crash insurance, intermediary constraints, and stock return predictability. In AFA 2013 San Diego Meetings Paper.

Chordia, T., Roll, R., & Subrahmanyam, A. (2000). Commonality in liquidity. Journal of Financial Economics, 56(1), 3-28.

Chordia, T., Sarkar, A., & Subrahmanyam, A. (2005). An empirical analysis of stock and bond market liquidity. Review of Financial Studies, 18(1), 85-129.

Christoffersen, P., & Jacobs, K. (2004). The importance of the loss function in option valuation. Journal of Financial Economics, 72(2), 291-318.

Chung, K. H., & Chuwonganant, C. (2014). Uncertainty, market structure, and liquidity. Journal of Financial Economics, 113(3), 476-499.

Duffie, D., Pan, J., & Singleton, K. (2000). Transform analysis and asset pricing for affine jump-diffusions. Econometrica, 68(6), 1343-1376.

Easley, D., & O'Hara, M. (2010). Liquidity and valuation in an uncertain world. Journal of Financial Economics, 97(1), 1-11.

Easley, D., & O'hara, M. (2010). Microstructure and ambiguity. The Journal of Finance, 65(5), 1817-1846.

Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. The quarterly journal of economics, 643-669.

Epstein, L. G., & Schneider, M. (2007). Learning under ambiguity. The Review of Economic Studies, 74(4), 1275-1303.

Epstein, L. G., & Schneider, M. (2008). Ambiguity, information quality, and asset pricing. The Journal of Finance, 63(1), 197-228.

Eraker, B., Johannes, M., & Polson, N. (2003). The impact of jumps in volatility and returns. The Journal of Finance, 58(3), 1269-1300.

Eraker, B. (2004). Do stock prices and volatility jump? Reconciling evidence from spot and option prices. The Journal of Finance, 59(3), 1367-1404.

Fang, F., & Oosterlee, C. W. (2009). Pricing early-exercise and discrete barrier options by Fourier-cosine series expansions. Numerische Mathematik, 114(1), 27-62.

Fleming, J., & Whaley, R. E. (1994). The value of wildcard options. The Journal of Finance, 49(1), 215-236.

Foresi, S., & Wu, L. (2005). Crash-o-phobia: A domestic fear or a worldwide concern? The Journal of Derivatives, 13(2), 8-21.

Gilboa, I., & Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. Journal of mathematical economics, 18(2), 141-153.

Gorton, G. B., Hayashi, F., & Rouwenhorst, K. G. (2013). The fundamentals of commodity futures returns. Review of Finance, 17(1), 35-105.

Goyenko, R. Y., Holden, C. W., & Trzcinka, C. A. (2009). Do liquidity measures measure liquidity? Journal of financial Economics, 92(2), 153-181.

Greenspan, A. (2004). Risk and uncertainty in monetary policy. The American Economic Review, 94(2), 33-40.

Gronwald, M. (2012). A Characterization of Oil Price Behavior: Evidence from Jump Models. *Energy Economics*, vol 34, no. 5, pp. 1310-1317.

Guiso, L., Sapienza, P., & Zingales, L. (2013). Time varying risk aversion (No. w19284). National Bureau of Economic Research.

Gukhal, C. R. (2001). Analytical Valuation of American Options on Jump-Diffusion Processes. Mathematical Finance, 11(1), 97-115.

Hameed, A., Kang, W., & Viswanathan, S. (2010). Stock market declines and liquidity. The Journal of Finance, 65(1), 257-293.

Harvey, C. R., & Whaley, R. E. (1992a). Market volatility prediction and the efficiency of the S & P 100 index option market. Journal of Financial Economics, 31(1), 43-73.

Harvey, C. R., & Whaley, R. E. (1992b). Dividends and S&P 100 index option valuation. Journal of Futures Markets, 12(2), 123-137.

He, Z., & Krishnamurthy, A. (2013). Intermediary Asset Pricing. American Economic Review, 103(2), 732-770.

Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. Review of financial studies, 6(2), 327-343.

Hirsa, A., & Madan, D. B. (2004). Pricing American options under variance gamma. Journal of Computational Finance, 7(2), 63-80.

Ho, T., & Stoll, H. R. (1981). Optimal dealer pricing under transactions and return uncertainty. Journal of Financial economics, 9(1), 47-73.

Huang, X., & Tauchen, G. (2005). The relative contribution of jumps to total price variance. Journal of financial econometrics, 3(4), 456-499.

Huang, J. Z., & Wu, L. (2004). Specification analysis of option pricing models based on time-changed Lévy processes. The Journal of Finance, 59(3), 1405-1440.

Jorion, P., & Stoughton, N. M. (1989). An empirical investigation of the early exercise premium of foreign currency options. Journal of Futures Markets, 9(5), 365-375.

Ju, N., & Miao, J. (2012). Ambiguity, learning, and asset returns.Econometrica, 80(2), 559-591.

Klibanoff, P., Marinacci, M., & Mukerji, S. (2005). A smooth model of decision making under ambiguity. Econometrica, 73(6), 1849-1892.

Klibanoff, P., Marinacci, M., & Mukerji, S. (2009). Recursive smooth ambiguity preferences. Journal of Economic Theory, 144(3), 930-976.

Kou, S. G. (2002). A jump-diffusion model for option pricing. Management science, 48(8), 1086-1101.

Kyle, A. S. (1985). Continuous auctions and insider trading. Econometrica: Journal of the Econometric Society, 1315-1335.

Lee, S. S., & Hannig, J. (2010). Detecting jumps from Lévy jump diffusion processes. Journal of Financial Economics, 96(2), 271-290.

Lee, Y., Hu, H., Chiou, J. (2010). Jump Dynamics with Structural Breaks for Crude Oil Prices. Energy Economics 32, 343-350.

Li, H., Wells, M. T., & Cindy, L. Y. (2008). A Bayesian analysis of return dynamics with Lévy jumps. Review of Financial Studies, 21(5), 2345-2378.

Longstaff, F. A., & Schwartz, E. S. (2001). Valuing American options by simulation: A simple least-squares approach. Review of Financial studies, 14(1), 113-147.

Maccheroni, F., Marinacci, M., & Ruffino, D. (2013). Alpha as Ambiguity: Robust Mean-Variance Portfolio Analysis. Econometrica, 81(3), 1075-1113.

Madan, D. B., & Seneta, E. (1990). The variance gamma (VG) model for share market returns. Journal of business, 511-524.

Madan, D. B., Carr, P. P., & Chang, E. C. (1998). The variance gamma process and option pricing. European finance review, 2(1), 79-105.

Medvedev, A., & Scaillet, O. (2010). Pricing American options under stochastic volatility and stochastic interest rates. Journal of Financial Economics, 98(1), 145-159.

Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. Journal of financial economics, 3(1), 125-144.

Næs, R., Skjeltorp, J. A., & Ødegaard, B. A. (2011). Stock market liquidity and the business cycle. The Journal of Finance, 66(1), 139-176.

Nomikos, N., & Andriosopoulos, K. (2012). Modelling energy spot prices: Empirical evidence from NYMEX. Energy Economics, 34(4), 1153-1169.

Ozsoylev, H., & Werner, J. (2011). Liquidity and asset prices in rational expectations equilibrium with ambiguous information. Economic Theory,48(2-3), 469-491.

Pan, J. (2002). The jump-risk premia implicit in options: Evidence from an integrated time-series study. Journal of financial economics, 63(1), 3-50.

Routledge, B. R., & Zin, S. E. (2009). Model uncertainty and liquidity. Review of Economic dynamics, 12(4), 543-566.

Ruijter, M. J., & Oosterlee, C. W. (2012). Two-dimensional Fourier cosine series expansion method for pricing financial options. SIAM Journal on Scientific Computing, 34(5), B642-B671.

Savage, L. J. (1954). The foundations of statistics.

Schwartz, E. S. (1997). The stochastic behavior of commodity prices: Implications for valuation and hedging. The Journal of Finance, 52(3), 923-973.

Shiller, R. J. (1981). The use of volatility measures in assessing market efficiency. The Journal of Finance, 36(2), 291-304.

Stock, J. H., & Watson, M. W. (2001). Vector autoregressions. The Journal of Economic Perspectives, 15(4), 101-115.

Trolle, A. B., & Schwartz, E. S. (2009). Unspanned stochastic volatility and the pricing of commodity derivatives. Review of Financial Studies, hhp036.

Vayanos, D., & Wang, J. (2012). Market Liquidity--Theory and Empirical Evidence (No. w18251). National Bureau of Economic Research.

Vives, X. (1995). Short-term investment and the informational efficiency of the market. Review of Financial Studies, 8(1), 125-160.

Appendix: Figures and Tables

Figure 1.1: The Effect of Ambiguity on Price Impact

In this figure, I present a numerical example based on my trading model with learning under ambiguity. Specifically, I plot three time series for λ_t defined by equation (13) for the level of ambiguity σ_A^2 taking values of 2, 4, and 6. In all three time series, the payoff variance $\sigma_0^2 = 2$, and the signal noise $\sigma_s^2 = 1$. In addition, I choose the risk aversion parameter $\rho = 2$ and the ambiguity aversion parameter $\theta = 8$ based on Ju and Miao (2012).

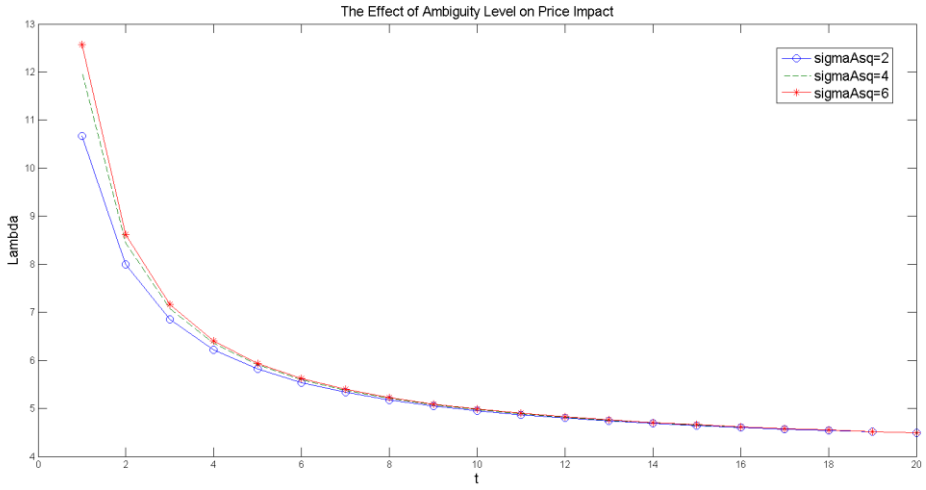

Figure 1.2: The Effects of Risk Aversion and Ambiguity Aversion on Price Impact

In this figure, I present two numerical examples based on my trading model with learning under ambiguity. In each panels, I plot three time series for λ_t defined by equation (13) for different parameter of interest. In both panels, the payoff variance $\sigma_0^2 = 2$, the level of ambiguity $\sigma_A^2 = 4$, and the signal noise $\sigma_s^2 = 1$. In the upper panel, I let the risk aversion coefficient ρ take values of 2, 3, and 4, while fixing the ambiguity aversion parameter $\theta = 8$. In the lower panel, I let the ambiguity aversion coefficient θ take values of 7, 9, and 11, while fixing the risk aversion coefficient is $\rho = 2$.

