Essays on the effects of index trading on asset prices

by

Joseph DeCoste

A thesis submitted in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

in

Finance

Faculty of Business
University of Alberta

Abstract

In the first essay, I explore whether excess demand in commodity futures markets affects the spot price of oil. I use a sign restricted vector autoregressive oil market model that explicitly includes futures markets. This model allows for the detection of futures demand effects which feedback into spot prices through a price signaling channel, in contrast to previous studies relying solely on an inventory channel. I find novel evidence that excess demand in futures markets drives over half of the short run variation in the spot price of oil, and can explain puzzling incidents of oil price behavior such as the 2008 boom and bust in oil prices and the 2014 oil price crash. I find that this relationship is much stronger after 2003, the period commonly associated with a rise in financialization and commodity index investment.

In the second essay, I test for the existence of excessive comovement amongst stocks in the S&P 500. Using a fuzzy regression discontinuity approach, I show that membership in the S&P 500 leads to significant positive excess comovement in the long term. I evaluate a traditional, liquidity based explanation and a friction based explanation, and find no evidence that liquidity drives excess comovement. I show that the lack of evidence for excess comovement shown in Chen, Singal, Whitelaw (2016) is due to heterogeneous effects on newly included firms versus established members. One potential explanation is that investors take time to fully integrate the new stock into the group immediately after inclusion, reducing observed increases in comovement in the short term. Another is that firm inclusion is related to a change in fundamentals. These results constitute new evidence of frictions when exposed to large classes of traders with correlated, non-fundamental demands, such as those populating the S&P 500.

In the third essay, I test for the existence of excess coskewness amongst stocks in the S&P 500. Using a combination of event study and fuzzy regression discontinuity approaches, I show that membership in the S&P 500 leads to significant negative excess coskewness in the long term, but positive excess coskewness in the short term, pointing to important transitory effects of inclusion that differ from persistent long term effects. These coskewness results point to price distortions caused by index membership, with implications for both market and allocative efficiency, and diversification benefits.

Acknowledgements

I greatly thank Professor Masahiro Watanabe, my supervisor, for his guidance, patience, and advice throughout my time in the Ph.D. program. His thorough, timely, and helpful feedback, alongside his encouragement and support, whenever it was needed, will always be appreciated.

I thank my thesis committee members, Professors Aditya Kaul, Felipe Aguerrevere, and Sahil Raina for their positive encouragement, support, and advice that helped significantly improve and shape this thesis. I also thank my external committee members, Professors Sebastien Fossati and Liyan Yang, and my committee chair Vikas Mehrotra for their advice, questions, and willingness to participate in my examination process.

Of course, my thesis would not be possible without the efforts of the much wider Alberta School of Business community, and I would like to thank my instructors, the Finance administrative staff, and my fellow Ph.D. students for their support. I also thank the wonderful team in our Business Ph.D. Office, Debbie Giesbrecht, Helen Wu, and Trish Reay, for their sincere and exemplary dedication to Ph.D. students at the University of Alberta.

I also appreciate the support and friendship of many of my fellow Ph.D. students, and I am particularly thankful to my Finance Cohort colleagues, Rodrigo Cardenas and Yusuke Tsujimoto. I also acknowledge the generous financial support the Ph.D. office and AIMCO provided during my Ph.D. studies.

Most of all, I am deeply grateful to my family, particularly my wife Stacey Gallant, for her love, sacrifice, support, patience, and encouragement throughout my long and winding journey.

Table of Contents

Page

List of Tables

List of Figures

Chapter 1

Introduction

The bundling of assets into indices encourages investors to trade the entire basket instead of evaluating the investment merits of individual constituents within that basket, effectively "financializing" the index of assets itself. This can create investment demand not linked to an asset's economic fundamentals, putting pressure on asset prices. In this thesis, I explore the effects of this index financialization on commodity spot prices (Chapter 2), equity comovement (Chapter 3), and equity coskewness (Chapter 4).

In the first essay, I explore the influence of commodity index trading in futures markets on commodity spot prices. Any influence is important because a commodity spot price distortion can constitute a distortion in the real economy by affecting firms' production decisions. This also has implications for the affordability of essential resources like food and energy for individual households.

To conduct my study, I utilize a sign restricted structural vector autoregressive model adapted from the workhorse model of Kilian and Murphy (2014). I use this model to separate the effects of excess futures market demand from changing economic fundamentals. Specifically, this model lets me capture excess demand transmitted from futures markets to spot markets through a price signaling channel that has not been explored before. This contrasts previous models relying on a detectable inventory response that may not exist when observable information about economic fundamentals is imperfect.

I find that excess futures market demand drives the largest portion of short-run oil price variation. This effect is primarily concentrated after 2003, during the period of rising commodity index trading, directly linking it with increasing financialization.

I also analyze the evolution of oil prices and show that futures market shocks explain major price swings, most notably during 2008 and 2014. These results constitute evidence of large impacts of financialization on commodity spot markets, empirically linking trading pressure in futures markets to spot prices through a price signaling channel.

My second essay studies the influence of financialization of equity indices on the return process of equities. I evaluate whether membership in the S&P 500 can lead to long-term changes in a stock's comovement with the index. This is important because a change in comovement for members indicates a distortion to the firm's fundamental return process, since a change in index status on its own should not change the firm's business fundamentals or returns. Any distortion of the firm's return process has implications for the informational efficiency of prices and the efficiency of capital allocation in the economy.

To conduct my study, I use a fuzzy regression discontinuity design. This quasi-experimental design allows me to effectively control inevitable differences between index and non-index firms and isolate the causal effect of index membership on comovement. I find that S&P 500 member firms indeed experience positive excess comovement with each other due to index membership. This is a new long-run finding and complements literature that debates the short-run change in comovement upon index inclusion.

In the third essay, I evaluate the effect S&P 500 membership has on coskewness with the index. As this is an unexplored area of research, I examine both the long-run and the short-run impact of index membership using a combination of an inclusion event study and a regression discontinuity design. Assets with higher coskewness have a higher likelihood of moving with the market during market upswings but a lower likelihood of moving with the market during market downswings, which is generally seen as beneficial to investors. I find that transitory effects increase coskewness when firms are first added to the index, but the long-run effect is negative. This points to a novel impact of S&P 500 membership inducing undesirable decreases in coskewness, with potential implications for prices, investor portfolio choice, and diversification benefits.

Chapter 2

Does excess futures market demand affect the spot price of oil?

2.1 Introduction

Can excess demand in commodity futures markets distort spot prices? This critical question has been at the center of a long, contentious debate since the dramatic boom and bust of commodities in 2008.[1] Traditional economic theory implies that the answer to this question is no. Yet, puzzling price swings in commodities after 2003 (Figure 2.1),[2] coinciding with the rise of index investing and historic increases in the comovement of commodities with broader financial markets, provide fuel for the argument that financialization could adversely affect the ability of prices to incorporate information accurately. My examination supports this argument; I find that not only is excess futures market demand a key determinant of commodity prices, it is also increasingly important in the recent period linked to financialization.

My research question is motivated by the theoretical work of Sockin and Xiong (2015), who show how upward pressure placed on futures prices by commodity index investors can be mistaken as a positive economic signal by goods producers.[3] These goods producers subsequently increase their demand for oil as an input in production.[4]

[1]Cheng and Xiong (2014) and Fattouh, Kilian, and Mahadeva (2013) take two different perspectives on the likelihood of futures market demand being problematic. Both critically review previous studies and agree that there is no definitive evidence either way.

[2]The beginning of the financialization of commodities markets is generally placed around 2003-2004. Cheng and Xiong (2014) associate this period with rapidly increasing capital investment from commodity index traders and growing gross and net positions in commodity futures contracts. Bouchouev (2020) associates the period with increasing liquidity to hedgers as futures markets outstrip the size of physical markets and a related shift towards Contango.

[3]Not to be confused with oil producers, goods producers consume oil as an input into their production of intermediate and final goods

[4]There is empirical evidence that firm managers use asset prices as signals in decision making,

This serves as a novel mechanism whereby pressure on futures prices from uninformed investors can feedback into spot market prices through a price signaling channel.

To study the influence of financialization on commodity prices, I examine the oil market, primarily for the richness of data available. Specifically, given the global nature of oil markets, I examine the price of Brent oil. I utilize a structural vector autoregressive oil market model based on Kilian and Murphy (2014) to deal with the highly endogenous nature of observable global oil market variables. I extend the model from the traditional fundamental shocks of oil supply, oil demand, and inventory demand, to include a non-fundamental futures demand shock. I utilize data on ICE Brent Futures contracts to incorporate this new shock.

The model relies primarily on a set of impact sign restrictions on economic responses derived from economic theory to identify and disentangle the largely unobservable shocks from the observable data. The fundamental shocks are relatively straightforward: a negative supply shock is a production disruption increasing the spot price of oil, a positive demand shock directly pressures the spot price upwards, while a positive inventory demand shock similarly pushes oil prices upwards, but on the back of increased inventories instead of increased oil consumption. A key novelty lies in the modeling of the futures demand shock to disentangle it from the fundamental demand shock. To do this, I combine the theory of Sockin and Xiong (2015) with the implications of the theory of storage. I show that, while a demand shock results in a decrease in the futures-spot spread, a futures demand shock is accompanied by an increase in spread. I also supplement the model with additional dynamic sign restrictions, realistic restrictions on the elasticity of oil demand in use, and a set of narrative restrictions to ensure the model can credibly explain several observable

and that using asset prices as signals can lead to poor decisions and be damage firms, in Dessaint, Foucault, Fresard and Matray (2019) and Brogaard, Ringgenberg and Sovich (2019).

exogenous events.

My results show that futures demand shocks, independent of fundamental information, explain over half of the short run variation in Brent oil prices. Furthermore, the importance of these shocks has increased over time, explaining just 6% of short run variation before 2003 but over half of short run variation after. I argue that this is consistent with an observed increase in financialization, specifically, rising levels of index investment. In contrast to the short run results, demand shocks are the most important driver of long run oil prices, with futures demand playing a much weaker long run role.

My analysis also reveals significant contributions of excess futures demand to dramatic oil price swings during my sample. Specifically, I find that excess demand in futures markets drove the largest portion of the run up and collapse in oil prices in 2008. Additionally, I find that the oil price crash in 2014, while triggered by supply shocks, was heavily exacerbated by futures trading. My results also point to earlier contributions to oil price behavior during the bond market crisis of 1994 and the Asian financial crisis of 1997. My model captures fundamental oil market behavior throughout the sample as well. Specifically, I capture supply shocks during the Venezuelan oil strike of 2002, the US invasion of Iraq in 2003, the Libyan Civil War in 2011, the Abqaiq-Khurais drone attacks on Saudi oil facilities in 2019, and the Saudi-Russia oil price war in 2020. I also capture a major demand shock depressing oil prices around the onset of the global Covid-19 pandemic.

Finally, I evaluate other possible explanations for my results. I show that it is unlikely that my futures demand shock is driven by changing interest rates, changing storage costs, or by differences in the speed of information diffusion to spot and futures markets. I also confirm that my futures demand shock captures financial market pressure, not hedging pressure.

For emphasis, the contributions of my study are not limited solely to oil markets. Financialization has changed the relationship of many indexed commodities with financial markets, and many commodities share similarly puzzling price behavior, most notably during 2008. Furthermore, the implications of my study are not limited to commodities and instead indicate the propensity of financialization to affect both financial markets and the real economy more generally.

The remainder of the paper is organized as follows: Section 2 provides a brief background to my study and elaborates on its contributions, while section 3 presents my structural vector autoregressive model, and section 4 describes the data I use. Section 5 presents my main results on the financialization of commodities, including the drivers of oil spot prices, the importance of futures demand, how this has changed over time with financialization, and how futures trading has affected real world oil prices throughout the sample period. Section 6 concludes.

2.2 Background and Contribution

My results contribute to several lines of literature. First, my evidence directly relates to the aforementioned debate on the causes of the 2008 commodity price boom and the extended debate on the financialization of commodities generally. Masters (2008) attributes the price boom to a bubble caused by increased demand for commodity exposure by commodity index traders, while Hamilton (2009) instead argues that emerging market demand justified increasing prices. Cheng and Xiong (2014) point out that growth in emerging market demand had already begun to slow as early as late 2007, while prices continued to rise another 40% through mid-2008 before collapsing, posing a puzzle to fundamental demand based explanations. [5] Early

[5] China's GDP growth peaked in mid-2007, world equity indices peaked in the fall of 2007, the US entered a recession in Dec 2007, Bear Stearns collapsed in the spring of 2008, and so on.

studies relying on granger causality tests do not detect speculative effects leading up to 2008, but have been criticized as suffering from simultaneity bias.[6][7][8] However, Tang and Xiong (2012) point to increasing correlations of index commodities with wider financial markets, and Henderson, Pearson, and Wang (2015) point to the results of a natural experiment linking uninformed investor flows to price changes. These papers suggest that financialization is indeed an essential determinant of commodity prices.