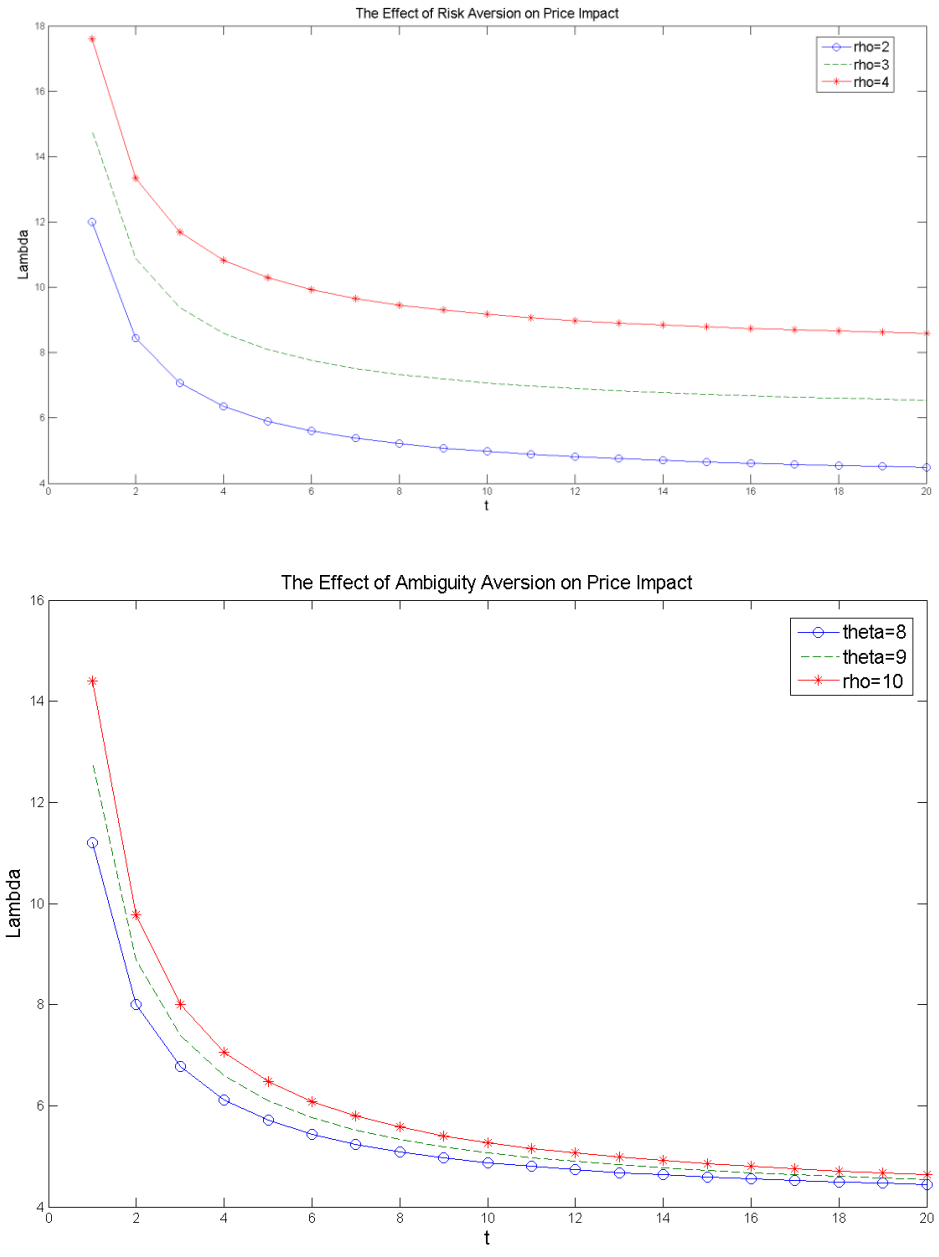

Figure 1.3: Monthly U.S. Stock Market Volatility

The VXO index is used for period after 1986, when the index becomes available. Monthly standard deviation of the daily S&P 500 index is used to construct the return volatility before 1986. The 18 dotted vertical lines refer to the volatility shocks identified following Bloom (2009) and are described in Table 1.1. A month is labeled as a shock in volatility if the HP detrended ($\lambda = 129600$) volatility in that month exceeds 1.96 times the standard deviation of the entire HP detrended volatility series.

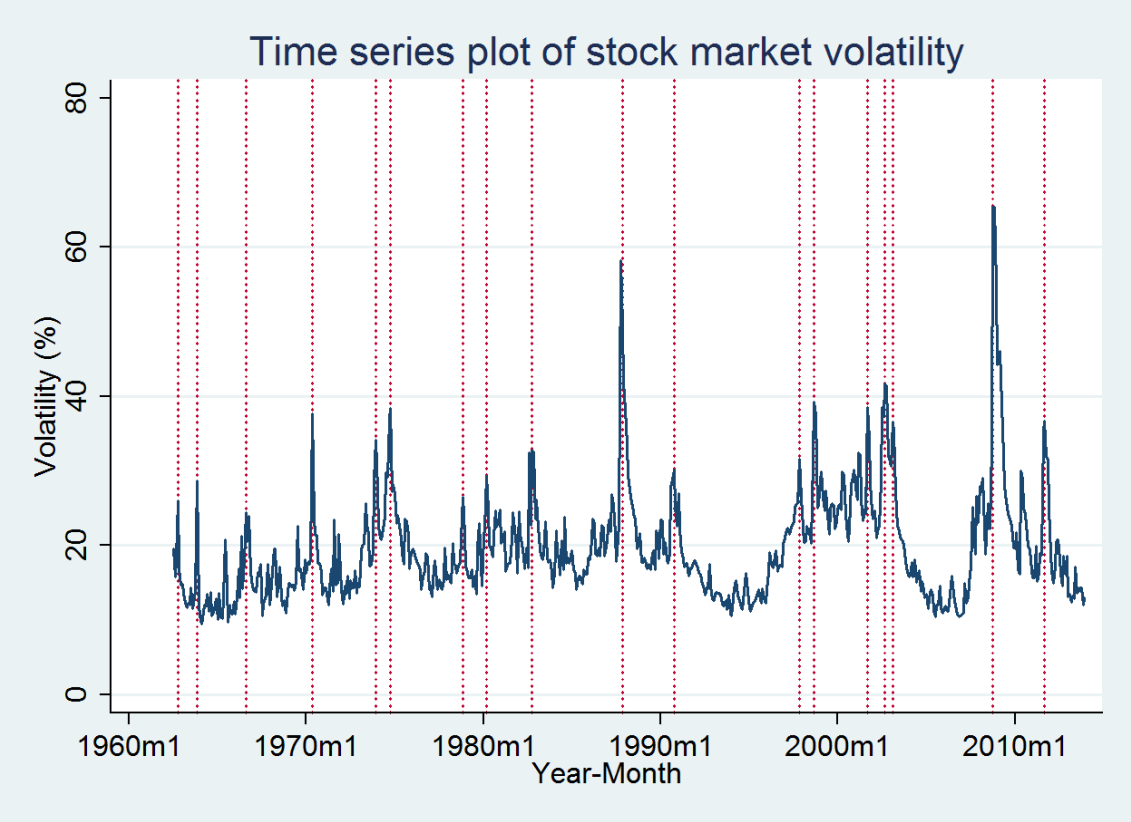

Figure 1.4: Stock Market Illiquidity

This figure shows the monthly time series plot of stock market illiquidity. I focus on the price impact aspect of market illiquidity. I select stocks traded on the NYSE to maintain homogeneity. For each stock, I use daily returns and volumes to construct the Amihud (2002) measure at a monthly frequency. The monthly market wide price impact measure is hence obtained by averaging the measure across individual stocks. The solid line is the original series, and the dashed line is the HP detrended ($\lambda = 129600$) series. The 18 dotted vertical lines refer to the volatility shocks identified following Bloom (2009) and are described in Table 1.1.

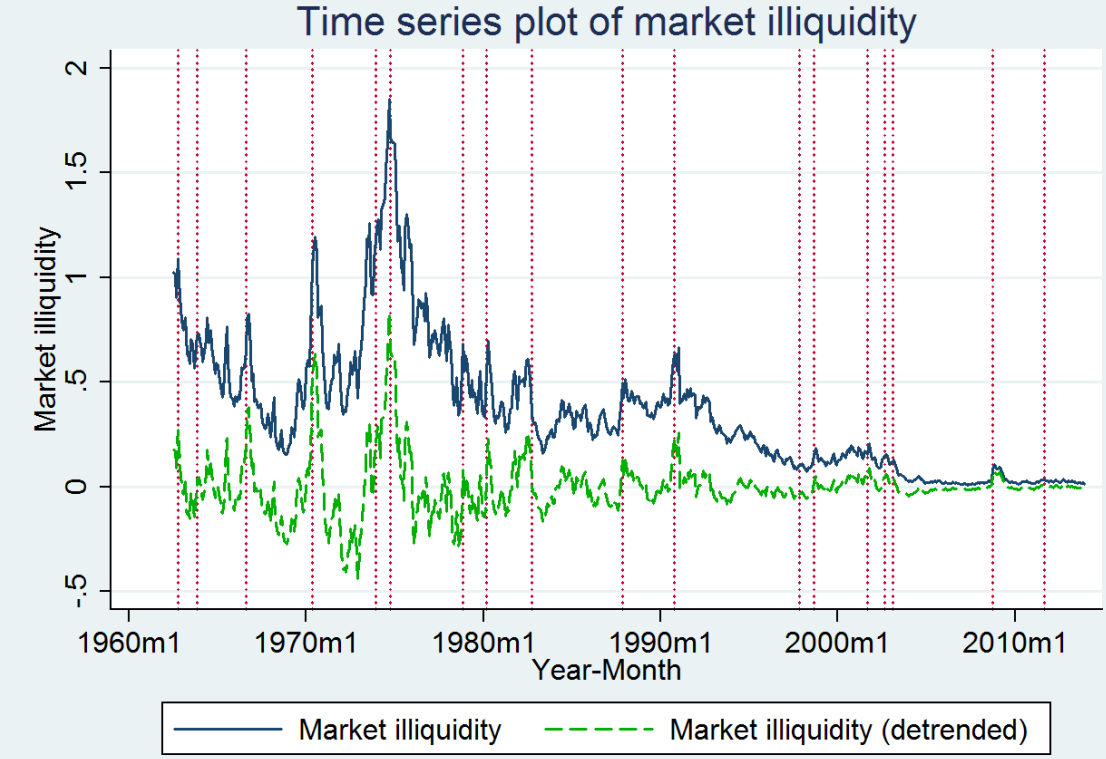

Figure 1.5: Response of Market Liquidity to Volatility Shocks

In this figure, I plot the response of market liquidity to shocks from volatility, from estimating the recursive VAR: (log(stock market return), stock market volatility shocks indicator, Federal Funds Rates, market liquidity). All variables are detrended. Dashed lines are 95% confidence interval bands around the response based on bootstrapped standard errors (100 runs). The four time series are at the monthly frequency starting from 1962.10 till 2013.12.

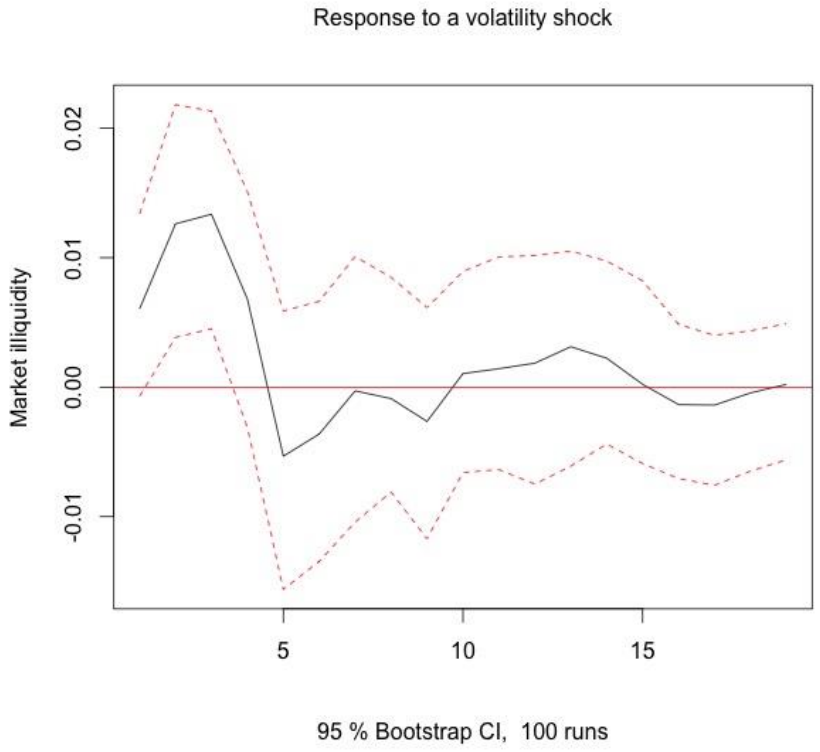

Figure 1.6: Response of Market Liquidity to Volatility Shocks for 1962--1989 and 1990--2013

In this figure, I split the sample period into 1962--1989, and 1990--2013. The upper figure is for the first periods and shows the response of market liquidity to shocks from volatility from estimating the recursive VAR: (log(stock market return), stock market volatility shocks indicator, Federal Funds Rates, market liquidity). All variables are detrended. Dashed lines are 95% confidence interval bands around the response based on bootstrapped standard errors (100 runs). The four time series are at the monthly frequency starting from 1962.10 till 2013.12.

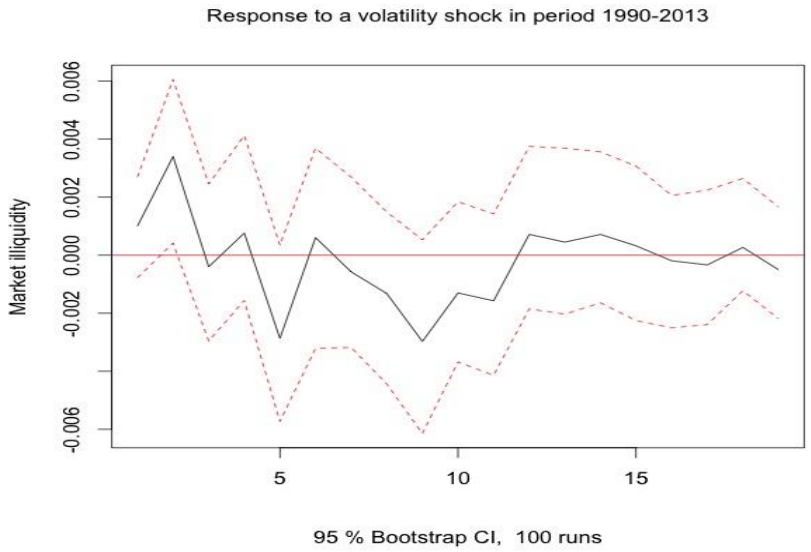

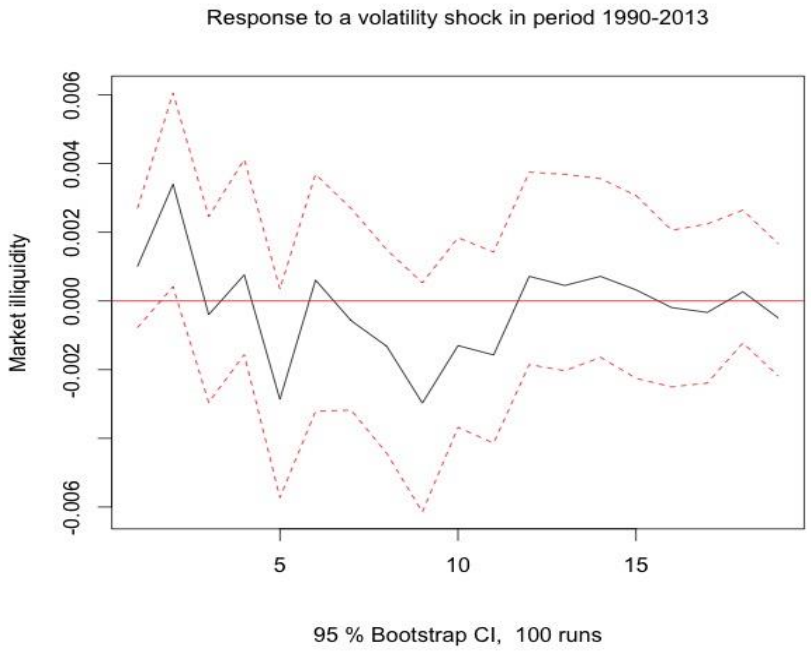

Figure 1.7: Response of Market Liquidity to Changes in Stock Price Level and FFR

In this figure, I plot the response of market liquidity to changes in stock price leve and FFR from estimating the recursive VAR: (log(stock market return), stock market volatility shocks indicator, Federal Funds Rates, market liquidity). Dashed lines are 95% confidence interval bands around the response based on bootstrapped standard errors (100 runs). The four time series are at the monthly frequency starting from 1962.10 till 2013.12.

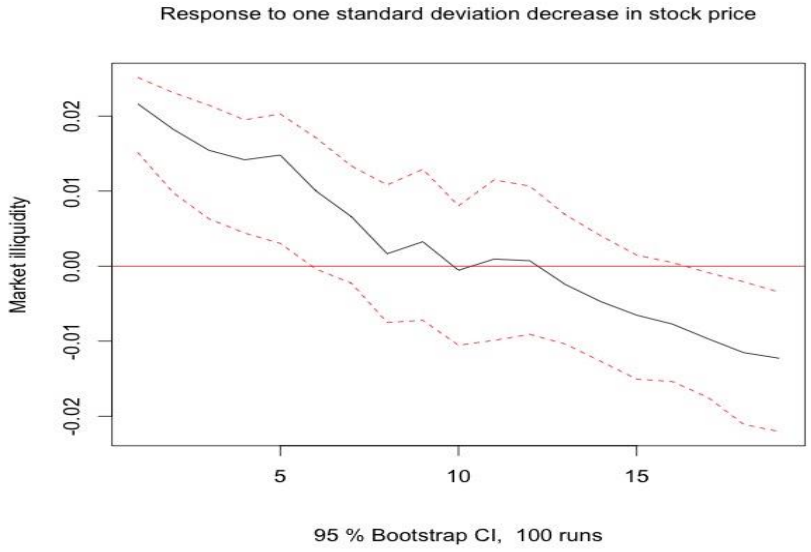

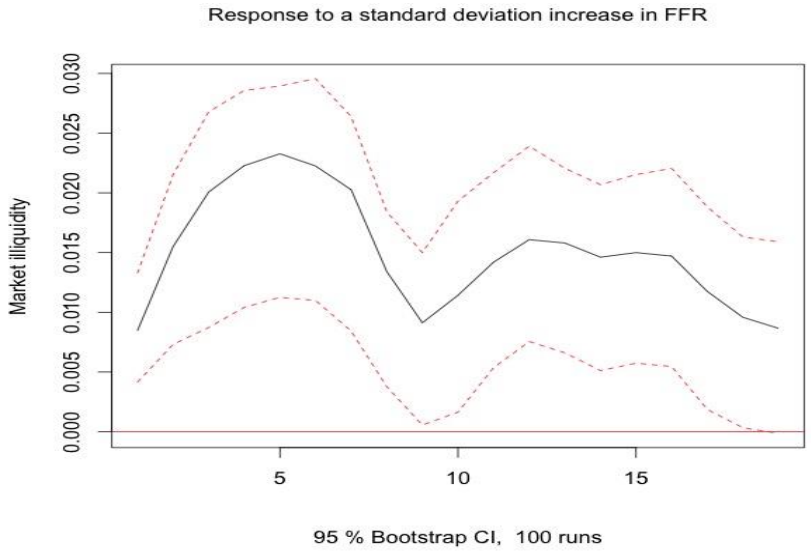

Figure 1.8: Response of Market Liquidity to Changes in Stock Price Level and FFR for 1962--1989 and 1990—2013

In this figure, I plot the response of market liquidity to shocks from volatility, stock return, and FFR from estimating the recursive VAR: (log(stock market return), stock market volatility shocks indicator, Federal Funds Rates, market liquidity). Dashed lines are 95% confidence interval bands around the response based on bootstrapped standard errors (100 runs). The four time series are at the monthly frequency starting from 1962.10 till 2013.12.

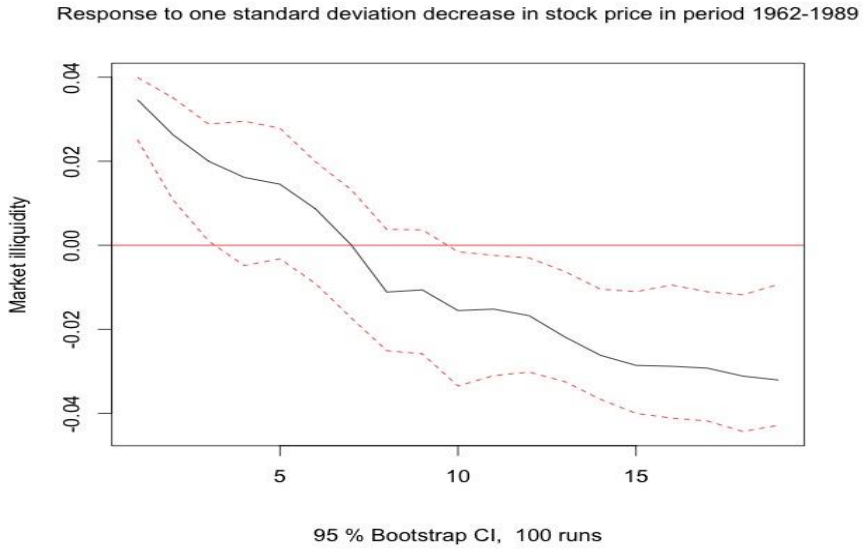

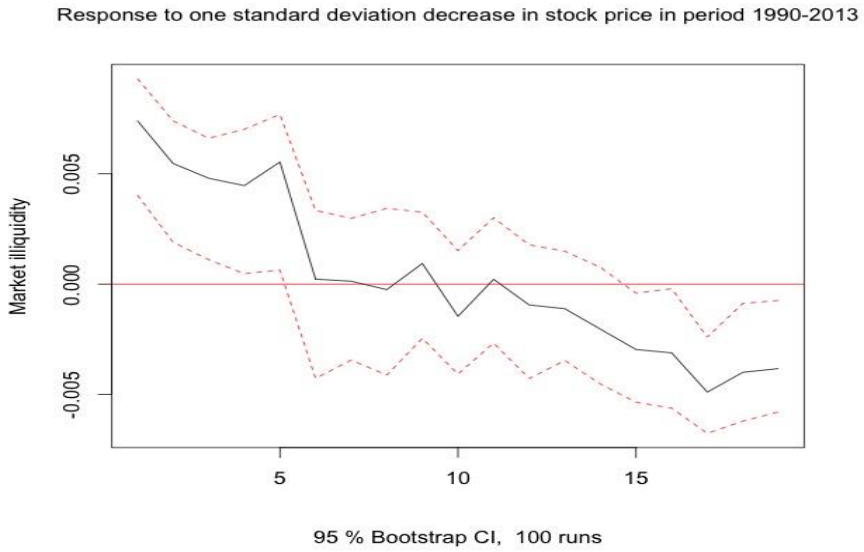

Figure 1.8 (*continued*)

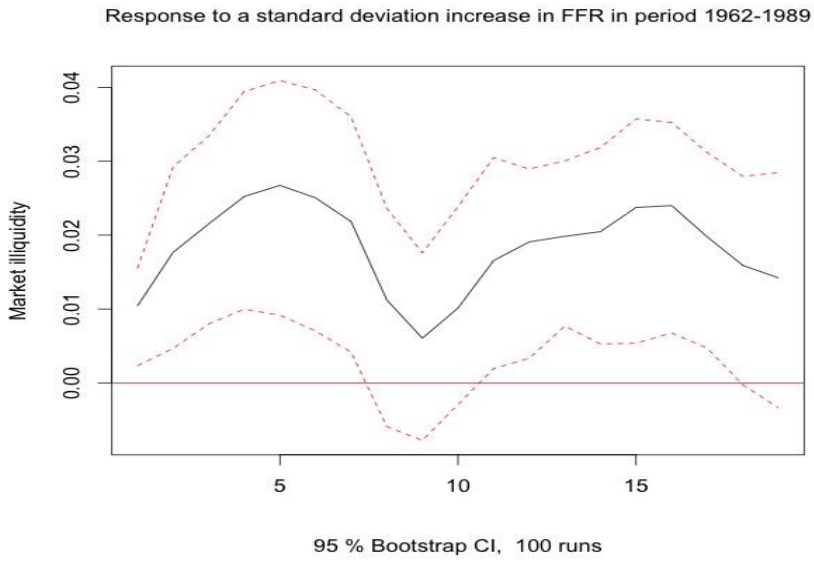

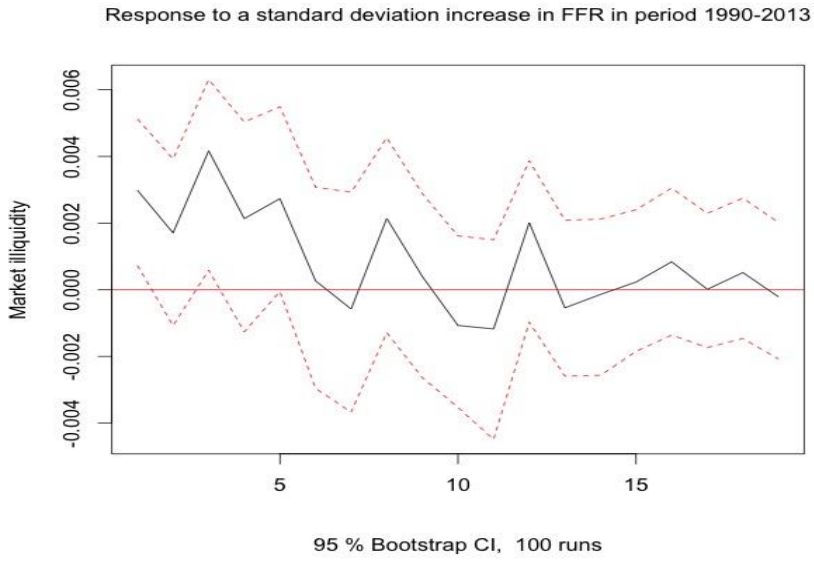

Figure 1.9: VAR Estimation of Volatility Shocks Impacts on Market Liquidity (Alternative Measures for Uncertainty Shocks)

In this figure, I present the response of market liquidity to shocks from volatility based on alternative measures for uncertainty shocks. In the first figure of the upper panel, I use the 0/1 indicator but each indicator is scaled by the actual volatility. In the second figure of the upper panel, I identify shocks using the first month instead of the month with the local maximum volatility. In the first figure of the lower panel, I exclude economic shocks. The VAR orderings are the same as the benchmark result in Figure 3. Dashed lines are 95% confidence interval bands around the response based on bootstrapped standard errors (100 runs). The four time series are at the monthly frequency starting from 1962.10 till 2013.12.