Second, my results add to a wider body of evidence linking the demand of index investors to asset prices generally. Shleifer (1986) and Jain (1987) link equity index inclusion to non-fundamental changes in equity prices, while Barberis, Shleifer, and Wurgler (2005) link equity index inclusion to increasing asset correlations. Chen, Singal, and Whitelaw (2016) point to changing fundamentals upon index inclusion to cast doubt upon inclusion studies, while chapter 3 of this thesis confirms increased return comovement for index members using a regression discontinuity design instead. Harford and Kaul (2005) show that correlated order flow drives strong common effects in the returns of S&P 500 index firms.[9]

Third, my results contribute to the literature on the effects of financial markets on the real economy. Bond, Edmans, and Goldstein (2012) argue the importance of accounting for the feedback effect of market prices on the real economy and show that doing so can help explain a number of puzzling phenomena. Along this line,

[6] See Cheng and Xiong (2014) for a comprehensive review of the identification problems in this area of research.

[7] Early studies by Irwin, Sanders and Mirrin (2009), Stoll and Whaley (2010), Irwin and Sanders (2012), and Hamilton and Wu (2015), among others, find no predictive link between investor flows and futures prices, while Singleton (2014) uses a similar approach with a different measurement of investor changes and finds that investor position changes do predict future price changes.

[8] I verify that there is no predictive link in my sample and present the results of my granger causality test, linking net changes in swap dealer flows from the CFTC's DCOT report, in Table A.5.

[9] See Kaul, Mehrotra, and Morck (2000), Chen, Noronha, Singal (2004), and Froot and Dabora (1999) for additional evidence

Dessaint, Foucault, Fresard, and Matray (2019) link firms' investment decisions to non-fundamental drops in the stock price of their product-market peers. Brogaard, Ringgenberg, and Sovich (2019) show that firms make worse decisions and exhibit lower performance when the index commodities they use in production experience higher degrees of financialization, implying that they reference commodity prices in decision making. My work complements these studies by providing evidence supporting a feedback loop of futures prices into economic decision making, consistent with the theory of Sockin and Xiong (2015).

Finally, my study builds upon an existing literature searching for speculative effects in oil prices through changes in inventories. Kilian and Murphy (2014) and Knittel and Pindyck (2016) measure the effect that speculative trading has on oil prices through an inventory response channel, capturing the extent to which changing expectations of future economic fundamentals affect commodity prices through changing demand for inventories. My model adds the flexibility to study how excess demand for long positions in futures contracts, unrelated to economic information, can affect spot prices. These shocks affect spot prices through a price signaling channel that induces a demand response instead of an inventory demand channel.

2.3 Empirical Design

My analysis aims to establish the relationship between excess demand in futures markets and the spot price for oil. However, a simple regression of the spot price of oil on measures of trading activity will suffer from severe endogeneity. First, there is an omitted variable problem, as additional variables, such as economic strength, could simultaneously drive the price of oil and trading.[10] Second, if the price of oil

[10]For example, if the economy is strong, people may have more money to invest in long only index funds, and oil demand may be higher.

rises for some other reason, it may attract speculators who observe the price increase and wish to trade on momentum, leading to a reverse causality problem. Thus, a change in the level of trading does not represent an independent or exogenous shock. To deal with these issues, I use a structural vector autoregressive (SVAR) model, adapted from Kilian and Murphy (2014), to integrate a futures demand shock into a model of global oil prices.

2.3.1 Structural Vector Autoregressive Model

To model global oil markets and analyze the drivers of global oil prices, I posit the following set of structural relationships:

$$B_0 y_t = \beta_0 + \sum_{i=1}^{24} \beta_i y_{t-i} + u_t \tag{2.1}$$

Where y_t represents a vector of endogenous variables, including percent change in global crude oil production, global real activity, the real price of crude oil, changes in inventories of crude oil, and the real price of crude oil futures. The model includes five independent structural shocks, represented by u_t, which drive observed changes in these variables. The shocks include flow supply and demand shocks, an inventory demand shock, my futures demand shock capturing pressure from excess demand on prices independent of economic fundamentals, and a residual. B_0 is an invertible matrix containing estimates which capture the contemporaneous relationship of the shocks with the variables and maps the structural shocks simultaneously to all of the observable reduced form variables y_t at the time they occur. For example, B_0 can map an independent demand shock at time t to immediate and simultaneous effects on observable economic activity and the oil price. Since the independent shocks are not directly observable, we must disentangle their influence from the relationships between observable variables. To do so, I run the following monthly, reduced form regressions with 24 lags (seasonal dummies are suppressed):[11]

$$y_t = A_0 + \sum_{i=1}^{24} A_i y_{t-i} + \epsilon_t \tag{2.2}$$

Since this regression fails to capture the contemporaneous relationship between the regressors, the ϵ_t are correlated across variables and do not reflect independent, exogenous shocks. Hence, an underlying shock could cause innovations to the real price of oil to be related to innovations in inventory levels, real activity, production, and so on. Specifically, the reduced form innovations and the structural shocks are linearly related as $u_t = B_0^{-1}\epsilon_t$. The structural shocks, u_t, are then disentangled and recovered from the reduced form innovations, ϵ_t, from (2.2), by estimating and imposing theoretically sound restrictions on B_0. I utilize the reduced form estimates from (2.2) and a suitably identified B_0 to recover the structural shocks, structural impulse response functions, and conduct a variance decomposition and historical decomposition of oil prices. Appendix B contains additional details on this process for the interested reader.

2.3.2 Identification

In this section, I define the observable effects of the structural shocks that I estimate and detail how the structural model, specifically B_0, is restricted to identify them. I use a combination of static and dynamic sign, elasticity, and narrative restrictions to identify the model. The shocks include fundamental shifts in oil supply, oil demand, and inventory demand, along with non-fundamental pressure from futures demand trading. I normalize all shocks to positively affect oil prices, as is the convention in

[11] Kilian (2009) illustrates the importance of including long lags to accommodate cycles within the market.

the literature, and summarize the impact sign restrictions in Table 2.1.

Flow Supply and Flow Demand Shocks

The first two shocks in the model are standard supply and demand shocks, which have traditionally been the focus of the oil market literature. A flow supply shock in the model shifts the supply curve to the left along the demand curve, resulting in a lower quantity of oil produced and a higher oil price. With less oil consumed in goods production, real activity decreases. I also restrict the response to a negative supply shock to be positive for the price of oil and negative for production and real activity for at least 12 months, in line with Kilian and Murphy (2014).

In a similar vein, a flow demand shock shifts the oil demand curve to the right along the oil supply curve, increasing the equilibrium level of production and the price of oil. In this shock, oil is consumed by final goods producers, so real activity increases. Contrary to earlier research, recent work by Kilian and Murphy (2014) points to demand factors playing a more prominent role than supply.

Inventory Demand Shock

Since oil is storable, an increase in demand for oil can occur without an increase in real activity. This gives rise to an inventory demand shock, visible by a similar shift in the demand curve for oil to the right along the supply curve, with increased production and a higher price, but accompanied instead by an increase in inventories and a decrease in real activity. Inventory demand can capture changing expectations about future supply and demand conditions for oil. Market participants may decide to hold their oil for future sale if an expected future oil shortfall indicates a rising price. Similarly, during times of increased uncertainty, producers may choose to hold more inventory to protect against expensive production disruptions.

Futures Demand Shock

I now define the futures demand shock, which is the primary interest of this paper. In short, this shock first increases futures prices but is misconstrued as a flow demand shock and exhibits all of the characteristics therein. It is disentangled from a flow demand shock by a spread response implying a non-fundamental increase in the futures price.

This shock is meant to capture the effect of excess demand for long positions in commodity futures unrelated to fundamental information. The most typical example would be a flow into a commodity index investment, driven by an exogenous desire for diversification or an exogenous change in wealth. Such index investors typically generate exposure to the underlying commodities, directly or indirectly, via futures markets (or OTC swap dealers, who ultimately pass the exposure along to futures markets), and these flows put upward pressure on futures prices. Conversely, a negative futures demand shock reflects downwards pressure on prices due to an excess decrease in demand for long positions. In the context of index trading, this could reflect a change in financial market demand for commodity exposure within a wider portfolio allocation strategy, or redemptions to meet liquidity needs.

Beyond the impact on futures prices, I follow the definition provided by Sockin and Xiong (2015) for a futures demand shock which can affect the demand for, and price of, oil in the spot market. In their model, goods producers view a rising futures price as a signal of higher economic strength, and thus, higher demand for their finished goods. They then purchase more oil to input into production, effectively increasing demand for oil, and putting upward pressure on spot market prices. Without incorporating information on futures prices, a futures demand shock is therefore observationally equivalent to a flow demand shock.

To disentangle the two shocks, I use information on futures market prices to impose restrictions on the futures-spot spread implied by the theory of storage.

To illustrate, consider the characterization of the futures-spot spread of Pindyck (1994):

$$f_{t,t+\tau} - p_t = -\psi(p_t, N_t, E[Q_{t,t+\tau}]) + \tau r_t \tag{2.3}$$

Which simply states that the spread between the log futures price and the log spot price is a function of the convenience yield ψ and the relatively stable interest rate r_t. Focusing on ψ, Pindyck (1994) shows that the convenience yield is increasing in price, since higher prices imply higher convenience, decreasing in inventories N_t, as higher inventories are subject to decreasing marginal benefit, and increasing with expectation of future tightness in the supply of oil, here represented by demand $E[Q_{t,t+\tau}]$.[12] While Kilian and Murphy (2014) do not use the spread in their model, this same theory underpins it.[13] Lombardi and Van Robays (2011) also use the spread to identify their "destabilizing" shock.

The relationship above implies that if there is a negative supply shock, both futures and spot prices will increase. However, the convenience yield will also increase, resulting in a decrease in the futures-spot spread in Eq(3). The same rationale holds for a positive flow demand shock. Applying the same logic to the inventory demand shock, we would expect a decrease in the futures-spot spread since an increase in the convenience yield defines an inventory demand shock. In contrast, if a futures

[12] To show this, Pindyck (1994) estimates a model for ψ, due to Brennan (1991), and finds that the model, with the above properties, is a good match to the data. Specifically, he estimates the model $\psi = \beta P_t(\frac{N_t}{Q_{t+1}})^{-\phi}$ where N_t is inventory, and Q_{t+1} represents quantity of the commodity of interest. He estimates that β and ϕ are positive, which means ψ is increasing in P_t and Q_t and decreasing in Q_{t+1}.

[13] For additional empirical support for the theory of storage, see Fama and French (1988), Brennan (1991), Ng and Pirrong (1994), and Gorton, Hayashi and Rouwenhurt (2013).

demand shock moves the futures price away from its fundamental value, the futures price will be higher than it would otherwise, increasing the futures-spot spread.

Elasticity Restrictions

It is important that my model yield credible elasticity estimates. A key innovation in Kilian and Murphy (2014) is the introduction of elasticity restrictions. They conclude that such restrictions are necessary to appropriately identify oil market models, playing a key role in invalidating models that overestimate the importance of supply shocks. Specifically, I impose bounds on the impact price elasticity of oil supply (between 0 and 0.10) and the impact price elasticity of oil demand in use (between -0.80 and 0).

Narrative Restrictions

I further sharpen the inference in my sign restricted model by utilizing a set of narrative restrictions on the structural shocks and historical decompositions to verify that the model agrees with several plausibly exogenous events established in the oil market literature. These exogenous events are rare and usually surround unexpected disasters, wars, or other political events. Following Antolin-Diaz and Rubio-Ramirez (2018), I remove any models which do not capture the major supply disruptions caused by the Venezuelan Oil Strike of 2002, the Invasion of Iraq in March 2003, and the Libyan Civil War in February 2011. All of these events caused significant production disruptions in major oil producing nations, and were arguably exogenous to oil supply and demand fundamentals.[14] Specifically, a model with no negative

[14] Antolin-Diaz and Rubio-Ramirez (2018) identify additional exogenous events which occur either before, or too early in our sample period to utilize. Notably, the Gulf War in 1990 has been beneficial. This event happens too early in my sample, where transition dynamics cannot be fully modeled (since my model has 24 months of lags), so it is left out.

structural supply shock during the event month, or where supply shocks are not the primary driver of unexpected changes in production, is considered incredible.

I also introduce a new narrative restriction related to the global Covid-19 pandemic, which arose in early 2020. This is a rare example of a visible exogenous shock affecting demand. Unlike the above supply shocks, the extent of the pandemic was revealed over several months; however, the period of March and April 2020 is when major international travel closures and widespread business disruptions occurred in much of the world. [15] Therefore, models in which demand shocks do not play an important, primary role in the determination of oil prices during March-April 2020 in particular, are considered incredible.[16]

Why investor flow is not included in the model

I intentionally omit an investor flow variable for several reasons. First, and most significant, is that including an investor flow variable is insufficient to identify a futures demand shock within a sign restricted SVAR model because the theoretical response of both the futures demand shock and the flow demand shock are the same—an increase in investor flow. Second, the sample size is insufficient to support an additional variable, and all other variables serve important identification purposes. Third, even for secondary analysis, there is insufficient data to include investor flow in a global oil market study. The best investor flow data is specific to the US oil market and WTI Crude futures, instead of Brent Crude futures, and covers a very limited period, making it inappropriate for a global oil market study.[17]

[15] The WHO declared the Covid-19 outbreak a pandemic on March 11th, 2020.