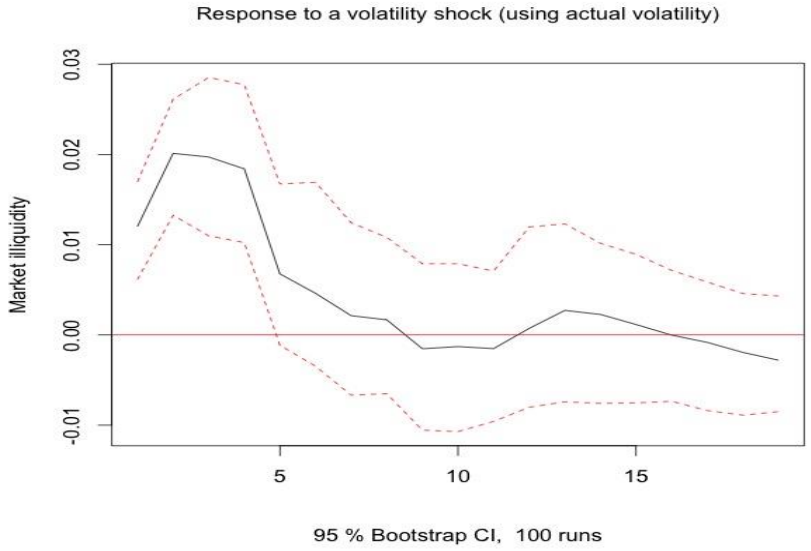

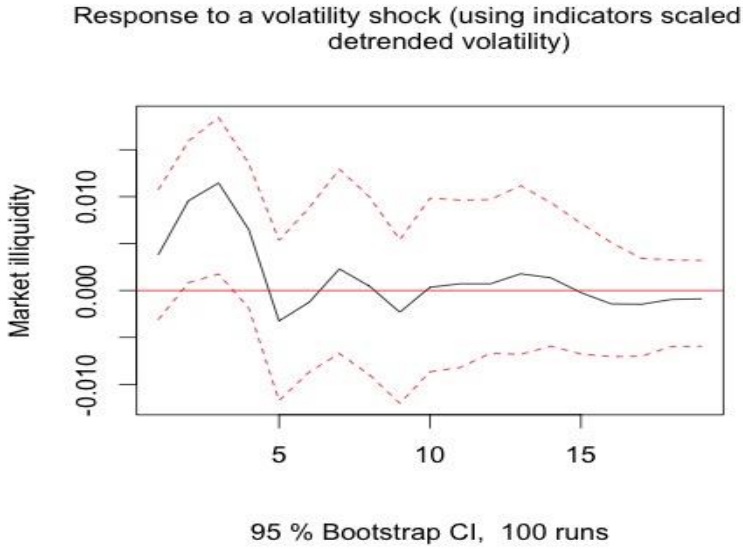

Figure 1.9 (*continued*)

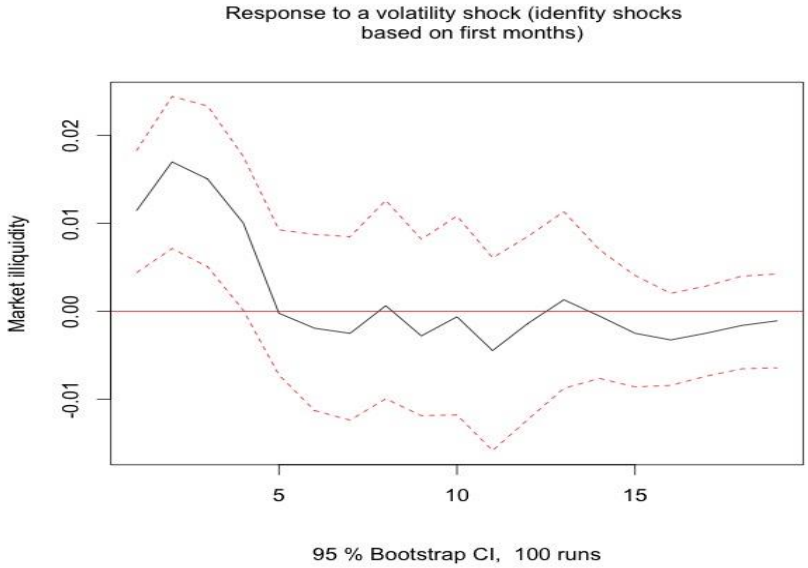

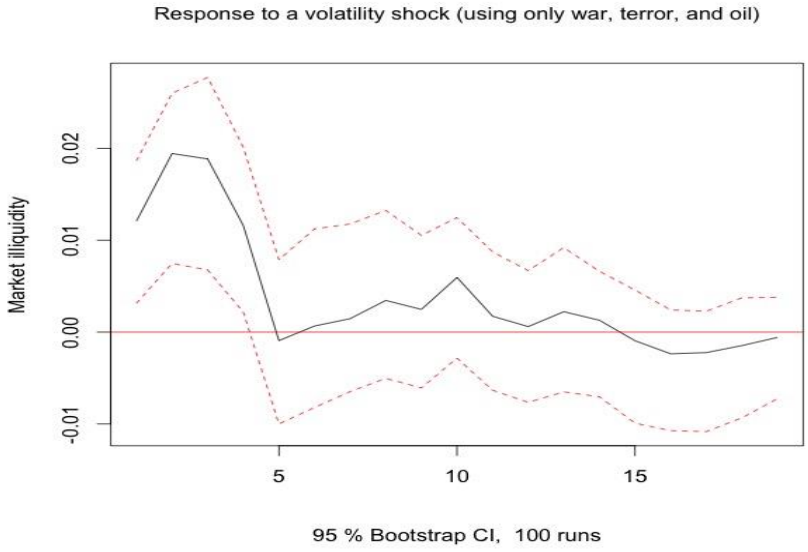

Figure 1.10: VAR Estimation of Volatility Shocks Impacts on Market Liquidity (Alternative Orderings)

In this figure, I present the response of market liquidity to shocks from volatility based on alternative VAR orderings. In the first figure of the upper panel, the VAR ordering is (stock market volatility shocks indicator, log(stock market return), Federal Funds Rates, market liquidity). In the second figure of the upper panel, the ordering is (stock market volatility shocks indicator, market liquidity). In the first figure of the lower panel, the ordering is (market liquidity, FFR, log(stock market return), stock market volatility shocks indicator). Dashed lines are 95% confidence interval bands around the response based on bootstrapped standard errors (100 runs). The four time series are at the monthly frequency starting from 1962.10 till 2013.12.

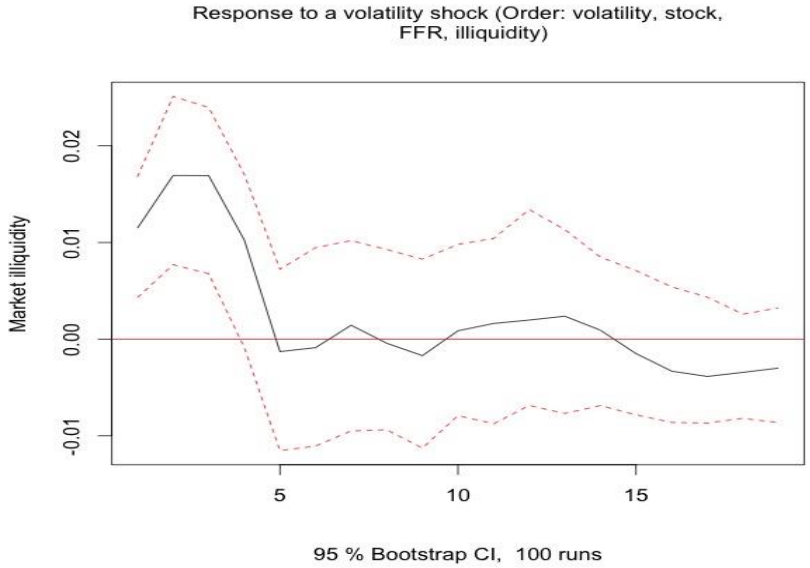

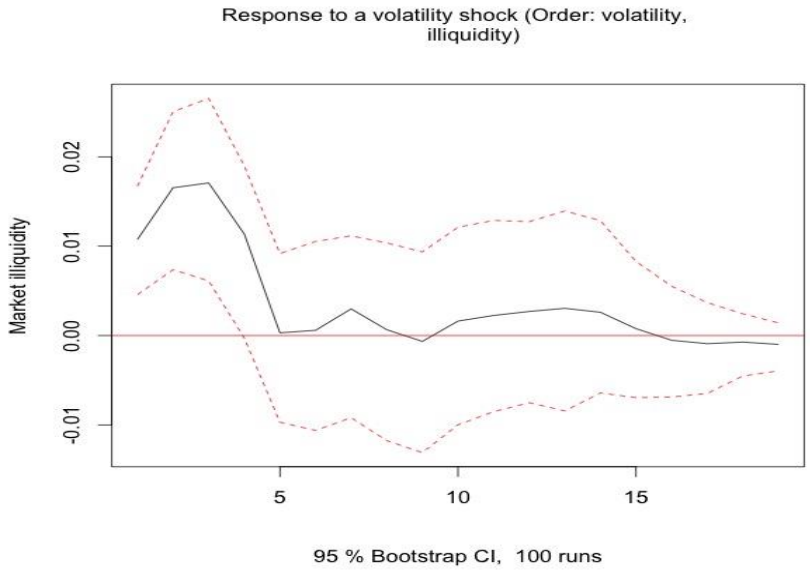

Figure 1.10 (*continued*)

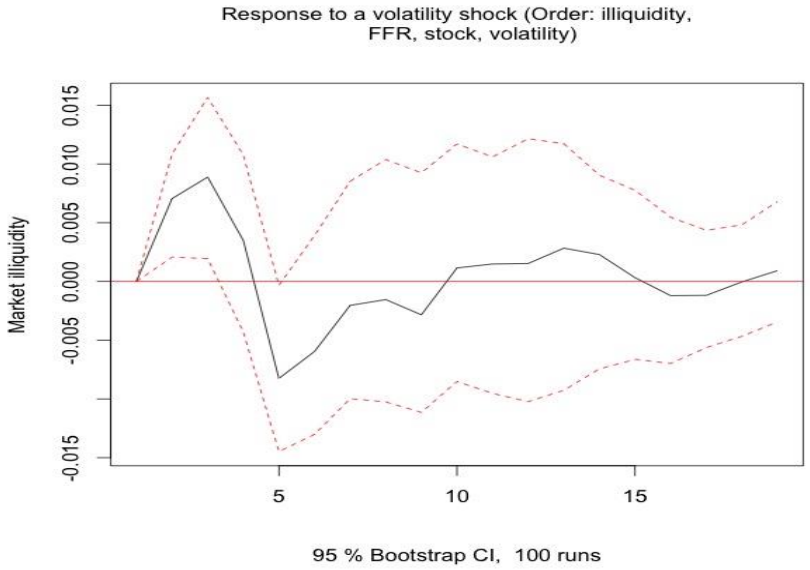

Table 1.1: Major Stock Market Volatility Shocks

This table shows the 18 identified shocks in the monthly stock market volatility series. The VXO index is used for the period after 1986, when the index becomes available. Monthly standard deviation of the daily S&P 500 index is used to construct the return volatility before 1986. A month is labeled as a shock in volatility if the HP detrended (lambda=129600) volatility in that month exceeds 1.96 times the standard deviation of the entire HP detrended volatility series.

Event	Max Volatility	First Volatility	Type
Cuban missile crisis	Oct-62	Oct-62	Terror
Assassination of JFK	Nov-63	Nov-63	Terror
Vietnam buildup	Aug-66	Aug-66	War
Cambodia and Kent State	May-70	May-70	War
OPEC I, Arab-Israeli War	Dec-73	Dec-73	Oil
Franklin National	Oct-74	Sep-74	Economic
OPEC II	Nov-78	Nov-78	Oil
Afghanistan, Iran hostages	Mar-80	Mar-80	War
Monetary cycle turning point	Oct-82	Aug-82	Economic
Black Monday	Nov-87	Oct-87	Economic
Gulf War I	Oct-90	Sep-90	War
Asian Crisis	Nov-97	Nov-97	Economic
Russian, LTCM default	Sep-98	Sep-98	Economic
9/11 terrorist attack	Sep-01	Sep-01	Terror
Worldcom and Enron	Sep-02	Jul-02	Economic
Gulf War II	Feb-03	Feb-03	War
Credit crunch	Oct-08	Aug-07	Economic
Sovereign debt crisis	Sep-11	Aug-11	Economic

Table 1.2: VAR Descriptive Statistics for (Stock, Volatility, FFR, Illiquidity)

The entries in Panel A show the p-values for pair-wise F-tests based on the null that variables in the cause column does not Granger cause effect variables. The results are computed from a VAR with 12 lags and a constant term over the 1962.07 to 2013.12 period. Variable "stock" refers to the detrended log price of S&P 500, "volatility" uses monthly VXO series after 1986 and stock market volatility before, "FFR" stands for Federal Funds rate, and "illiquidity" stands for market illiquidity based on the Amihud (2002) measure.

A. Granger-Causality Tests				
	Effect Variables			
Cause Variables	Stock	Volatility	FFR	Illiquidity
Stock	0	0.18	0.04	0.7
Volatility	0	0	0.13	0
FFR	0.22	0.53	0	0.01
Illiquidity	0	0.28	0.03	0
B. i. Variance Decomposition of Stock (Percentage Points)				
Forecast horizon	Stock	Volatility	FFR	Illiquidity
1	1	0	0	0
4	0.75	0.2	0.01	0.04
8	0.77	0.18	0.03	0.02
12	0.79	0.16	0.03	0.02
B. ii. Variance Decomposition of Volatility (Percentage Points)				
Forecast horizon	Stock	Volatility	FFR	Illiquidity
1	0.08	0.92	0	0
4	0.07	0.93	0	0
8	0.07	0.92	0	0.01
12	0.07	0.91	0	0.01
B. iii. Variance Decomposition of FFR (Percentage Points)				
Forecast horizon	Stock	Volatility	FFR	Illiquidity
1	0	0	1	0
4	0.01	0.01	0.97	0.01
8	0.04	0.03	0.89	0.04
12	0.07	0.04	0.82	0.07
B. iv. Variance Decomposition of Illiquidity (Percentage Points)				
Forecast horizon	Stock	Volatility	FFR	Illiquidity
1	0.1	0.04	0.01	0.85
4	0.08	0.11	0.11	0.7
8	0.09	0.09	0.19	0.63
12	0.09	0.09	0.22	0.6

Figure 2.1: Moneyness and Maturity Distributions

This figure shows the frequency distributions for moneyness (K/S0) and maturity (T, days), for both our matched put and call samples for period 2005.05.02--2007.12.31. The underlying asset is the S&P 100 index. We pair an OEX option with an XEO option with the same strike and maturity, if their quotes are posted within 60 seconds. We show the results for the put sample in the upper two panels and that for the call sample in the lower two panels. We show the results for moneyness in the first column and that for maturity in the second column.

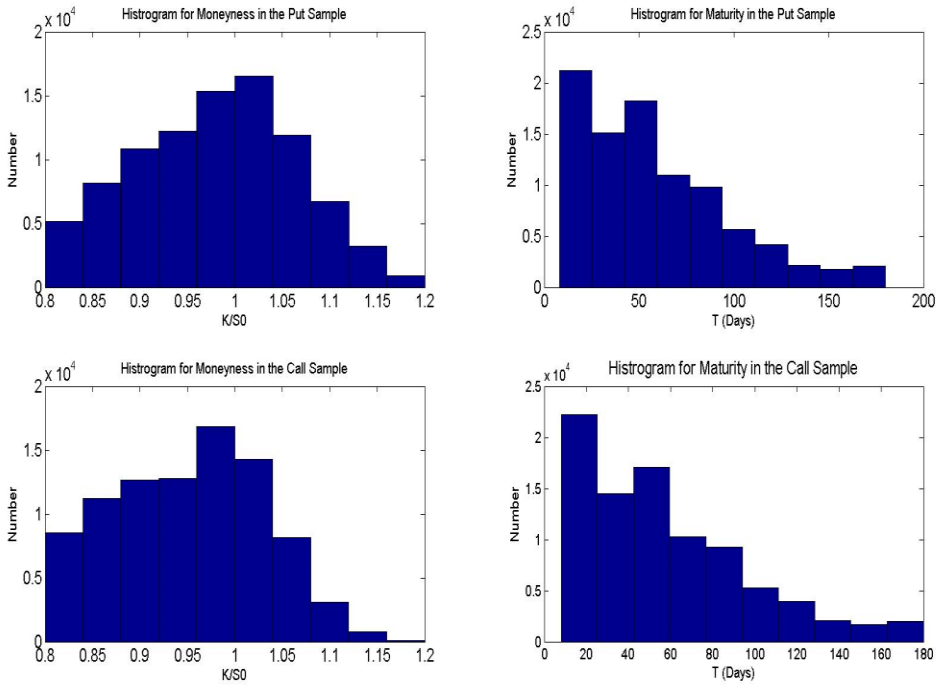

Figure 2.2: Market EEP Plotted Against Moneyness and Maturity

This figure plots EEP/S0 (%) against moneyness (K/S0, first three pictures) and maturity (T, days, last three pictures) for our matched put sample for period 2005.05.02--2007.12.31. The underlying asset is the S&P 100 index. We pair an OEX put with an XEO put with the same strike and maturity, if their quotes are posted within 60 seconds.

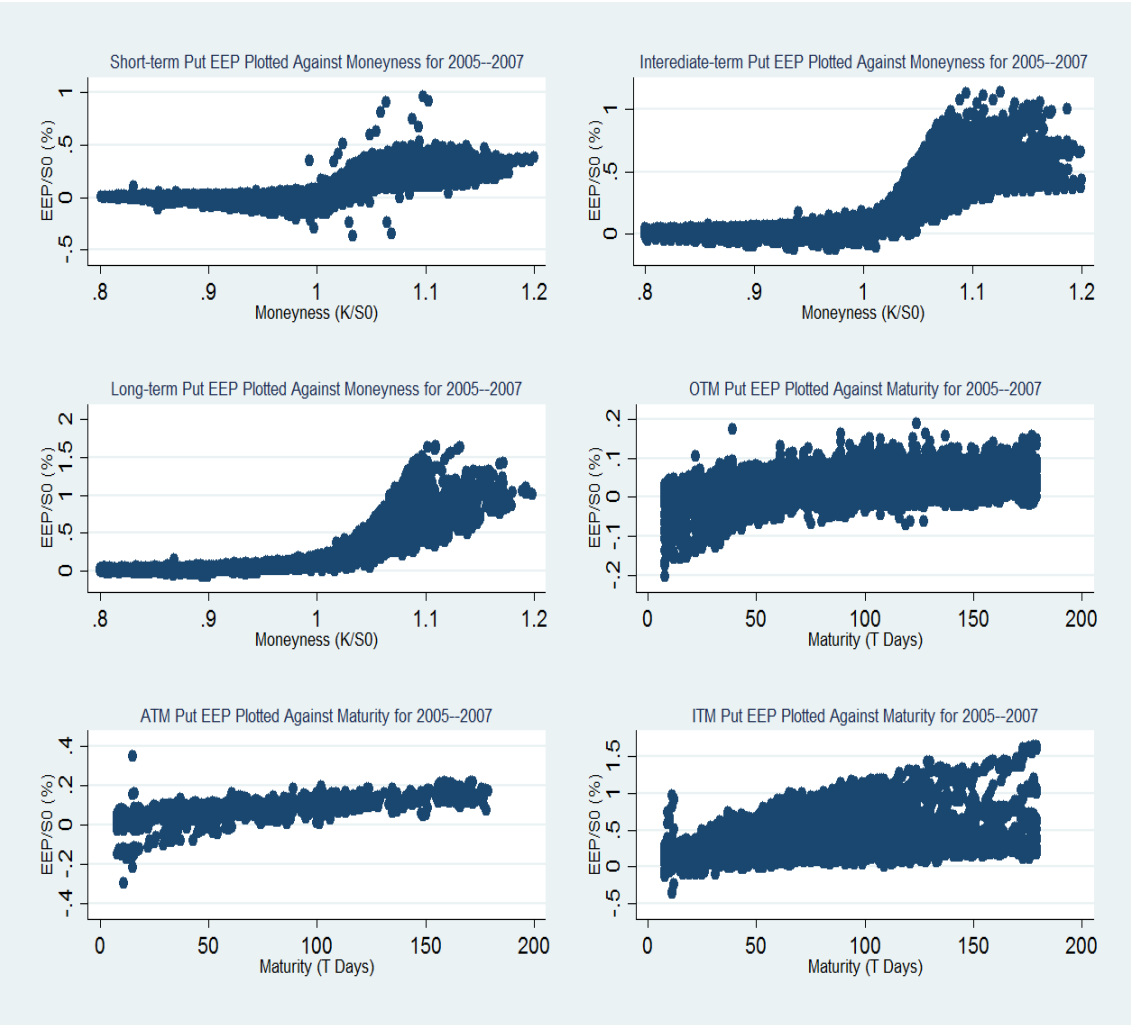

Figure 2.3: Time Value minus Spread for OEX Puts for 2005--2007

This figure shows the relations between time value minus the bid-ask spread and moneyness (left) and maturity (right) for OEX puts for 2005.05—2007.12. Time value is the difference between the market price of an OEX put and its intrinsic value. The vertical axes in both figures are in dollars.

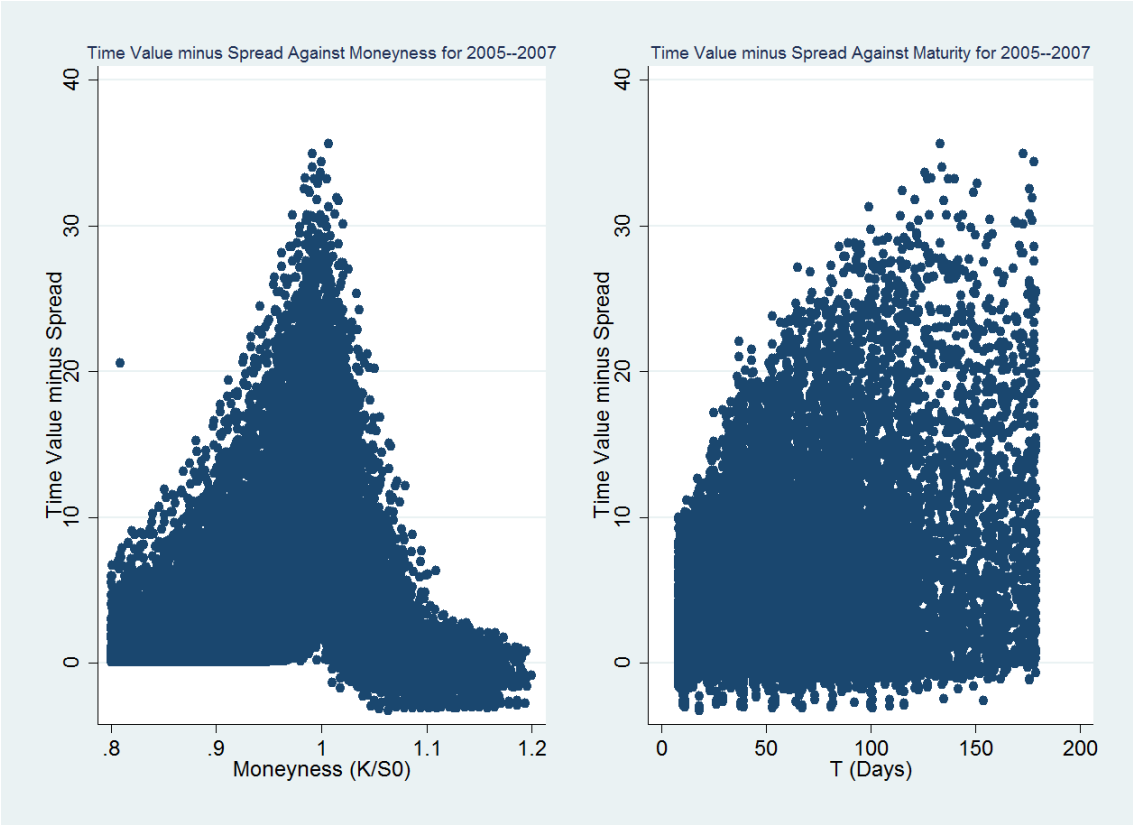

Table 2.1: Descriptive Statistics

This table summarizes our matched put sample for 2005.05.02 to 2007.12.31. The underlying asset is the S&P 100 index. We pair an OEX put with an XEO put with the same strike and maturity, if their quotes are posted within 60 seconds. In Panel A, we report the summary for the overall sample. Here the prices are scaled by the underlying level and expressed in basis points. We also report the summary of moneyness (K/S0), and maturity (days). "S.D." refers to standard deviation. "I.Q.R." refers to interquartile range. In Panel B, we report group-based average EEP, scaled by European option prices and the underlying levels respectively. Here, OTM refers to K/S0<0.99, ATM refers to 0.99<=K/S0<1.01, and ITM refers to 1.01<K/S0. Short refers to T<30, Intermediate (Int.) stands for 30<=T<90, and long is for 90<=T.

	Mean	Median	S.D.	I.Q.R.
Panel A: Summary of the Overall Sample				
Am/S0 (Basis Points)	335.11	128.13	389.31	428.42
Eu/S0 (Basis Points)	319.30	125.19	369.14	402.09
EEP/S0 (Basis Points)	16.44	4.13	26.18	21.40
K/S0	0.98	0.99	0.08	0.13
T (Days)	58.63	51.00	39.17	52.00
Number of observations	91083			

Panel B: Summary of EEP for Moneyness and Maturity Groups									
	OTM&Short	ATM&Short	ITM&Short	OTM&Int.	ATM&Int.	ITM&Int.	OTM&Long	ATM&Long	ITM&Long
EEP/EU(Basis Points)	612.01	1006.34	344.21	280.45	401.97	596.17	313.67	449.37	817.94
EEP/S0(Basis Points)	1.04	3.41	17.61	2.01	7.45	38.61	3.33	12.52	63.74
Number of observations	9376	4186	10265	28260	3718	17949	9688	1108	6533

Table 2.2: Comparative Statics of Market EEP

This table reports the linear regression estimates for market EEP (EEP/S0*100) on option characteristics, including moneyness, moneyness-squared, and maturity (years), as well as interest rates (R*100), and dividend yields (q*100). The estimations are based on our matched put and call samples for 2005.05.02 to 2007.12.31. The underlying asset is the S&P 100 index. We pair an American option with a European option with the same strike and maturity, if their quotes are posted within 60 seconds. All explanatory variables are standardized. Robust t statistics clustered by trading dates are reported in parenthesis. *** indicates $p < 0.01$, ** indicates $p < 0.5$, and * indicates $p < 0.1$.

	(1)
	Put
VARIABLES	EEP/S0*100
K/S0	-1.636***
	(-9.355)
(K/S0)2	1.808***
	(9.983)
T/360	0.049***
	(6.009)
R*100	0.031***
	(3.688)
q*100	-0.014***
	(-5.112)
Constant	0.158***
	(33.409)
Observations	91,083
Adjusted R-squared	0.737

Table 2.3: Parameter Estimates

In Panel A, we present parameters estimates for the set of models studied in this paper, including the Black-Scholes (BS), the Heston stochastic volatility model (SV), the Merton jump (MJ), the variance gamma (VG), and the finite moment log stable (FMLS). The estimation is based on minimizing root mean squared pricing errors (RMSE) constructed with OTM European calls and puts on a daily basis. The sample period is 2005.05.02 to 2007.12.31. The underlying is the S&P 100 index. In Panel B, we provide the estimated wild card period volatility, where the wild card value is determined according to Fleming and Whaley (1994). We assume market call EEP in our sample period is driven entirely by the wild card option; the wild card volatility is thus the parameter inverted from the cross section of call EEP. We report the mean and standard error (s.e.) of each parameter time series. We also report that of the RMSE time series for each model, scaled by the underlying level.