[16] Results are qualitatively similar, including credibly capturing most of these events, even when narrative restrictions are omitted, as presented in Appendix A, Figures A.1-A.5, and Table A.2.

[17] Kilian and Zhou (2020) outline the issues of sample size, discuss the maximum number of variables, and point to the use of US local market variables as a fundamental mistake in oil market SVARs.

Finally, the desirability of using investor flow data is unclear. Findings relying on available investor flow data are subject to criticism, as the data depends on trader classifications which may not align with the actual trading motives of participants. Cheng and Xiong (2014) argue that the trading activity of traditional commercial hedgers is consistent with significant amounts of speculative trading, and that all trader groups appear to trade speculatively at the margin. Additionally, they argue that no trader group can be treated as plausibly exogenous.

Omitted Variable Bias

The points made above allude to one of the potential limitations of the SVAR methodology generally. Since the sample size in an SVAR can only accommodate a limited number of variables, there is a possibility of an omitted variable related to the real price of oil, which biases the results. However, Kilian and Zhou (2020) review several studies and conclude that the standard oil market framework I build upon is remarkably robust to omitted variables.

2.4 Data

I utilize monthly measures of global oil production, global real economic activity, global real price of oil, global oil inventories, and oil futures prices that are standard to the oil market literature.[18] Specifically, my measure of global oil production is log global production including lease condensates available from the Energy Information Administration. I use the Dry Cargo Shipping Rate Index to measure real economic activity, developed in Kilian (2009) and Kilian and Zhou (2018), and maintained by the Federal Reserve Bank of Dallas. To compute the global real price of oil, I deflate

[18] See Kilian and Murphy (2014) and Kilian and Zhou (2020) for a more comprehensive discussion of variable selection in oil market SVAR models

the Brent oil spot price, available from the EIA, by the US consumer price index, available from the Bureau of Labor Statistics, and take the log. I proxy changes in global crude oil inventories by scaling US oil stocks by the ratio of OECD oil stocks to US oil stocks, both reported by the EIA. Finally, I include the 3-month ICE futures contract on Brent oil,[19] similarly deflated by CPI and log transformed, consistent with the use of Brent as the global spot price of oil.

The scope of my study is global, reflecting the global nature of oil markets. This explains my choice of variables; and the use of Brent oil instead of WTI in particular.[20] Kilian and Zhou (2020) outline the importance of using global variables in oil market modeling to avoid omitted variable bias and warn against using local US market measures, particularly the WTI price. My period of study is July 1989 to September 2020, the period during which all necessary data is available.

2.5 Empirical Results

2.5.1 Main Results

I begin my analysis by establishing whether futures demand shocks impact spot prices. The answer, evident in Figure 2.2, is yes. Specifically, in the bottom row, the futures demand shock is characterized by a positive short-run spot price response peaking at around 5%-10% between months 1 and 5 before gradually diminishing over time.

[19] Results are similar using 6-month futures instead, as shown in Appendix A, Figures A.12-A.14 and Table A.3.

[20] Refiners acquisition cost for imported crude oil from the EIA is often used as the spot price variable in oil market studies because it captures the global nature of oil markets, and data is available much earlier than the Brent Spot price. Using Brent as the spot price instead is necessary to accurately calculate the spread as it is the basis for futures contracts. The desire to study futures markets already precludes using earlier data in the analysis.

[21] The error bands show that, while there is some variation in the magnitude of the responses, a positive spot price response to excess futures market demand is evident across admissible models, even when uncertainty in the reduced form parameters is accounted for. Further, the futures demand shock is associated with a short run increase in real activity. This response implies that excess trading in futures markets feeds back into spot prices by affecting the real economy, and demand in particular, consistent with Sockin and Xiong (2015).

Next, I conduct a forecast error variance decomposition of the spot price of oil to evaluate the relative importance of the shocks, and determine how substantial futures demand shocks are as a driver of spot prices. Strikingly, Table 2.3 shows that futures demand trading is responsible for around half of the variation in the spot price of oil over the short run. It also contributes, albeit much more weakly, to long run variation in oil prices. Specifically, futures market shocks account for 65.4% of explained short run variation in oil prices, compared to the 16.8% explained by demand shocks, which are the next strongest contributor to 1 to 15 months' oil price variation. This reverses in the long run, with demand shocks accounting for 50.70% of oil price variation and supply shocks explaining 22.60%, versus 16.60% explained by futures demand shocks. Inventory demand shocks explain 9-10% across horizons. These results indicate that financial market shocks, transmitted through futures market trading, have significant effects on real oil prices, which are relatively persistent, but eventually resolve themselves and yield to supply and demand fundamentals in the long run determination of prices. These results are in stark contrast to traditional models, which do not account for the effects of financial market trading and subsequently assign virtually all oil price changes to supply and demand fundamentals.

[21]The size of the shock is standardized to one standard deviation for evaluating impulse response functions.

Finally, I examine the influence of each shock through time by conducting a historical decomposition of the spot price of oil. This entails recovering the time series of structural shocks, applying the impulse response functions to each shock, and cumulating their effect over time. Figure 2.3 shows the analysis; the magnitude and variation in the futures demand shock through time qualitatively reinforces the importance of excessive futures trading on prices.

Collectively, these results imply that excessive trading in futures markets can influence the real economy and significantly distort spot market prices, particularly in the short run.

2.5.2 Financialization and Futures Demand Shocks

I now turn my focus to the link between futures demand shocks and the rise of financialization. Given the documented rise in the popularity of index investment in commodities after 2003, a natural question is whether these futures demand shocks have become more influential over time.

To answer this question, I split the sample and estimate the SVAR separately before and after 2003. I then conduct a forecast error variance decomposition of the spot price during each period to compare the relative importance of each shock over time. To accommodate the reduced sample size, I reduce the number of lags in the model from 24 to 12. This likely truncates some dynamics and risks introducing bias. However, the comparative nature of this particular analysis serves to mitigate this, and full sample results with reduced lags provide qualitatively similar results to the longer lag model.[22]

Table 2.3 shows the results of this variance decomposition for each sub-period

[22] Using 12 months lags instead of 24 for the full sample analysis yields qualitatively similar results in Table A.1, suggesting that any bias introduced here is not severe.

and suggests that futures demand shocks have become increasingly important over time. Specifically, during the period Jan. 2003- Sept. 2020, 57.7% of short run, and 15.3% of long run variation in the oil price is explained by futures demand shocks. In contrast, before 2003, only 6.72% of short run, and 6.1% of long-run variation is due to futures demand. This is a dramatic increase in the influence of futures market trading over time, consistent with the observed rise of index investment and suggesting its rising influence on both futures and spot market prices.

2.5.3 Subperiod Results

Next, I extract what insights this model can provide about the drivers of oil prices during different historical periods throughout the sample. I utilize historical decompositions of the spot price of oil to show the evolution of oil prices surrounding each event and evaluate the cumulative percentage change in the oil price associated with each shock.

The 2008 boom and bust of oil prices

I first explore the drivers of oil prices during the 2008 global financial crisis. Figure 2.4 provides a clear picture of the dynamics during this time. There is a clear slowdown and slight reversal in the demand effect in early 2008, around the time of Bear Sterns' collapse. Meanwhile, oil prices continued to rise through to mid-2008, clearly driven in the model by futures demand shocks. Specifically, futures demands shocks are responsible for pushing spot oil prices up by almost 50% from the fall of 2007 to their peak in mid-2008, while price changes due to demand are relatively flat through to spring, before declining. These results reconcile the previously puzzling finding of spot oil prices being driven by demand shocks, despite signs of economic slowdown in late 2007 and early 2008, as pointed out by Cheng and Xiong (2014). These futures demand shocks coincide with the large inflows into commodity index funds pointed out by Cabellero, Farhi, and Gourinchas (2008), as investors reallocated out of real estate.

Futures Demand shocks during other periods

The literature so far focuses extensively on 2008; however, I find new evidence that futures demand shocks are influential during other major oil price events as well. Specifically, I find that futures demand shocks exacerbated the oil price crash of 2014. Figure 2.5 shows that, while the crash was initially triggered by a glut of supply consistent with contemporary belief and subsequent literature, the puzzling continuation in decline is primarily due to futures demand shocks. This result contrasts with explanations from previous models that attribute the extended decline, and indeed the largest part of the overall decline, to demand shocks. Similar to 2008, such demand based explanations are puzzling; both US and Chinese GDP were stable during the period, making such a historic, demand driven price drop questionable.

I also find signs of futures demand shocks contributing to oil prices during the bond market crisis of 1994. Specifically, Figure 2.6 shows futures demand driving increases in the oil price beginning in the spring of 1994 and continuing through to the end of the year, consistent with a shift in financial markets away from bonds and into other assets. Additionally, in Figure 2.7, I find signs of excess futures demand during the 1997 Asian financial crisis. Specifically, there is apparent upward pressure on oil prices around the beginning of the period, consistent with an outflow of investment from Asian equity, debt, and currency markets into other asset classes; this reverses as the period resolves. During this time, the primary downward pressure on prices comes from flow demand shocks, capturing the real economic effect of the crisis. One caution here is that, since futures demand shocks are much more important after 2003, the magnitude of the responses before 2003 may be overestimated. Nevertheless, it is interesting, but perhaps not surprising, that futures demand shocks play a more prominent role during financial turmoil, transmitting large financial market shocks to the real economy.

Exogenous events

Next, I illustrate the ability of the model to capture more general oil market behavior by presenting the estimated responses of oil prices during the set of exogenous oil market events.[23] Figure 2.8 illustrates the historical decomposition of changes in oil production during the 2002-2003 and 2011 events. The model captures the significant supply shocks of the Venezuelan oil strike in December 2002, the Invasion of Iraq in the spring of 2003, and the Libyan Civil War in February 2011, and attributes these supply shocks as the primary drivers of production changes during that time.

Figure 2.9 shows the decline in oil prices due to demand shocks starting in Jan. 2020, when Covid-19 was revealed in China. This decline continued throughout early 2020 as the pandemic worsened, with a particularly large shock throughout March and April, when the WHO declared a global pandemic and more countries began imposing restrictions. While other shocks have transitory effects, the demand shock accounts for a sizeable permanent drop in the oil price and is the primary determinant of oil prices during the period.

[23]While these events are used to restrict the model, the model manages to capture them quite well even when the restriction is relaxed, while other key results remain similar. Specifically, Figures A.1-A.5 in Appendix A show similar IRFs and historical decompositions, while Table A.2 shows a qualitatively similar variance decomposition. Thus, narrative restrictions in the model do not drive the results but instead serve to validate them.

Other periods of interest

Finally, I examine oil price behavior around selected recent events that may interest the reader and present the results in Appendix A. When analyzing oil market events like this, it is essential to remember the large size of global oil markets and the resulting difficulty in directly attributing any observed behavior if an event is small, no matter how visible such an event might be in the news. Nevertheless, starting with Figure A.6, I find the assassination of Jamal Khashoggi on October 2nd, 2018, to be associated primarily with a small demand shock to oil prices, and a very weak futures demand shock. The event occurs during the beginning of the month, so much of the volatility surrounding it is likely resolved within the month. Next, I show in Figure A.7, that the September 14th, 2019 Abqaiq-Khurais drone attacks on Saudi oil facilities is captured as a supply shock, temporarily decreasing oil production during the month. Finally, sharing Figure 2.9 with the onset of Covid-19, I capture the Saudi-Russian oil standoff, with a large supply shock putting downwards pressure on the price of oil when Saudi Arabia flooded the market with oil in response to Russian refusal to cut production. The supply shock reverses in April after Saudi Arabia and Russia reached an agreement.[24] [25]

[24] Another notable oil market incident during this period is the negative price of the front month WTI futures contract which dropped from $17.85 to -$37.63 on April 20th, 2020 due to limited US storage and a liquidity squeeze at expiry. This event occurred and was resolved in a very short period, and only on the front month WTI contract. This event does not show up at a monthly level and is unrelated to Brent oil, which was stable during this period. Brent Oil also does not share the same storage issues as WTI, and is cash-settled.

[25] The assassination of Qasem Soleimani on January 3rd, 2020, also falls within this period, but its contribution, if any, to global oil prices is impossible to disentangle from the early onset of Covid-19 in China, also occurring in January of 2020.

2.5.4 Comparison with previous models

In this section, I outline the differences between my results, obtained with the inclusion of a futures demand shock, and the results obtained by a similar model without it.[26] To enhance comparability, I replicate the results of Kilian and Murphy (2014) in two ways: the first covers their original period using their original data, and the second covers my sample period and uses my data, but follows their restrictions to obtain their original set of shocks. I present these results in Appendix A. Compared with the historical decomposition in Figures A.8 (original) and A.10 (current sample), the first novelty of my results is that, when I include a futures demand shock in the model, it supplants the importance of demand shocks in the short run. This is not surprising because they are observationally equivalent to and could be captured by traditional demand shocks in the absence of identifying restrictions for the futures market shock.

The second novelty of my results is the ability to credibly explain the 2008 run up and collapse of oil prices. Models that do not include a futures demand shock fail to explain oil price behavior during 2008, as shown in Figure A.9 (original) and Figure A.11 (current sample). While these models conclude that demand is largely responsible for a prolonged increase in oil prices after 2003, the estimated contribution of demand during early 2008 is relatively flat, and the oil price boom and bust during this period remains a puzzle. In contrast, I capture credible pressure from futures market trading, consistent with the observed shift of flows into commodity index funds during that time, and a flat and reversing trend of demand, consistent with observed signs of economic reversal.

[26]Kilian and Murphy (2014) name their inventory demand shock a speculative demand shock, not to be confused with my futures demand shock. In later papers, the authors call it an inventory demand shock.

2.5.5 Alternative Explanations

One might worry that my futures demand shock captures something other than trading pressure from financial markets. This concern is especially poignant since shocks are defined solely by the restrictions imposed on the model and the economic theory behind those restrictions. Here, I consider alternative explanations that potentially fit the restrictions outlined in Table 2.1, and rule out alternative mechanisms.

Interest Rates and Storage Cost

First, I evaluate interest rate changes as a potential driver of the futures demand shock. One of the critical identifying assumptions behind the shock is an increase in the futures-spot spread. Aside from changing fundamentals outlined previously, the futures-spot spread also increases whenever interest rates increase, reflecting the changing value of future cash flows. While changing interest rates may match the spread restrictions of the futures market shock, whether rate changes should lead to positive responses for production, real activity, and spot prices is ambiguous at best. As increases are usually implemented to cool an economy, it is more likely that the contemporaneous impact on the economy is negative, even if they are generally implemented in response to economic strength. Nevertheless, I show in Table 2.4 that changes in both Fed fund futures and Euribor futures are uncorrelated with the futures demand shock. Furthermore, during the most dramatic period of oil price rises in early 2008, interest rates were actually declining, providing additional evidence against interest rate changes driving the shock.[27]

Next, while increasing storage costs could also increase the futures-spot spread, the remaining responses of the futures demand shock are likely incompatible. A shock to

[27] Additionally, interest rates were essentially flat from 2009-2015, but the futures demand shock in Figure 2.3 is not.

storage cost is more credibly captured as a negative inventory demand shock, directly increasing the cost and thus decreasing the demand to store inventories. Testing this, I find in Table 2.4 that the LOOP Futures Contract, a proxy of oil storage cost in the US, is unrelated to the futures demand shock.

Information Diffusion

The next question I examine is whether my shock is simply capturing faster information diffusion in futures market prices relative to spot market prices, which could cause futures prices to lead spot prices, temporarily increase the spread on positive news, and decrease the spread on negative news.

There are two clear reasons to discount this explanation. First, recent studies do not support futures markets consistently leading spot markets.[28] Second, studies of information diffusion to futures and spot markets generally focus on the possibility of daily or intraday differences, not differences at the monthly level of my analysis. For example, Silvapulle and Moosa (1999) use daily data to show that spot and futures prices react simultaneously to new information but point out that using daily data may conceal any relationship at higher frequencies. They argue that intraday data, when available, may be better suited to finding speed differences. Even early research favoring a strong lead in futures markets, by Schwarz and Szakmary (1994), finds that the relationship weakens beyond a day, and disappears at a monthly level. These arguments suggest that a consistent lead from futures markets, driven by speed differences in information diffusion, is not a likely driver of my futures demand shock at the monthly level.

[28] Among the more recent papers finding that both markets play an important and temporally varying role in price discovery are Silvapulle and Moosa (1999), Kaufmann and Ullman (2009), Figuerola-Ferretti, and Gonzalo (2010), Peri, Baldi, and Vandone (2013), Dolatabadi, Nielsen, and Xu (2014).

In Figure 2.10, I confirm that there is no significant difference in the speed of diffusion of news into spot and futures prices, even at a daily level, by looking at the reaction of both prices to surprise announcements of changes in the fed funds rate by the FOMC. This alleviates concerns that my results are driven by increases in the spread upon impact due to information being diffused differently across markets.

Hedging Pressure

I have so far focused on investor demand driving the futures-spot spread away from the fundamental relation in equation (2.3). I now evaluate the other side of the futures market, and ask whether futures demand shocks are driven by changes in hedger demand instead of investor demand, both of which can affect the spread. Specifically, I consider whether an increase in the futures-spot spread associated with my futures demand shock could arise because hedgers decrease their demand for short positions in futures contracts, which would provide similar upward pressure on futures prices if futures demand stays the same.[29]

To test this, I check the correlation of the futures demand shock with changes in the net futures position of producers and merchants, obtained from the CFTC Commitment of Traders (COT) report.[30] In Table 2.4, I show no relationship between changes in the net futures position of producers and merchants and the futures demand shock, indicating that hedging activities do not drive the shock.

Theoretically, hedging demand is most likely to be captured as an inventory demand shock. Acharya, Lochstoer, and Ramadorai (2013) show that producers hedge their current inventories and part of their future production. Hedging demand should then be related to changing demand for inventories or changing expectations of fu-

[29] Note, the academic and industry narratives after 2003 have focused on explaining clear increases in demand for long positions, not decreases in short positions, so this exercise may be academic.

[30] Results are the same when using the shorter series of Disagreggated COT data from the CFTC.

ture production. Kilian and Murphy (2014) show the latter is incorporated into oil prices through inventory demand as well, since expected future tightness in supply implies a higher future spot price, making storage profitable.[31] This line of reasoning suggests that my inventory demand shock should be able to capture hedging demand. Therefore, hedging demand is theoretically unlikely to drive my future shock.

Stepping back, none of the competing explanations which could explain the spread response of the futures demand shock are particularly credible. For the most part, they are not consistent with the other economic responses associated with the shock, and empirical tests do not suggest that they are related.

2.6 Conclusion

I find consistent evidence that excess demand for commodity futures contracts is an important determinant of spot market oil prices and has increased in importance over time alongside increased financialization. My analysis shows that excess demand strongly influences short-run oil prices but yields in the long run to traditional economic mechanisms, particularly demand. The new evidence I provide demonstrates that excessive futures market demand was not only a major driving force behind the 2008 rise and fall of oil prices, but also during the dramatic fall of oil prices in 2014 and, potentially with weaker effect, during earlier financial crises in 1994 and 1997. These results also reconcile previously puzzling results, which identified demand as a driver of stronger oil prices in early 2008, despite a weakening economy, and as one driver of weaker prices in late 2014, despite a relatively stable global economy.

[31] Both hedging demand and demand for inventories have been linked to economic uncertainty. For example, Hamilton and Wu (2014) show that hedging demand is higher during times of economic uncertainty. Likewise, Alquist and Kilian (2010) observe an increase in demand for inventories during times of uncertainty. This further supports an inventory demand response capturing hedging demand.

Finally, my study provides empirical support for the theory of Sockin and Xiong (2015), whereby futures market prices feedback to the real economy as signals in managerial decision making, with potentially distortive effects. This indicates the importance of accounting for financial market shocks in oil market models, and models of the real economy in general. It also points to the need for caution in interpreting the informational content of financial prices when making economic decisions.

Figure 2.1: Real Price of Brent Oil

Displays the real spot price of brent oil reported by the Energy Information Administration (EIA) between Jul. 1991 and Sept. 2020. Price in US Dollars. The horizontal line coincides with the year 2003 and marks a common start point of the period of financialization in commodity markets defined by the popularity of commodity futures market investment.

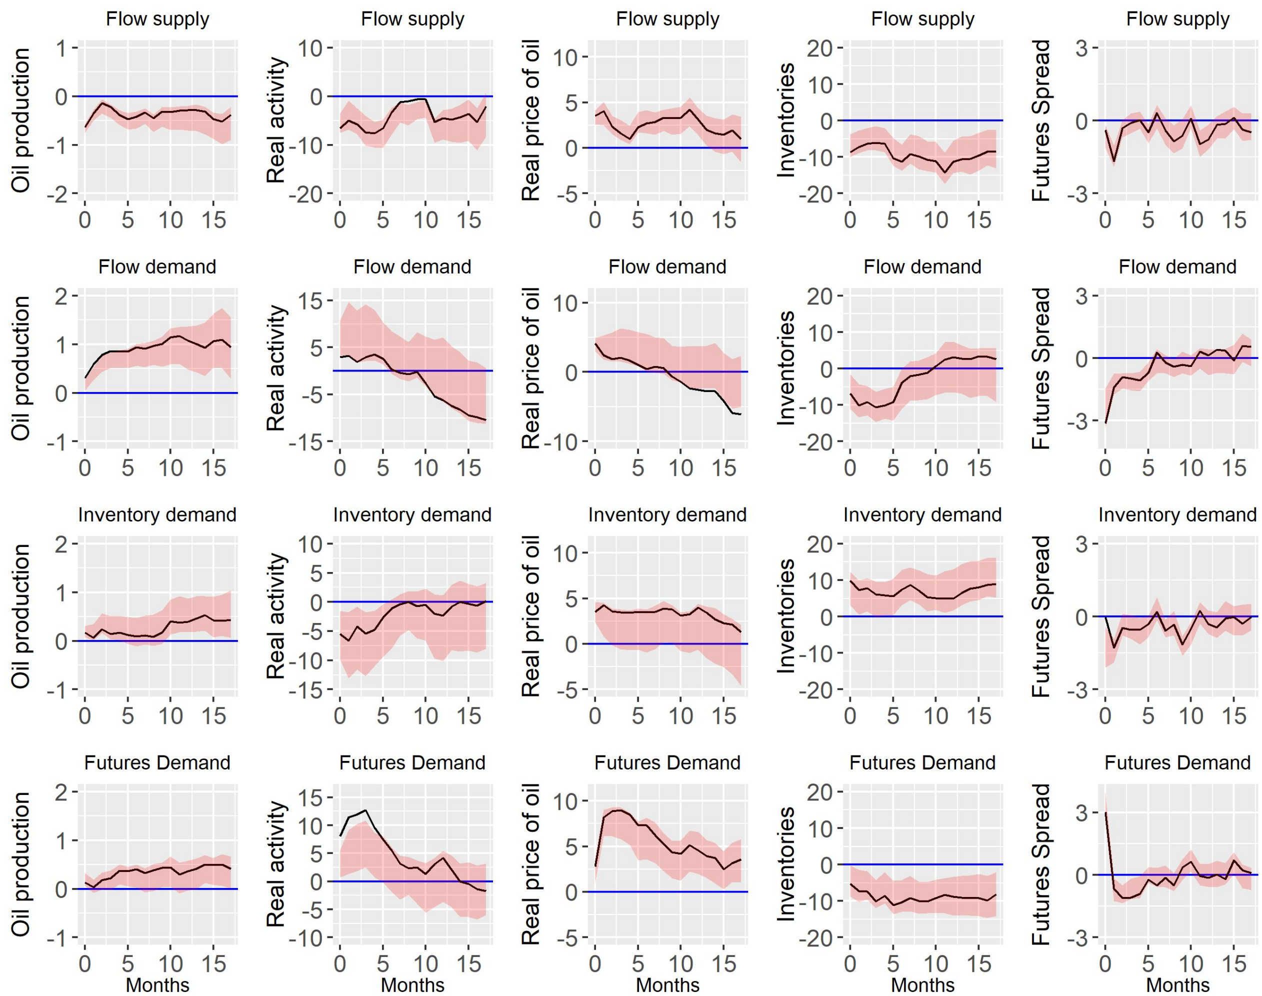

Figure 2.2: Structural IRFs

Structural Impulse Response Functions showing the response of each variable to a one standard deviation innovation to each structural shock. Responses are the cumulative % change for production, real activity, and the spot price, and cumulative level change for inventories. The Spread response is the difference in the futures and spot responses. The red band illustrates the 68% error band from the posterior distribution of the IRF's. Obtained following the methodology described in section 3, and in Appendix B.