Panel A: Model Parameters			
		Mean	s.e.
BS	σ_w	0.125	0.003
	RMSE (Basis Points)	48.721	0.610
SV	κ	7.867	0.399
	θ	0.032	0.001
	σ_V	0.486	0.016
	ρ	-0.5	0.017
	V_0	0.031	0.002
	RMSE (Basis Points)	25.143	0.342
MJ	σ_w	0.081	0.002
	λ	0.66	0.025
	δ_J	0.082	0.001
	μ_J	-0.131	0.001
	RMSE (Basis Points)	22.249	0.221
VG	σ_w	0.117	0.002
	ν	0.374	0.007
	θ	-0.157	0.011
	RMSE (Basis Points)	23.142	0. 233
FMLS	σ_w	0.082	0.002
	α	1.65	0.003
	RMSE (Basis Points)	22.357	0.220
Panel B: Wild Card Period Volatility			
		Mean	s.e.
FW model	σ_{wc}	0.131	0.002
	RMSE (Basis Points)	42.309	0. 834

Table 2.4: Alternative Models' Performances in Pricing the Cross Section of EEP

This table summarizes the RMSE generated by all five models. Pricing errors are defined as: $errModel = \frac{EEP^{Market} - EEP^{Model} - EEP^{Wild\ Card}}{S0} \times 10^4$. Model EEP are computed via the Longstaff and Schwartz (2001) algorithm, and the wild card option values are determined using the Fleming and Whaley (1994) approach, where both computations are using the estimated parameters that are summarized in Table 3. RMSE is constructed daily, using the entire cross section of EEP. We compare model performances based on the RMSE metric in accordance with Christoffersen and Jacobs (2004), because we estimate parameters based on minimizing RMSE. The sample consists of both European puts and American puts on the S&P 100 index for the period 2005.05.02--2007.12.31. We pair a European put with an American put with identical strike price and maturity, if their quotes are posted within 60 seconds. In Panel A, we report the mean and the standard error of mean for each time series of RMSE. In Panel B, we report the mean for the RMSE time series conditional on moneyness and maturity. In Panel C, we report the mean of the mean pricing error time series, where each mean pricing error is constructed using the entire cross section of EEP. Here, OTM refers to K/S0<0.99, ATM refers to 0.99<=K/S0<1.01, and ITM refers to 1.01<K/S0. Short refers to T<30, Intermediate (Int.) stands for 30<=T<90, and long is for 90<=T. *** indicates significant at the 1 percent level, based on a paired t-test.

Table 2.4 (*continued)*

	BS	SV	MJ	VG	FMLS
Panel A: Summary of the RMSE for the Overall Sample (Basis Points)					
Mean	12.113	15.138	9.857	10.856	8.797***
Std. Err	0.14	0.17	0.09	0.1	0.08

Panel B: Average RMSE Conditional on Moneyness and Maturity (Basis Points)									
	OTM&Short	ATM&Short	ITM&Short	OTM&Int.	ATM&Int.	ITM&Int.	OTM&Long	ATM&Long	ITM&Long
BS	2.610	18.078	10.635	3.157	3.541	19.729	7.717	6.040	34.715
SV	2.611	18.076	12.034	3.162	3.869	26.576	7.710	6.868	49.748
MJ	2.611	18.078	10.019	3.012	3.554	15.662	7.312	5.048	26.366
VG	2.611	18.075	10.925	3.047	3.504	17.601	7.429	5.512	29.509
FMLS	2.610	18.069	9.744***	2.984	3.127***	13.859***	7.243	4.162***	21.776***

Panel C: Average Mean Pricing Errors Conditional on Moneyness and Maturity (Basis Points)									
	OTM&Short	ATM&Short	ITM&Short	OTM&Int.	ATM&Int.	ITM&Int.	OTM&Long	ATM&Long	ITM&Long
BS	-0.047	1.328	8.151	1.457	2.621	16.784	3.193	5.197	29.298
SV	-0.052	1.456	9.805	1.452	3.077	23.713	3.145	6.156	44.711
MJ	-0.094	1.507	7.157	1.160	2.685	11.747	2.098	4.093	18.414
VG	-0.106	1.458	8.576	1.209	2.506	14.545	2.446	4.305	22.533
FMLS	-0.114	1.333	6.237***	1.104	2.094***	7.2319***	1.887	3.090***	7.8856***

Table 2.5: Transaction Cost Saving Option Values --- A Numerical Experiment

This table shows the mean and its standard error (s.e.) of the value of the transaction cost saving option embedded in American puts through a set of numerical experiments based on our augmented Longstaff-Schwartz algorithm. All reported values are scaled by the underlying level (100). All model parameters are based on the mean estimates in Table 2.3. The benchmark environmental parameters are S0=100, K = 105, T = 60, spread = 0.9 (bid = option value*spread), and p = 0.25. In each panel, we alter one of the environmental parameters while setting others to be the benchmark values. In all panels, we assume r=0.04 and no dividends.

Panel A: The Effect of Moneyness (S0=100, T=60, spread = 0.9, p=0.25)						
K/S0		0.9	0.95	1	1.05	1.1
BS	mean	0.00	0.00	0.20	28.60	25.28
	s.e.	0.00	0.00	0.00	0.01	0.00
SV	mean	0.00	0.01	0.27	16.24	25.46
	s.e.	0.00	0.00	0.00	0.02	0.00
MJ	mean	0.04	0.08	0.15	18.56	25.46
	s.e.	0.00	0.00	0.00	0.01	0.00
VG	mean	0.04	0.13	0.39	13.15	25.37
	s.e.	0.00	0.00	0.00	0.01	0.00
FMLS	mean	0.04	0.08	0.36	24.71	25.30
	s.e.	0.00	0.00	0.00	0.01	0.00
Panel B: The Effect of Maturity (S0=100, K=105, spread = 0.9, p=0.25)						
T		30	60	90	120	150
BS	mean	18.28	26.91	22.51	16.37	12.73
	s.e.	0.01	0.01	0.01	0.01	0.01
SV	mean	27.26	17.95	8.44	4.65	2.89
	s.e.	0.01	0.02	0.01	0.01	0.01
MJ	mean	12.77	20.58	27.22	21.94	19.18
	s.e.	0.00	0.01	0.01	0.01	0.01
VG	mean	13.02	13.24	13.05	24.42	2.48
	s.e.	0.00	0.01	0.00	0.01	0.01
FMLS	mean	12.76	21.48	27.31	19.07	9.91
	s.e.	0.00	0.01	0.01	0.01	0.01

Table 2.5 (*continued*)

Panel C: The Effect of (Percentage) Bid-Ask Spread (S0=100, K=105, T=60, p=0.25)						
Spread		0.8	0.85	0.9	0.95	1
BS	mean	70.05	48.62	27.94	9.29	0.00
	s.e.	0.02	0.02	0.01	0.01	0.00
SV	mean	59.73	35.00	16.83	3.95	0.00
	s.e.	0.04	0.03	0.02	0.01	0.00
MJ	mean	44.13	32.57	15.55	9.19	0.13
	s.e.	0.02	0.02	0.01	0.00	0.00
VG	mean	25.91	18.93	12.94	6.51	0.20
	s.e.	0.01	0.00	0.00	0.00	0.00
FMLS	mean	41.20	35.01	24.26	9.84	0.02
	s.e.	0.02	0.02	0.01	0.00	0.00
Panel D: The Effect of Liquidation Probability (S0=100, K=105, T=60, spread=0.9)						
p		0	0.25	0.5	0.75	1
BS	mean	0.00	28.39	32.66	36.83	35.82
	s.e.	0.00	0.01	0.01	0.01	0.01
SV	mean	0.00	16.64	13.72	11.98	11.33
	s.e.	0.00	0.02	0.02	0.01	0.01
MJ	mean	0.00	16.79	31.73	39.56	47.66
	s.e.	0.00	0.01	0.01	0.01	0.01
VG	mean	0.00	12.86	25.55	38.21	50.86
	s.e.	0.00	0.00	0.00	0.00	0.01
FMLS	mean	0.00	20.29	33.57	38.21	50.84
	s.e.	0.00	0.01	0.01	0.00	0.00

Table 2.6: The Value of the Transaction Cost Saving Option in OEX Puts

This table shows the pricing results for the transaction cost saving option values in OEX puts for 2005.05-2007.12. The values are determined based on our augmented Longstaff-Schwartz algorithm. The probability of liquidation p used here is the ratio of sell quantity (identified with the Lee and Ready (1991) algorithm) over open interest and is constructed daily based on the same 9 moneyness-maturity groups as that in Table 4. In Panels A and B, we construct the pricing error for each option series as $errModel = \frac{EEP^{Market} - EEP^{Model} - EEP^{Wild\ Card} - CostValue}{S0} \times 10000$. Based on this pricing error, we construct RMSE and MPE daily and report their time series means and stand errors in Panels A and B. We report in Panel C the estimates for the regression of pricing errors (excluding the transaction cost option value) on the transaction cost option value. Robust t statistics clustered by trading dates are reported in parenthesis. *** indicates $p < 0.01$, ** indicates $p < 0.5$, and * indicates $p < 0.1$.

Panel A: RMSE					
	BS	SV	MJ	VG	FMLS
mean	8.942	10.513	5.972	6.612	4.711
std.err	0.13	0.16	0.08	0.09	0.07
Panel B: MPE					
	BS	SV	MJ	VG	FMLS
mean	3.871	6.419	2.613	2.724	0.941
std.err	0.14	0.17	0.12	0.13	0.11
Panel C: Regressing Pricing Errors on The Transaction Cost Saving Option Values					
	(1)	(2)	(3)	(4)	(5)
VARIABLES	errBS	errSV	errMJ	errVG	errFMLS
Cost Value	0.312***	0.304***	0.339***	0.327***	0.362***
	(22.581)	(21.711)	(23.556)	(23.016)	(24.552)
Constant	5.048***	6.943***	3.180***	3.883***	1.519***
	(22.724)	(30.762)	(17.365)	(19.415)	(9.305)
Observations	91,083	91,083	91,083	91,083	91,083
Adjusted R-squared	0.254	0.203	0.291	0.287	0.313

Table 2.7: Orthogonality Test

This table shows the remaining explanatory power of option characteristics on pricing errors, defined as $errModel = \frac{EEP^{Market} - EEP^{Model} - EEP^{Wild\ Card} - Cost\ Value}{S0} \times 10000$. The sample consists of both European puts and American puts on the S&P 100 index for the period 2005.05.02--2007.12.31. Option characteristics include moneyness (K/S0), moneyness squared (K/S0)2, maturity (T/360), interest rate (R), and dividend yield (q). All explanatory variables are standardized. Robust t statistics clustered by trading dates are reported in parenthesis. *** indicates $p < 0.01$, ** indicates $p < 0.5$, and * indicates $p < 0.1$.

VARIABLES	(1) errBS	(2) errSV	(3) errMJ	(4) errVG	(5) errFMLS
K/S0	-32.430***	-48.156***	-23.612***	-24.821***	-24.499***
	(-9.297)	(-6.860)	(-9.698)	(-10.245)	(-5.669)
(K/S0)2	35.655***	55.880***	27.361***	30.907***	26.621***
	(9.877)	(7.435)	(10.188)	(10.812)	(5.870)
T/360	1.808***	2.689***	1.051***	1.258***	0.409
	(6.218)	(5.883)	(3.518)	(4.515)	(0.861)
R	0.548***	1.219***	0.026	0.047	0.011
	(3.294)	(6.864)	(0.242)	(0.672)	(0.104)
q	-0.460***	-0.634***	-0.252**	-0.334**	-0.147
	(-3.589)	(-3.600)	(-2.176)	(2.507)	(-1.391)
Constant	3.871***	6.419***	2.613***	2.724***	0.941***
	(27.651)	(31.759)	(21.775)	(20.954)	(8.554)
Observations	91,083	91,083	91,083	91,083	91,083
Adjusted R-squared	0.336	0.417	0.217	0.267	0.078

Figure 3.1: Time Series of Absolute Log Price Changes of Crude Oil

This figure plots the time series of absolute log price changes of front month contract for crude oil futures. The prices are recorded every 5-seconds for the entire period of 2006—2014.

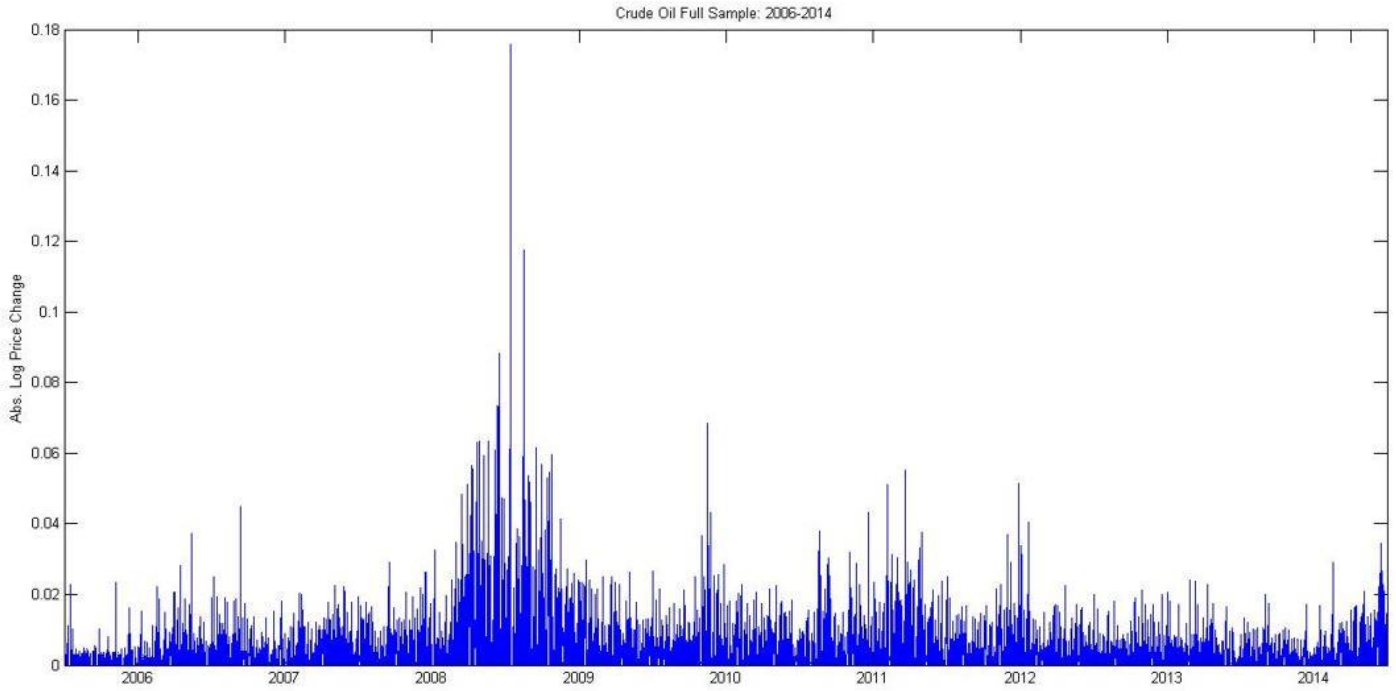

Figure 3.2: Time Series of Absolute Log Price Changes of Natural Gas

This figure plots the time series of absolute log price changes of front month contract for natural gas futures. The prices are recorded every 5-seconds for the entire period of 2006—2014.