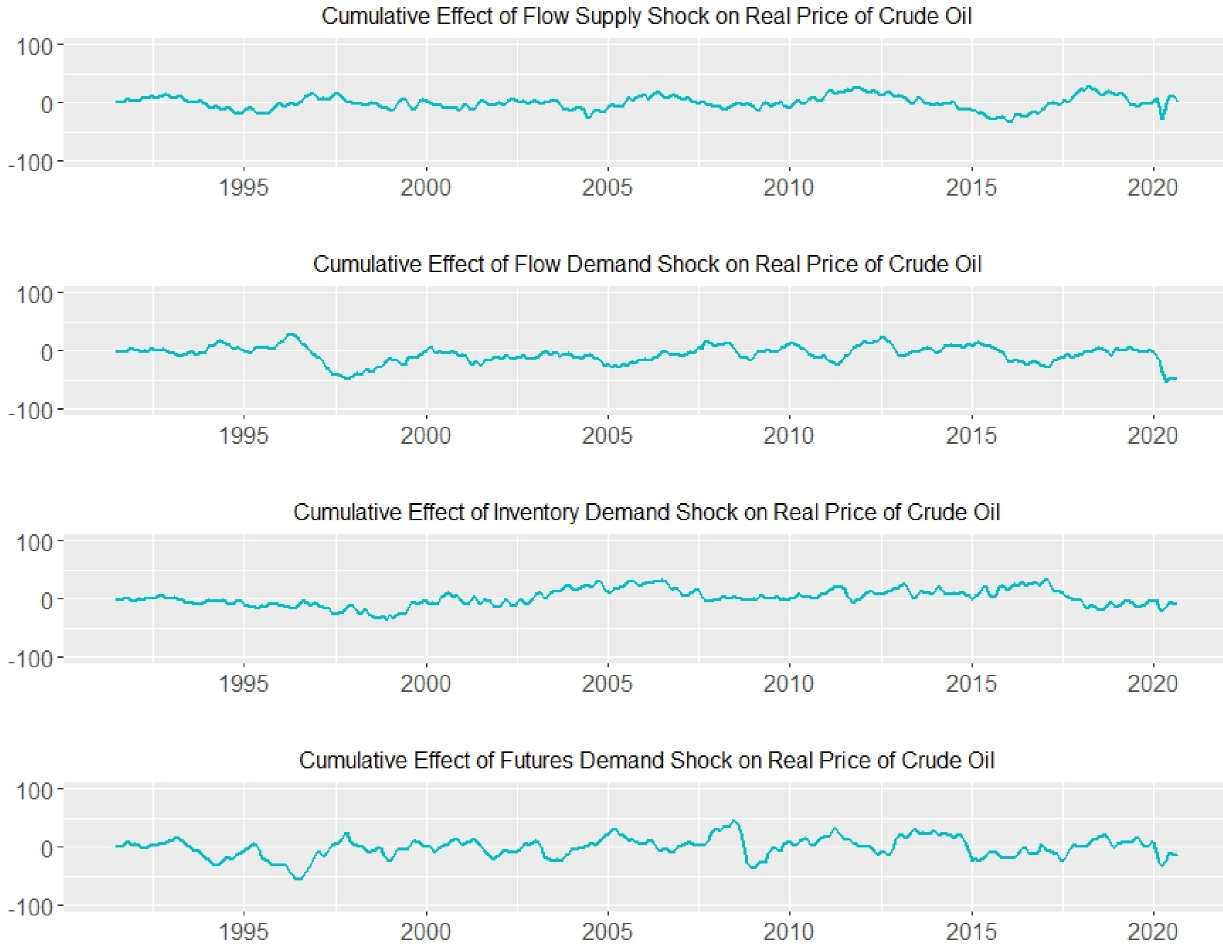

Figure 2.3: Historical Decomposition

Historical decomposition of the real spot price of brent oil from July 1991 to September 2020 showing the cumulative percentage change in spot price due to flow supply, flow demand, inventory demand, and futures demand shocks, respectively.

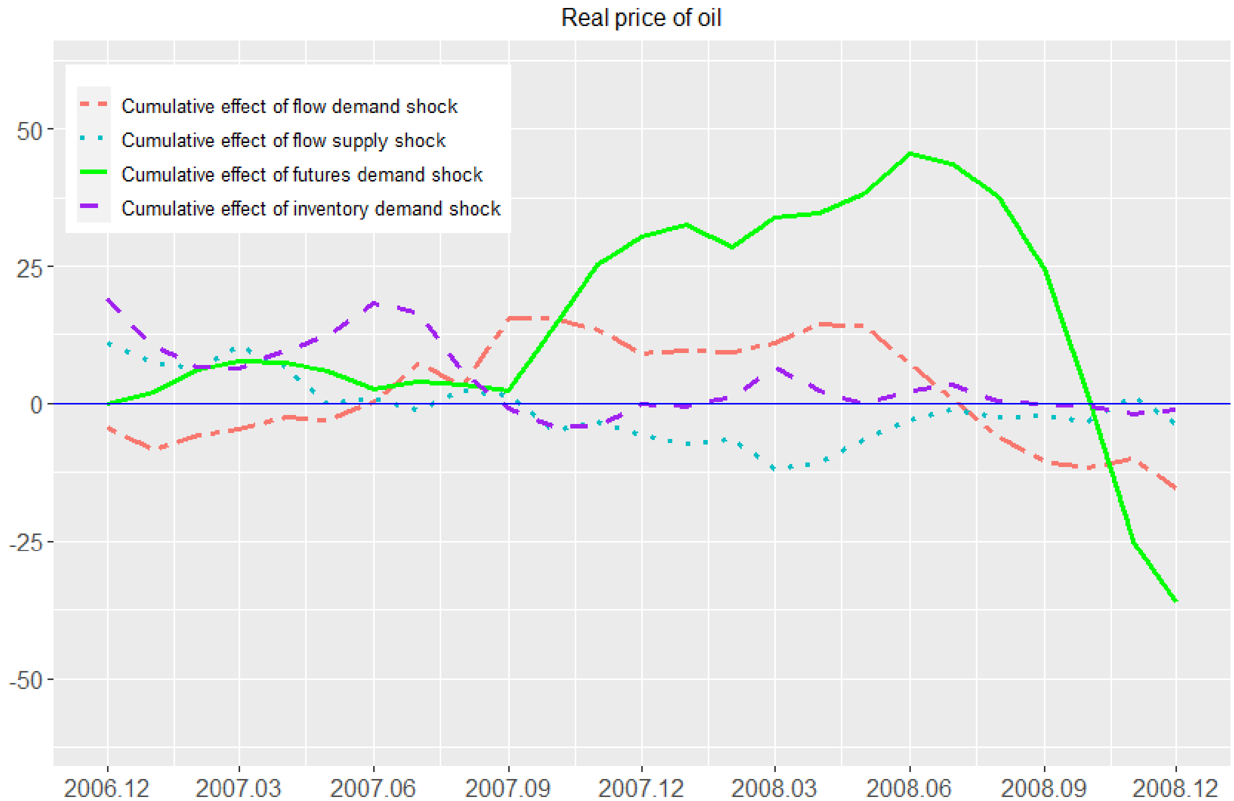

Figure 2.4: Historical Decomposition-2008 Financial Crisis

Historical decomposition of the real spot price of brent oil showing cumulative percentage change in spot price due to flow demand and futures demand shocks, respectively, between Jan. 2007 and Jan. 2009.

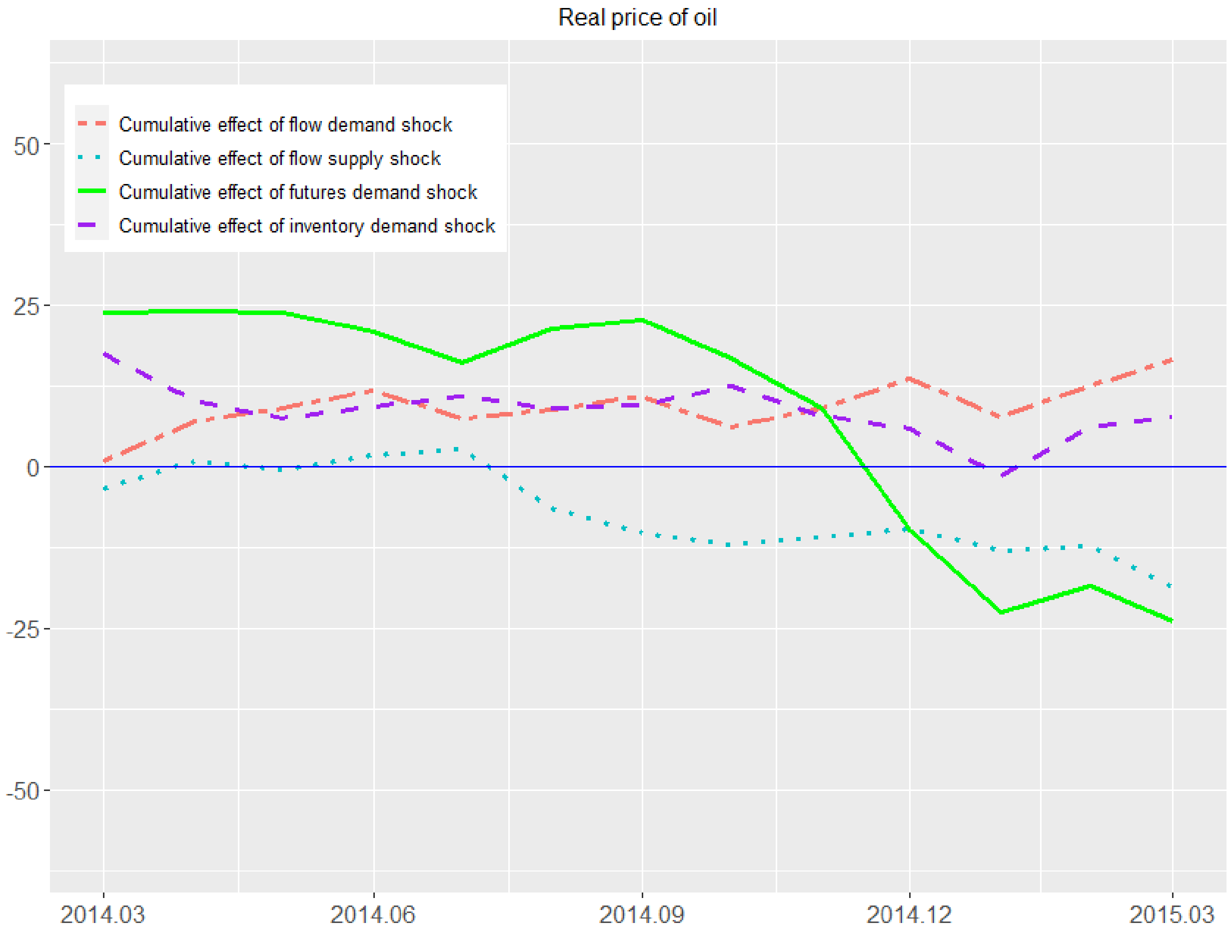

Figure 2.5: Historical decomposition 2014

Historical decomposition of the real spot price of brent oil showing cumulative percentage change in spot price due to each shock during 2014.

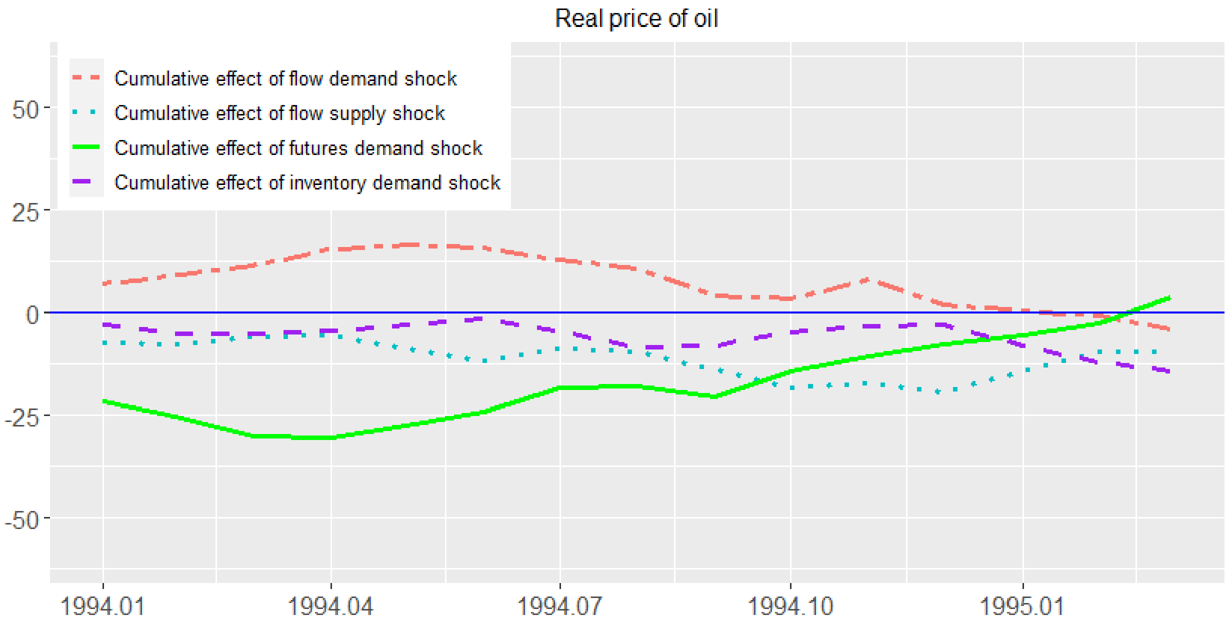

Figure 2.6: Historical decomposition 1994

Historical decomposition of the real spot price of brent oil showing cumulative percentage change in spot price due to each shock during the Bond Market Crisis of 1994.

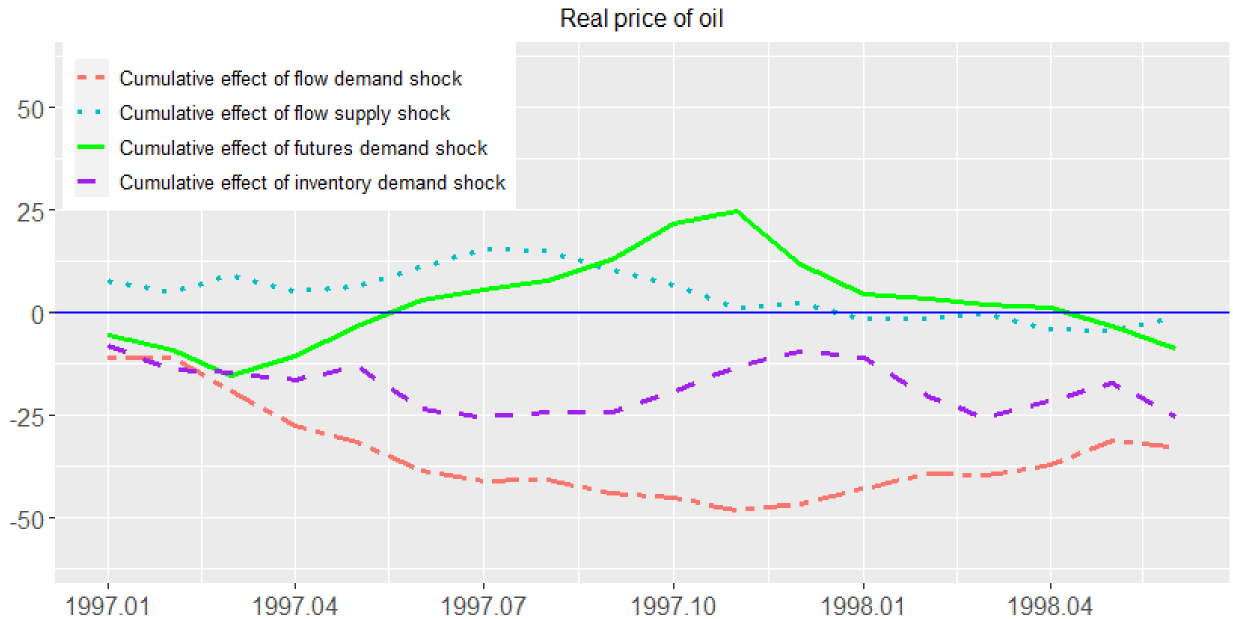

Figure 2.7: Historical decomposition 1997-1998

Historical decomposition of the real spot price of brent oil showing cumulative percentage change in spot price due to each shock during the Asian Financial Crisis of 1997.

Figure 2.8: Historical decomposition of exogenous supply shocks

Historical decomposition of global oil production, showing cumulative percentage change in global production due to each shock during the period of the Venezualan Oil Strike (Dec. 2002), Invasion of Iraq (Mar.-Apr. 2003), and Libyan Civil War (Feb. 2011).

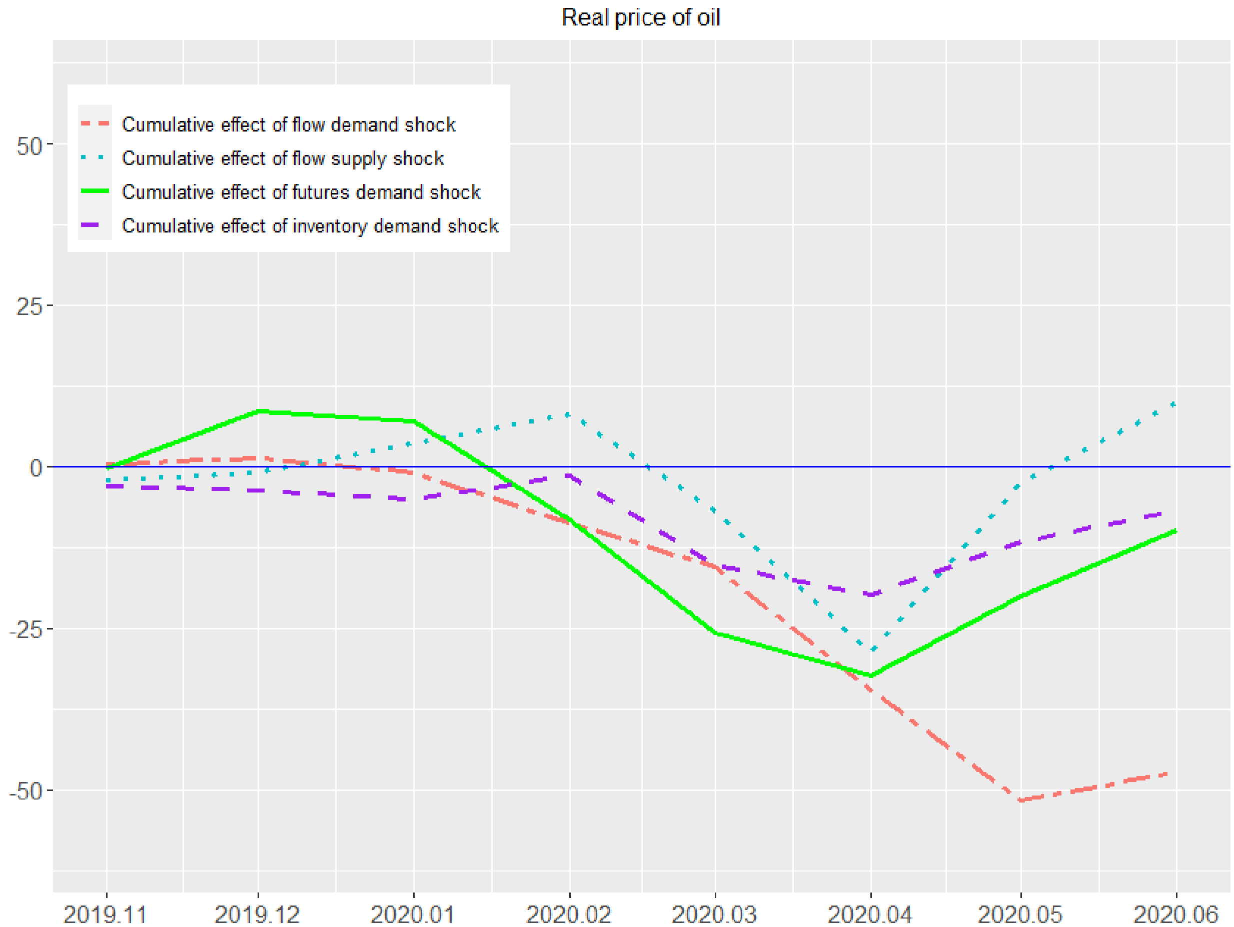

Figure 2.9: Historical decomposition Covid-19 Outbreak

Historical decomposition of the real spot price of brent oil showing cumulative percentage change in spot price due to each shock during the onset of the Global Covid-19 Pandemic and the Saudi-Russian Oil price war.

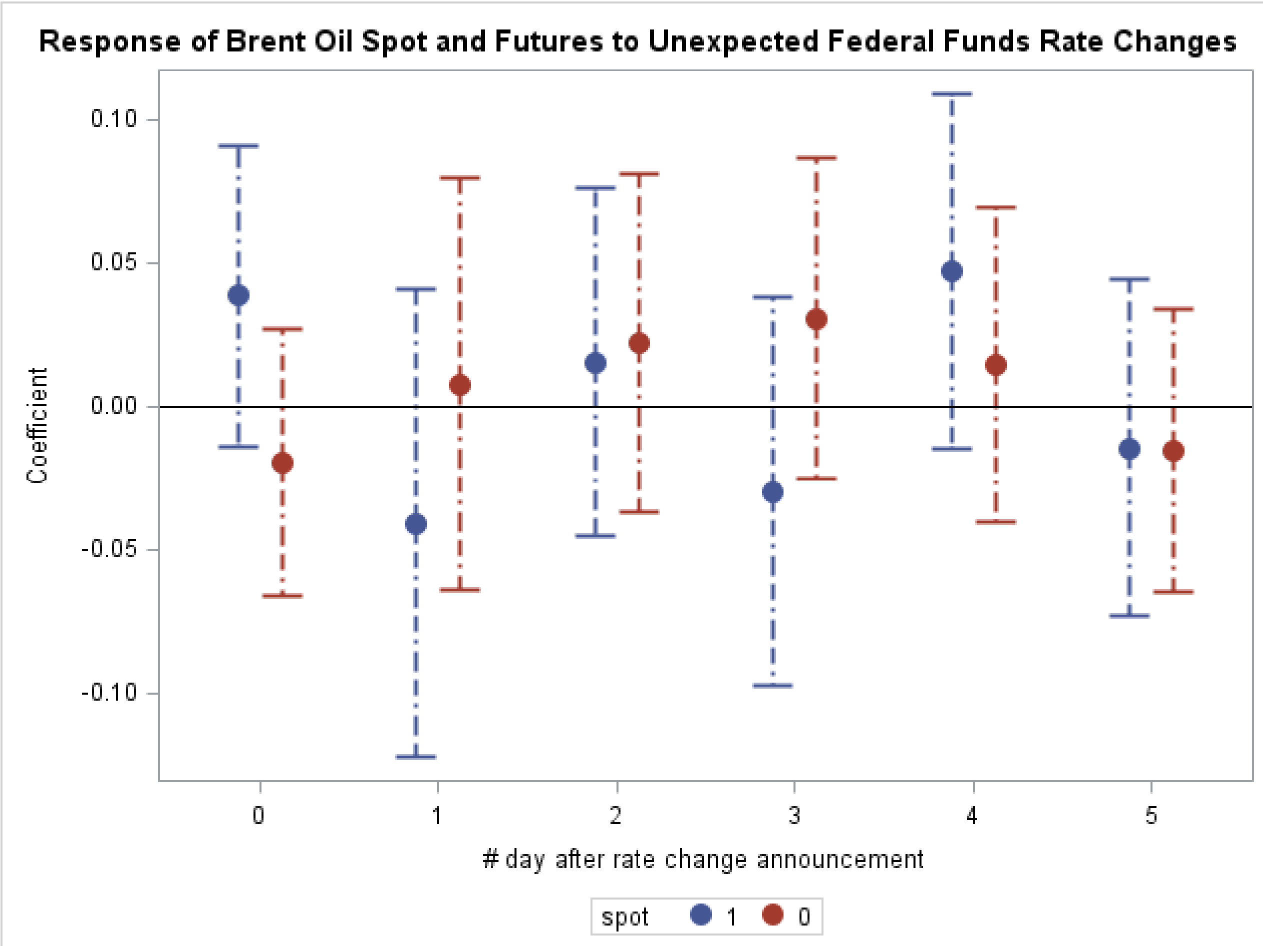

Figure 2.10: Response to unexpected Fed Fund rate changes

Displays point estimates and 95% confidence intervals for the price reaction of Brent spot and 3mo Brent futures to surprise FOMC announcements about federal funds rate changes on the announcement date and over the post-announcement window. Calculation of surprise FOMC announcements follows Bernanke and Kuttner (2005). Sample period for announcements is June 1993-April 2020.

Table 2.1: Identifying Sign Restrictions

Impact sign restrictions on the structural impulse responses of each structural shock. Following standard practice in the literature, all shocks are normalized to have a positive effect on the price of oil. Hence, the supply shock of interest is a negative supply shock.

	Negative Flow Supply	Flow Demand	Inventory Demand	Futures Demand
Production	-	+	+	+
Real Activity	-	+	-	+
Real Price Oil	+	+	+	+
Inventory			+	
Futures Price	+	+	+	+
Spread	-	-	-	+

Table 2.2: Forecast Error Variance Decomposition of the Real Spot Price of Oil

Variance decomposition of the real spot price of oil reflecting the percentage of variation at each monthly horizon attributable to each respective shock. Estimated using main model with restrictions from Table 2.1 with 24 lags estimated over July 1991-Sept 2020.

Horizon	Supply	Demand	Inv.	Spec.	Resid.
1	19.23	26.05	20.03	30.62	4.07
2	16.25	21.72	17.44	41.16	3.44
3	12.62	16.22	12.88	55.33	2.95
4	11.13	12.57	9.96	62.67	3.68
5	11.17	10.35	8.27	65.58	4.63
6	9.82	9.18	7.19	65.88	7.93
7	8.75	8.68	6.37	65.97	10.22
8	8.25	8.24	6.28	66.86	10.37
9	8.01	8.24	6.47	66.95	10.34
10	7.99	8.61	6.78	65.85	10.77
11	8.28	8.84	6.85	64.19	11.84
12	8.5	9.85	6.68	61.96	13.01
13	8.03	11.43	7.17	59.1	14.27
14	7.55	12.33	7.49	56.15	16.48
15	7.11	13.46	7.25	52.6	19.58
% of Explained	8.84%	16.79%	9.01%	65.4%	
600	9.21	20.66	4.1	6.78	59.24
% of Explained	22.60%	50.70%	10.10%	16.60%	

Table 2.3: Subperiod Forecast Error Variance Decomposition of the Real Spot Price of Oil

Variance decomposition of the real spot price of oil reflecting the percentage of variation at each monthly horizon attributable to each respective shock. Estimated using main model with restrictions from Table 2.1 with 12 lags estimated over the respective subperiods.