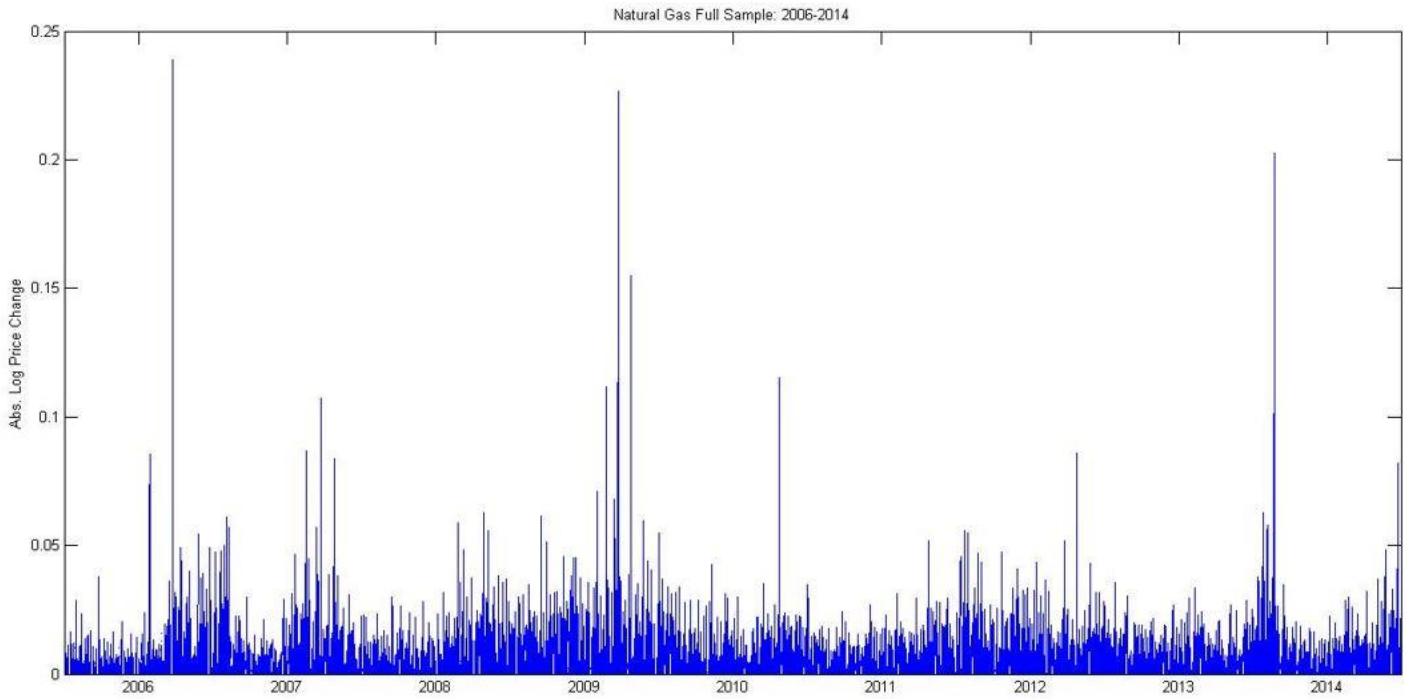

Figure 3.3: Tail-ends of the Crude Oil and Natural Gas Log Returns' Distribution

This figure plots the tail-ends of a 5-second log-return distribution for crude oil and natural gas. The log-returns were calculated from front month contract prices sampled every 5-seconds for the entire period of 2006-2014; overnight returns are discarded. The crude oil log-return sample is depicted in the top plot, while the natural gas is shown below.

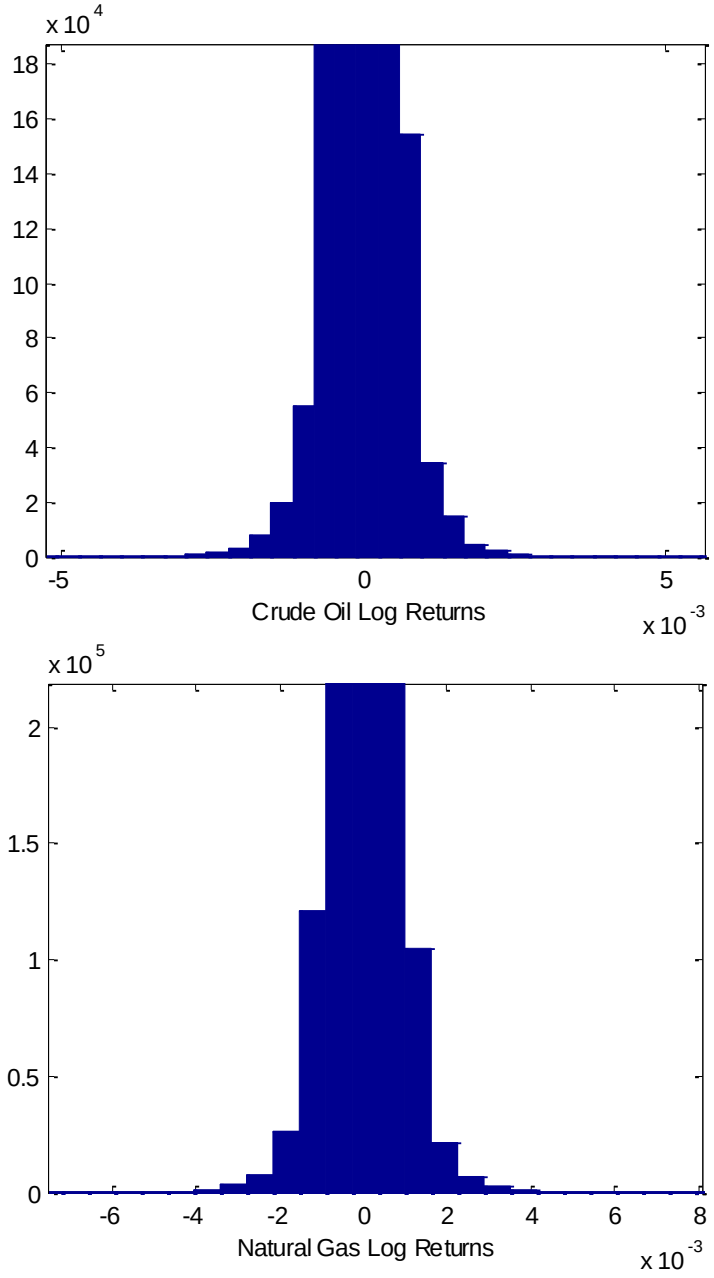

Figure 3.4: Jump Contribution to the Realized Variances of the Crude Oil and Natural Gas Log Returns' Distribution

This figure plots the jump component contribution to the realized variances for the 5-second log-returns for crude oil and natural gas. The log-returns were calculated from front month contract prices sampled every 5-seconds for the entire period of 2006-2014; overnight returns are discarded. The reported estimates are the mean values of $QV^{LJ} + QV^{SJ}$ averaged over the range of $\propto \in [4,10], \propto \in \mathbb{Z}$ and $1.5\varepsilon, 3\varepsilon$ and 4.5ε for each year in our sample. The crude oil jump contribution to realized variances is depicted in blue while the natural gas is shown in red.

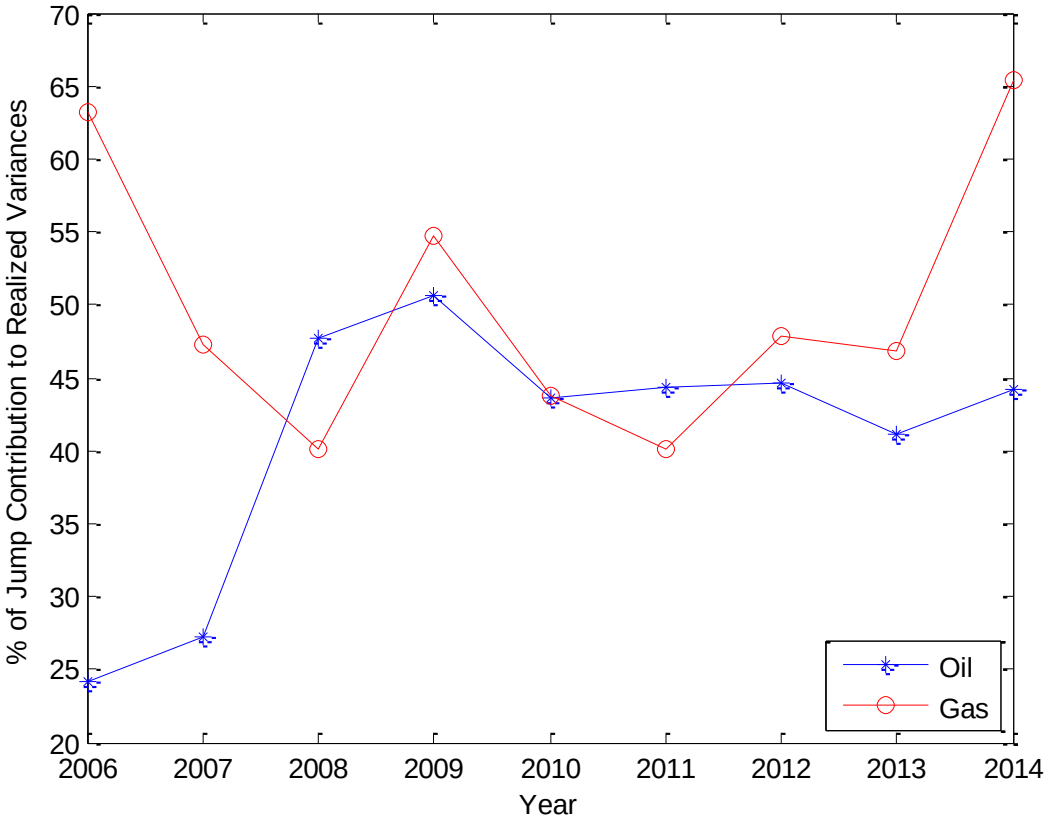

Figure 3.5: Finite Large Jump Contribution to the Realized Variances of the Crude Oil and Natural Gas Log Returns' Distribution

This figure plots the finite, large jump component contribution to the realized variances for the 5-second log-returns for crude oil and natural gas. The log-returns were calculated from front month contract prices sampled every 5-seconds for the entire period of 2006-2014; overnight returns are discarded. The reported estimates are the mean values of QV^{LJ} averaged over the range of $\propto \in [4,10], \propto \in \mathbb{Z}$ and $1.5\varepsilon, 3\varepsilon$ and 4.5ε for each year in our sample. The crude oil finite activity, large jump contribution to realized variances is depicted in blue while the natural gas is shown in red.