		Jul. 1990-Dec. 2002			
Horizon	Supply	Demand	Inven.	Spec.	Resid.
1	1.22	11.60	50.65	21.24	15.29
2	1.28	13.87	43.87	26.85	14.12
3	0.93	29.93	32.17	22.74	14.24
4	1.19	35.35	29.53	21.07	12.86
5	0.94	43.67	26.50	18.72	10.17
6	0.75	54.19	23.54	13.87	7.65
7	0.88	58.68	22.16	11.26	7.02
8	1.42	56.03	24.50	9.73	8.32
9	2.30	49.81	27.78	8.19	11.93
10	3.26	46.21	27.84	7.17	15.52
11	3.34	43.33	29.28	6.51	17.54
12	3.19	41.12	29.47	6.03	20.18
13	3.05	41.60	28.82	5.75	20.78
14	3.03	42.01	28.48	5.50	20.97
15	3.22	41.78	28.43	5.29	21.29
% of Explained	4.09%	53.07%	36.11%	6.72%	
600	6.69	40.92	27.07	4.82	20.50
% of Explained	8.4%	51.5%	34.1%	6.1%	
		Jan. 2003-Sept. 2020			
Horizon	Supply	Demand	Inven.	Spec.	Resid.
1	19.19	22.44	15.41	36.28	6.68
2	15.37	18.85	12.43	47.56	5.79
3	14.80	12.68	8.55	60.15	3.82
4	17.10	9.14	7.29	63.00	3.47
5	17.77	7.50	7.71	62.91	4.11
6	15.63	6.42	8.75	61.55	7.65
7	14.99	6.00	9.59	60.11	9.32
8	14.45	5.64	9.83	61.05	9.04
9	13.88	5.44	9.96	62.01	8.72
10	13.60	5.77	10.22	62.14	8.27
11	12.94	5.99	10.60	62.02	8.46
12	11.96	7.77	11.39	60.11	8.78
13	11.75	10.14	11.45	58.09	8.57
14	12.09	12.91	11.20	55.38	8.42
15	12.70	15.17	10.92	53.00	8.20
% of Explained	13.8%	16.5%	11.9%	57.7%	
600	46.87	10.94	10.24	12.27	19.68
% of Explained	58.4%	13.6%	12.7%	15.3%	

Table 2.4: Futures Demand Shock Correlation with other Variables

This table presents pearson and kendall correlations of the estimated futures demand shocks and the 30-day fed fund futures price (July 1991-Sept. 2020), Euribor futures price (Dec. 2012-Sept. 2020), LOOP Storage Futures price (Nov. 2015-Sept. 2020), and DCOT net flow of producers/merchants into WTI futures contracts(Jan. 2012-Sept. 2020). Also reported is the t-stat under the null hypothesis that the variables are uncorrelated.

	Pearson Correlations			
	Fed Fund	Euribor	LOOP	Prod/Merc
Spec. Futures Shock	0.006	-0.032	0.031	0.051
	(0.112)	(-0.312)	(0.233)	(0.523)
	Kendall Correlations			
	Fed Fund	Euribor	LOOP	Prod/Merc
Spec. Futures Shock	0.043	-0.041	0.038	0.035
	(1.202)	(-0.579)	(0.411)	(0.523)

Chapter 3

Comovement and S&P 500 Membership

3.1 Introduction

The theory of Barberis, Shleifer, and Wurgler (2005) predicts that membership in a popular group of assets, such as the S&P 500, can expose firms to correlated exogenous demands from certain classes of investors that primarily inhabit the index, and introduce a non-fundamental group factor to asset returns. They support their theory with evidence of an increase in comovement with the S&P 500 (the index) in a short period after firms are added to the group. However, Chen, Singal, and Whitelaw (2016) show that this "excess" comovement can be explained as a manifestation of momentum in stocks at the time of inclusion. Consequently, these studies provide insight on the effect of inclusion on return dynamics in the short run, but they do not address the important long run effect that membership entails for comovement.

The study of comovement is of interest to researchers because it has direct implications for asset pricing. For prices to be informationally efficient, they should reflect all available information about a firm's expected future discounted cash flows, with returns and comovement reflecting changes in these fundamental expectations. However, the potential existence of excess comovement directly violates this efficiency, which stands as a pillar of traditional asset pricing theories and is essential to the efficient allocation of capital within the economy. This paper aims to identify whether such excess comovement exists by estimating the average effect of S&P 500 membership on comovement over the long term.

To identify the causal effect of S&P 500 membership on comovement, I utilize a fuzzy regression discontinuity design to exploit locally exogenous variation in the probability of inclusion arising from meeting the size criteria for index membership in a narrow window around the inclusion threshold. Intuitively, firms that are very close in size should not be systematically different given the stochastic nature of firm size, yet meeting the size requirement leads to a discrete change in the probability of inclusion, which can be used to measure the treatment effect of such inclusion.

The advantage of using my design is the ability to recover long term estimates by using all firm-month observations in the sampling period. In contrast, short term estimates generated using event studies are limited to measuring effects on a sub-sample of firms during a short window around inclusion. Such estimates may be unable to identify a group component in returns driving excess comovement if there are temporary confounding effects on the return process while the stock integrates into the index.[1] Time-based identification strategies may also be challenged by selection on characteristics associated with time variation in betas, as Chen et al. (2016) document with momentum. My fuzzy RD design allows estimation of the long term effect of membership, unconfounded by potential short run dynamics and time-based identification challenges.

I find that members of the S&P 500 experience an increase in comovement with the index, as measured by beta, of around 0.20. This indicates a significant group component in the return process for S&P 500 stocks which could create a disconnect between prices and fundamental values. I find no evidence of a short term increase in beta when I limit my sample to firms within one year of inclusion or deletion, consistent with previous short term studies, pointing to potential confounding effects during the short term integration process and the inherent differences in the long

[1] For example, persistent exogenous demand to buy the stock as part of the rebalancing process could introduce a temporary component to returns that is uncorrelated with firm fundamentals and overall group demand. Madhavan (2003) outlines the price pressure that rebalancing can have on stock prices after index inclusion.

term and short term estimates.

To ensure the robustness of my results, I estimate the effect using multiple window sizes, control specifications, and subsamples. Results are similar in all cases, with only minor differences in magnitudes. I specifically rule out a confounding effect from faster information diffusion due to liquidity gains from membership, both by controlling for liquidity differences and directly evaluating the reaction of index and non-index stock prices around surprise FOMC announcements. Finally, I show that the probability of inclusion that I isolate appears to be exogenous and is unrelated to various observable firm characteristics around the size threshold.

The rest of the paper is organized as follows: Section 2 reviews relevant literature and compares my findings with previous results, section 3 describes the data, section 4 provides details on the fuzzy RDD methodology I use, section 5 presents and discusses results and section 6 concludes.

3.2 Literature Review

The belief that index membership can generate returns unrelated to firm fundamentals can be traced back to the work of Shleifer (1987) and Jain (1986), which document an effect of index inclusion on prices even though inclusion should contain no fundamental information about a firm.[2] Several theories have arisen to explain such potential asset class effects, including that of Barberis et al. (2005), which claims that habitat investors (who limit attention to a certain set of stocks) and category investors (who view the index as a single entity to trade) introduce correlated demand. Reinforcing this, Harford and Kaul (2005) link correlated order flow to strong common effects in the returns of S&P 500 stocks. Despite these theories and evident patterns

[2] More recent evidence is given by Kaul, Mehrotra, and Morck (2000) and Chen, Noronha, and Singal (2004).

of comovement in asset returns, direct evidence of long term excess comovement is limited. [3]

My study is most comparable to that of Chen et al. (2016). Using a sample of firms added to the index from 1976 to 2012, they measure the simple change in a firm's beta with the index in the periods before and after inclusion. While they show that comovement increases on inclusion, they also find that a control sample matched on momentum experiences a similar increase, and conclude that increases in beta at inclusion are simply a manifestation of a relationship between momentum and increasing beta.

In contrast, my finding of significant excess comovement is generated using the previously described fuzzy RD design. This makes my sample significantly different, as it uses all S&P 1500 firm-month observations, not just those around inclusion, from 1995-2017. To directly compare time periods, I study a sample from 2001-2012 which overlaps one of their sub periods, and my results hold, indicating that study period differences do not drive the difference in results. When I limit my sample to firms within one year of addition or deletion from the index to more closely mimic the short term nature of their sample, the difference in our results is reconciled, pointing to short term heterogeneity and potential problems using a limited sample to estimate long term membership effects, as previously outlined.

My results also contribute to a larger debate on whether increased financialization, in general, leads to frictions that can negatively impact the efficient pricing of assets. One notable example is the literature surrounding the effects of indexing in commodity markets.[4] The creation of commodity indexes, and the rising popularity

[3]See Pindyck and Rotemberg (1990) and Tang and Xiong (2012) for signs of increasing correlations in grouped commodities and Pindyck and Rotemberg (1993), and Froot and Dabora (1999) for signs of increasing correlations amongst grouped equities.

[4]Some of the important papers in this debate include Hamilton (2009), Irwin and Sanders (2012), Sanders and Irwin (2010), Tang and Xiong (2012), Singleton (2013), Kilian and Murphy (2014),

of commodity index funds to track them, has been accused of creating frictions resulting in speculative bubbles in important commodities in 2007-2008. Speculation, arising after this precipitous price drop, blames the rise of commodity index funds for exposing commodities to large correlated exogenous demand from category investors. Sockin and Xiong (2015) build a model illustrating how such deviations from fundamental value can arise, particularly at times of high economic uncertainty. Various methods have been used in the literature to empirically test this and fundamental supply/demand explanations, with conflicting results. By studying comovement effects in the S&P, this paper sheds light on potentially more general consequences of financialization in other markets.

3.3 Data

To conduct my analysis, I construct monthly index constituency data from the Compustat S&P index file. This data includes all time periods for all firms included in an S&P index since 1964, along with the specific index identifier. I limit my sample to S&P 1500 stocks, as they form the pool of stocks that meet the specific minimum liquidity and ownership characteristics required for S&P 500 inclusion. For example, dual class, foreign headquartered firms, low liquidity firms, or firms with concentrated or otherwise unusual ownership structures are ineligible. Specific requirements are listed on the S&P website. The S&P 1500 has only existed since the introduction of the S&P 600 small cap index in 1994. Therefore, the period I study is January 1995 to Dec 2017.

Security level data and industries are merged from the monthly CRSP file to make a sample of firm-month observations. Comovement calculations rely on index return

Henderson, Pearson, and Wang (2014), and Sockin and Xiong (2015). There are many more.

data from CRSP daily index files, and I obtain daily stock returns and volume data from the CRSP daily and monthly stock files. To measure monthly exposures of stocks to Fama-French factors (Fama and French, 2015) and Carhart's momentum factor (Carhart, 1997), I use data from Ken French's data library and WRDS, respectively. I retrieve recession data from NBER and classify industries according to the standard Fama-French 12 industry definitions available from Ken French's data library.

The measure of comovement I use follows the standard beta calculation used in the literature:

$$Return_{id} = \omega + \beta_{im} Index_{kd} + \epsilon_{id} \tag{3.1}$$

This is the regression of daily excess returns of stock i on the daily excess returns for index k for each month m, estimated monthly using a rolling 12 month regression period. I capture comovement as the parameter β_{im}, and this β_{im} is the month m beta for firm i with index k, which I use as the monthly comovement outcome variable. The index of interest is the S&P 500 due to its unique exposure to large classes of investors. In all cases, I remove the contribution of the individual stock being regressed to the index return to avoid spurious results, though, given the size of the firms around the discontinuity and the fact that the index is value-weighted, this makes no practical difference to the comovement measure or later results.

Table 3.1 shows descriptive statistics for the full sample. In the entire sample of firms, a few differences are evident. As expected, S&P 500 firms are on average larger in scale. They also have lower lagged returns, consistent with generally lower levels of risk for larger firms. Non-S&P 500 firms in the full sample are also less liquid, as indicated by the Amihud (2002) measure of illiquidity. While this is a price impact measure, Chordia, Goyal, Sadka, Sadka, and Shivakumar (2009) show it has a strong link to speed of information diffusion in their study of the post-earnings announcement drift. S&P 500 firms also have lower direct trading costs as measured using effective tick, calculated following Goyenko, Holden, and Trzcinka (2009), which is a monthly proxy for effective spread. Factor exposures show slight differences, with the SMB exposure of non-SP500 firms predictably larger.

3.4 Fuzzy Regression Discontinuity Design

I conduct my analysis using a fuzzy regression discontinuity design. This quasi-experimental design relies on the sample being similar on both sides of a given treatment threshold, effectively making treatment randomly assigned in a window around it. To illustrate the methodology, consider the following regression on the full sample of firms in the S&P 1500.

$$Y_{it} = \alpha + \tau SP500_{it} + \epsilon_{it} \tag{3.2}$$

Here, τ captures the inclusion in the S&P 500 on comovement or another outcome variable of interest. Such a specification is plagued by endogeneity, as the full sample of firms included in the S&P 500 is likely much different than firms not in the index. To the extent that firm characteristics differ for the two groups, and those differing characteristics are associated with comovement, such a specification has the risk of biasing the estimate of τ. This is because inclusion would not be independent of ϵ_{it}, violating the condition for unbiased OLS estimators.