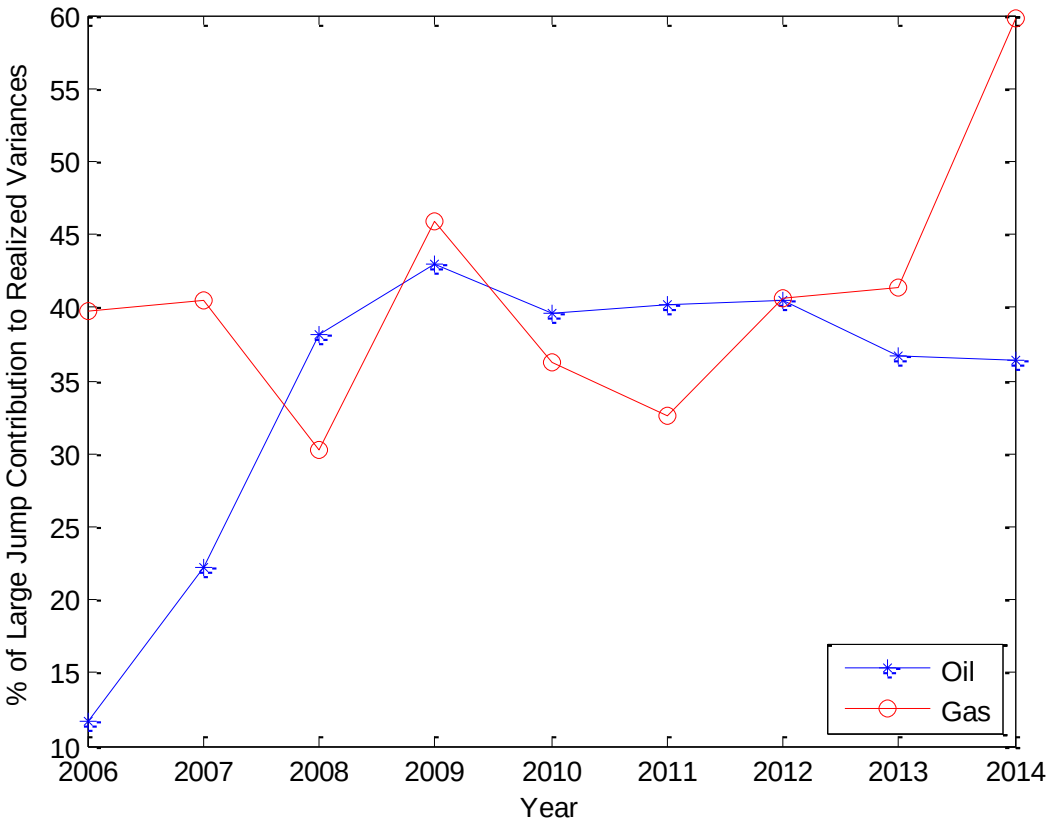

Figure 3.6: Infinite Small Jump Contribution to the Realized Variances of the Crude Oil and Natural Gas Log Returns' Distribution

This figure plots the infinite activity, small jump component contribution to the realized variances for the 5-second log-returns for crude oil and natural gas. The log-returns were calculated from front month contract prices sampled every 5-seconds for the entire period of 2006-2014; overnight returns are discarded. The reported estimates are the mean values of QV^{SJ} averaged over the range of $\propto \in [4,10], \propto \in \mathbb{Z}$ and $1.5\varepsilon, 3\varepsilon$ and 4.5ε for each year in our sample. The crude oil infinite activity, small jump contribution to realized variances is depicted in blue, while the natural gas is shown in red.

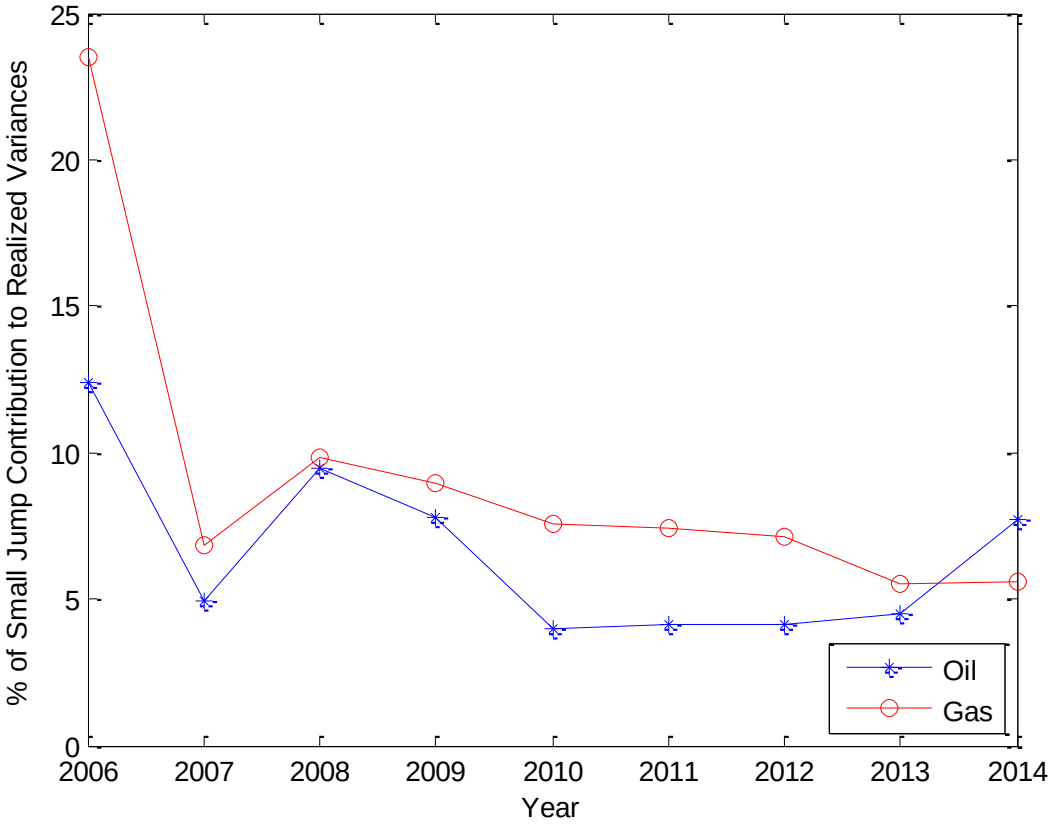

Table 3.1: Descriptive Statistics

This table describes the data sample. We use front month contract prices, sampled at 5 second intervals for all of 2006-2014, for crude oil (panel A) and natural gas (panel B) to calculate log-returns. We describe the log-returns per year and for the full sample and provide the minimum, mean and maximum values. We also include skewness and kurtosis to assess the normality of the sample. Stdev and Skew equal the 5 second return sample values$\times 10^3$, Kurt equals the 5 second sample return estimate $\times 10$ and Min and Max equal the respective minimum and maximum returns $\times 10^6$.

Panel A: The Crude Oil Sample					
	Stdev	Skew	Kurt	Min	Max
2006	0.3715	2.2529	0.3230	-0.0205	0.0372
2007	0.2955	1.9098	1.0981	-0.0251	0.0447
2008	0.5373	2.7950	3.2791	-0.0735	0.0884
2009	0.5626	37.4286	12.6030	-0.0611	0.1759
2010	0.3035	3.2304	4.3364	-0.0432	0.0684
2011	0.3446	-8.5513	2.9341	-0.0552	0.0379
2012	0.2531	-2.3796	4.3691	-0.0405	0.0515
2013	0.1981	-8.2468	2.1346	-0.0238	0.0242
2014	0.2267	-11.4922	2.3992	-0.0344	0.0266
Full sample	0.3652	15.8792	10.6130	-0.0735	0.1759
Panel B: The Natural Gas Sample					
	Stdev	Skew	Kurt	Min	Max
2006	0.7603	40.1795	12.4470	-0.0494	0.2388
2007	0.4820	12.9553	6.0704	-0.0869	0.1075
2008	0.4704	5.0064	1.4749	-0.0482	0.0626
2009	0.7076	49.6684	13.9520	-0.0458	0.2267
2010	0.4435	18.1808	5.2586e	-0.0369	0.1151
2011	0.3540	1.8909	1.3112	-0.0296	0.0519
2012	0.4884	4.5726	2.6315	-0.0549	0.0860
2013	0.3160	3.4707	1.5864	-0.0358	0.0332
2014	0.4534	-87.6051	42.9520	-0.2028	0.0589
Full sample	0.5101	24.1141	15.8580	-0.2028	0.2388

Table 3.2: Test for Jumps in General

This table shows the results for testing the null of no jump. The test statistic $S_J = B(4, \infty, 2\Delta_n)/B(4, \infty, \Delta_n)$ converges to 2 under the null of no jumps, where Δ_n denotes the sampling frequency (5 seconds). We calculate the asymptotic variance for S_J under the null and the associated 95 percent critical threshold C_{95}^{Normal} for a range of α. The null is rejected if $S_J < C_{95}^{Normal}$. We use both the crude oil and the natural gas samples for the entire period of 2006-2014. The series is sampled every 5 seconds.

Panel A: The Crude Oil Sample							
S_J	1.004						
α	4	5	6	7	8	9	10
C_{95}^{Normal}	1.982	1.975	1.967	1.958	1.947	1.935	1.924
Panel B: The Natural Gas Sample							
S_J	0.994						
α	4	5	6	7	8	9	10
C_{95}^{Normal}	1.985	1.980	1.975	1.969	1.961	1.954	1.945

Table 3.3: Tests for Finite Activity or Infinite Activity Jumps

This table shows the results for testing the null of finite activity. The test statistic $S_F = B(4, \alpha, 2\Delta_n)/B(4, \alpha, \Delta_n)$ converges to 2 under the null of no jumps, where Δ_n denotes the sampling frequency (5 seconds). We calculate the test statistics and its asymptotic variance for under the null and the associated 95 percent critical threshold C_{95}^{Normal} for a range of α. The null is rejected if $S_J < C_{95}^{Normal}$. We use both the crude oil and the natural gas samples for the period of 2006-2014. The series is sampled every 5 seconds.

Panel A: The Crude Oil Sample							
α	4	5	6	7	8	9	10
S_F	1.213	1.241	1.283	1.267	1.263	1.250	1.242
C_{95}^{Normal}	1.937	1.913	1.887	1.855	1.819	1.779	1.741
Panel B: The Natural Gas Sample							
α	4	5	6	7	8	9	10
S_F	1.156	1.163	1.165	1.146	1.125	1.123	1.106
C_{95}^{Normal}	1.949	1.933	1.915	1.893	1.868	1.843	1.813

Table 3.4: Tests for the Presence of Brownian Motion

This table shows the results for testing the null of presence of Brownian motion. The test statistic $S_W = B(1.5, \alpha, \Delta_n)/B(1.5, \alpha, 2\Delta_n)$ converges to $2^{0.25} = 1.189$ under the null of presence of Brownian motion, where Δ_n denotes the sampling frequency (5 seconds). We calculate the test statistic and its asymptotic variance under the null and the associated 95 percent critical threshold C_{95}^{Normal} for a range of α. The null is rejected if $S_W < C_{95}^{Normal}$. We use both the crude oil and the natural gas samples for the entire period of 2006-2014. The series is sampled every 5 seconds.

Panel A: The Crude Oil Sample							
α	4	5	6	7	8	9	10
S_W	1.244	1.237	1.232	1.230	1.228	1.227	1.226
C_{95}^{Normal}	1.189	1.189	1.189	1.189	1.189	1.189	1.189
Panel B: The Natural Gas Sample							
α	4	5	6	7	8	9	10
S_W	1.202	1.195	1.191	1.190	1.188	1.187	1.186
C_{95}^{Normal}	1.189	1.189	1.189	1.189	1.189	1.189	1.189

Table 3.5: Percentage of Total Return Variation from Jumps

This table shows the results for decomposing the total return variation, the quadratic variation ($QV = B(2, \infty, \triangle_n)$), into percentages generated by the continuous component ($QV^c = B(2, \alpha, \triangle_n)/QV$), the large jump component ($QV^{LJ} = U(2, \epsilon, \triangle_n)/QV$), and the small jump component ($QV^{SJ} = 1 - QV^C - QV^{LJ}$). In the second column, we report the range of ϵ, and the range of α is reported as usual. In the last column, we report the average QV in terms of various α. The row "mean" report the average QV in terms of different ϵ. We use both the crude oil and the natural gas samples for the entire period of 2006-2014. The series is sampled every 5 seconds.

Panel A: The Crude Oil Sample									
	ϵ/α	4	5	6	7	8	9	10	mean
QV^C		61.29%	62.57%	63.41%	64.11%	64.73%	65.33%	65.88%	63.90%
QV^{LJ}	1.5α	36.59%	35.56%	34.67%	33.81%	32.89%	31.94%	30.86%	33.76%
	3α	32.89%	30.86%	28.53%	26.91%	24.37%	22.40%	20.11%	26.58%
	4.5α	28.53%	25.36%	22.40%	19.37%	17.14%	15.87%	14.65%	20.47%
mean		32.67%	30.59%	28.54%	26.70%	24.80%	23.40%	21.87%	26.94%
QV^{SJ}	1.5α	2.12%	1.87%	1.91%	2.08%	2.37%	2.74%	3.26%	2.34%
	3α	5.82%	6.57%	8.06%	8.98%	10.90%	12.27%	14.01%	9.51%
	4.5α	10.18%	12.07%	14.19%	16.53%	18.12%	18.81%	19.47%	15.62%
mean		6.04%	6.84%	8.05%	9.20%	10.47%	11.27%	12.25%	9.16%
Panel B: The Natural Gas Sample									
	ϵ/α	4	5	6	7	8	9	10	mean
QV^C		55.95%	57.73%	58.87%	59.73%	60.50%	61.09%	61.77%	59.38%
QV^{LJ}	1.5α	41.13%	39.84%	38.91%	37.87%	36.59%	35.41%	34.24%	37.71%
	3α	36.59%	34.24%	31.82%	29.29%	27.09%	24.92%	23.20%	29.59%
	4.5α	31.82%	28.25%	24.92%	22.16%	19.89%	17.68%	16.65%	23.05%
mean		36.51%	34.11%	31.88%	29.78%	27.86%	26.01%	24.70%	30.12%
QV^{SJ}	1.5α	2.92%	2.43%	2.22%	2.40%	2.91%	3.50%	3.99%	2.91%
	3α	7.46%	8.03%	9.30%	10.98%	12.41%	13.99%	15.03%	11.03%
	4.5α	12.23%	14.02%	16.20%	18.11%	19.61%	21.23%	21.58%	17.57%
mean		7.54%	8.16%	9.24%	10.50%	11.64%	12.90%	13.54%	10.50%