The challenge is to ensure that $SP500_{it}$ and ϵ_{it} are independent. Previous studies attempt to solve this by comparing the same firm within a very narrow window around the time treatment occurs. In contrast, the RDD design I use in this paper exploits the differences around the size-based inclusion point, allowing my sample to include all months in the time period for all firms. If the window is narrow enough, and if there is sufficient randomness in the variable determining the assignment to the treatment group- in this case, size rank- then treatment is likely due to these random shocks to the firm's size around the discontinuity, rather than to systematic differences between the groups.[5] This similarity of firms around the inclusion point can be exploited by limiting the sample window in regression (5) to a sufficiently narrow threshold.

Further refinement of the specification can be achieved by controlling for any remaining differences in the assignment variable around the threshold, allowing a better approximation of the conditional mean at the discontinuity on each side of the threshold. The standard way to do this is to specify an appropriate polynomial spline for the assignment variable around the discontinuity point, to act as a control function (Gelman and Imbens, 2019). The specification then becomes:

$$Y_{it} = \alpha + \tau SP500_{it} + \lambda f(S_{it}) + \epsilon_{it} \tag{3.3}$$

Where $f(S)$ can be a polynomial spline of any order. For example, the specification for a simple linear spline would be:

$$\lambda f(S) = \lambda_1 S_{it} + \lambda_2 SP500 * S_{it}$$

Controlling for this polynomial within the regression specification allows the coefficient τ to capture the discontinuous effect of being in or out of the S&P 500 index at the discontinuity point. This control function approach also allows for wider windows by flexibly controlling any small differences in characteristics related to the

[5]To clearly illustrate, it is hard to imagine why over time the firm ranked 501 based on size would be systematically different than the firm ranked 500, given the stochastic nature of firm size.

assignment variable that arise between firms.

The above specification can act as a baseline for RDD on an index if the inclusion criteria are perfectly known and observable, as outlined by Appel, Gormley, and Keim (2019). However, the S&P 500 index has particular challenges which must be overcome. While inclusion in the S&P 500 is determined primarily by size rank, as the index represents the largest 500 firms in the US, there are two potential caveats. First, the S&P 500 has a stated goal to minimize turnover. Firms will not be automatically removed or added for small changes in size even if the rankings change. The second caveat is that the S&P 500 can make exceptions to the size ranking methodology on the basis of promoting a representative index inclusive of all industries. Thus, being ranked in the top 500 does not immediately guarantee inclusion, and treatment is no longer random around the discontinuity point.

To deal with this complication, I use a fuzzy discontinuity design. One can view the treatment in the previous sharp RDD case to be an increase in the probability of inclusion in the SP500 index, from 0 to 1, when a firm meets the size rank criteria. While treatment status may no longer be exogenous when other criteria are introduced, being above the size threshold still is, and so too is the resulting increase in the probability of inclusion from meeting the size criteria. Being above the threshold increases the probability of inclusion by less than one if other criteria are used. The purpose of fuzzy RDD is to isolate this exogenous variation in probabilities and use it to estimate the effect of SP500 membership on comovement.

The fuzzy RDD can be implemented in two stages. In the first stage, we can estimate the following regression for the probability of a firm being included in the SP500 index:

$$SP500_{it} = \alpha + \lambda f(S_{it}) + \tau 1[S_{it} \geq \overline{S_t}] + \epsilon_{it} \tag{3.4}$$

Here $1[S_{it} \geq \overline{S_t}]$ is an indicator variable, equal to 1 if security i sizerank in time t is greater than the cutoff for inclusion, and τ captures the discontinuous increase in probability at the sizerank threshold. The fitted values from this regression reflect variation in the probability of inclusion due to meeting the size criteria and any explained variation due to sizerank differences captured by the control function, $f(S)$. These fitted values can then be used in regression (3) in place of the endogenous SP500 status. This fuzzy regression discontinuity design specification can be used to estimate the effect of membership when an index's inclusion criteria are less transparent, as outlined by Appel, Gormley, and Keim (2019).

To increase the efficiency of my estimates, I follow Lee and Lemieux (2010), Cellini, Ferreira, and Rothstein (2010), and Cunat, Gine, Guadalupe (2012) and supplement the polynomial control function with additional controls to capture several sources of variation directly. First, I include industry fixed effects to control industry-level heterogeneity, which may arise due to S&P's selection procedure. I also include year and industry by year fixed effects in case comovement within industries varies over time. Next, the turnover minimization objectives of the firm imply a potential difference in returns, size, and momentum across firms in and out of the index.[6] I therefore include additional controls to capture lag return, size, operating performance, and factor exposure differences. The primary benefit of including such controls is efficiency, and I conduct robustness checks to ensure that results are not driven by specification.

With the inclusion of controls and fixed effects, the main specification that I estimate is the following:

[6]Consistent with this intuitive prediction in the face of turnover minimization by index makers, Chen et al. (2016) show that additions are generally winners with high momentum relative to deletions.

$$Y_{it} = \alpha + \tau P(SP500_{it}) + \lambda f(S_{it}) + \delta TO_{it} + \psi_I + \phi_Y + \psi_I * \phi_Y + \epsilon_{it} \tag{3.5}$$

Where Y_{it} is the outcome variable, ψ_I, and ϕ_Y are industry and year fixed effects, respectively. Turnover control variables such as lagged returns, firmsize, tobinq, and assets are included within TO_{it}. The polynomial spline $f(S_{it})$ is over the forcing variable, size rank. $P(SP500)$ is the estimated probability that firm i will be in the S&P 500 at time t, which is estimated using the first stage:

$$SP500_{it} = \alpha + \lambda f(S_{it}) + \tau 1[S_{it} \geq \overline{S_t}] + \delta TO_{it} + \psi_I + \phi_Y + \psi_I * \phi_Y + \epsilon_{it} \tag{3.6}$$

Estimation of the first stage is conducted as a logit regression, with $SP500_{it}$ being a binary response variable. An apparent discontinuity in probability of S&P 500 membership can be seen in Table 3.2, with firms above the sizerank threshold more likely to be S&P 500 members. The second stage is conducted using OLS. To ensure consistent estimation while using a non-linear first stage with a linear second stage, I follow the procedure outlined by Wooldridge (2002) and run an intermediate OLS regression where SP500 is regressed against the fitted values from the first stage, using the same controls. The fitted values from this intermediate regression are the probability estimates used in the second stage.[7] The standard errors for inference are the standard 2SLS errors adjusted for heteroscedasticity and clustered at the firm level to account for within-firm correlation.

This is applied to the discontinuity sample, constructed by ranking firms by market capitalization each month, and selecting firms within a set range above or below the

[7] Adams, Almeida, and Ferreira (2009) use the same approach to 2SLS with a probit first stage in their study of founder-CEO effects on firm performance.

size rank inclusion threshold, depending on the width of the window studied.

I run results over different window sizes and specifications for turnover controls (including no turnover controls). I also use multiple specifications for the polynomial spline in the analysis to allow flexibility in slope and functional form and ensure results are robust to misspecification of the control function.

3.5 Results

I run specification (5) across several window sizes and the full sample of firms, and present the results in Table 3.3. In all cases, there is a clear increase in comovement for firms more likely to be included in the S&P 500. Estimates are also economically meaningful, with the estimate on the primary +/-50 discontinuity sample being 0.20. This means that a firm experiences excess comovement with the S&P 500 of 0.20 by being a member of the group and becoming exposed to group returns. The effect is robust across window sizes, which attenuates concern that it is driven by an abnormal group of firms around the discontinuity.

Such an increase in comovement in the absence of differing fundamentals has clear implications for asset pricing because the increase points to a group component in asset returns, consistent with Barberis et al. (2005). The result supports their theory that membership in the S&P 500 exposes firms to correlated exogenous demands from certain classes of investors who primarily inhabit the index. The pressure of such demands seems to give rise to simultaneous price movements in stock prices, which are reflected in comovement.

The delinking of prices and firm fundamentals suggested by such an increase in comovement points to important violations in several traditional asset pricing relationships. Informational efficiency relies on prices acting as an unbiased predictor of a firm's fundamental value, which is no longer the case if they reflect non-fundamental information. Efficient allocation of resources in the economy relies on prices reflecting the fundamental merits of the projects being invested in as well.

Feedback effects are also possible if exposure to a group component in returns introduces additional uncertainty about a stock's future return distribution unrelated to fundamentals, as future returns become linked to a distribution of exogenous group demand. If the market prices this uncertainty, it could increase expected returns on the security, increasing the firms cost of capital and influencing financing and investment decisions of firms, further affecting allocational efficiency.

3.5.1 Liquidity

The results in Table 3.4 control for liquidity differences to alleviate concerns within the comovement literature, raised by Barberis et al. (2005), that inclusion in the S&P 500 index may increase daily measures of comovement, mechanically, through an increase in liquidity. The traditional concern is that S&P 500 membership increases liquidity and increases the speed with which prices incorporate macroeconomic news, magnifying comovement on the day of the announcement relative to low liquidity firms, who realize some of the return on subsequent days instead. However, as Table 3.4 shows, results are very similar when we remove liquidity controls. The specifications differ in the inclusion of liquidity controls and the calculation of beta. Dimson (1979) adjusted betas, which sum comovement coefficients in a three-day window centered on the day of interest, are used in columns (3) and (4) to more directly capture any comovement due to liquidity related return persistence.

To better understand why index membership does not lead to liquidity driven comovement, I directly test for an increase in the speed of information diffusion after macroeconomic news is released, for S&P 500 firms, at a daily level. To do so, I follow Bernanke and Kuttner (2005) and estimate the reaction of S&P 500 and Non-S&P 500 stocks to surprise changes in interest rates announced by the FOMC, using the same sample of S&P 1500 stocks that I use in the rest of the paper. I use fed fund futures contracts to differentiate expected and surprise changes in interest rates around the FOMC event dates. I then regress index returns against those surprise changes in a window on and after the event date. The resulting point estimates and confidence intervals are presented in Figure 3.1a for the full sample of firms and Figure 3.1b for the discontinuity sample of firms with qualitatively similar insights. Both index groups experience a similar negative and significant return on the day of the news announcement, consistent with the results of Bernanke and Kuttner (2005), but returns on days after the announcement show no return spillover, with returns insignificant from zero and insignificantly different across groups.

This points to a similar average speed of information diffusion for stocks in and out of the index when measured daily, where it might affect our daily comovement measure. This is in line with research on intraday reaction to FOMC announcements by Zebedee, Bentzen, Hansen, and Lunde (2008) and FOMC minutes by Jubinski and Tomljanovich (2013), which find that the S&P 500 index and the full CRSP sample of stocks, respectively, seem to fully incorporate new macro news within 15 minutes of an announcement. Thus, any average benefits of increased liquidity on information diffusion coming from index membership are likely realized on the same trading day, at least for stocks within the S&P 1500.[8] This more directly rules out liquidity as a likely driver of excess comovement for S&P 500 stocks.

[8]It is worth a reminder here that the same minimum liquidity is required for firms to be included in the S&P 1500 and S&P 500, so we are not studying highly illiquid stocks.

3.5.2 Subsample results

My finding of significant excess comovement in S&P 500 member firms contrasts with insignificant short term inclusion effects documented by Chen et al. (2016). The difference points to important heterogeneous effects across time or the specific sample of firms studied. Therefore, I analyze several subsamples to better understand potential heterogeneity in the average effect and make comparisons to the findings of Chen et al. (2016) more direct.

The study of inclusion effects by Chen et al. (2016) covers a period spanning from 1971 to 2012, which differs from the sample period in my analysis. However, a period from 2001-2012 overlaps both studies and allows direct comparison. While they find no increase in comovement at inclusion, my RDD findings yield a significant increase, similar to my full sample, ruling out differing time periods as a driver of differing results. The congruency of the subsample results with my full period result also indicates robustness to the time period studied and some homogeneity in the effect across time.

I then focus on sampling differences inherent in each methodology. Using an event study methodology, Chen et al. (2016) focus only on new additions in a short window around inclusion in the index. They effectively study the transitory impact of index inclusion on comovement instead of the long term average treatment effect. To replicate this, I split my sample into observations for firms that have been added or deleted from the index within one year, and those which have not been added or deleted within one year. The RDD estimates in column (3) of Table 3.5 show no sign of an increase in comovement for firms within one year of inclusion, consistent with Chen et al. (2016). On its own, this estimate should be interpreted with caution given the large reduction in sample size for this narrow subsample, and the fact that fuzzy RDD's generally require large samples. However, combined with findings by Chen et al. (2016), this does point to heterogeneity in the average effect for new versus established members.

One possible explanation for a difference in the short term effect of inclusion versus the long term effect of membership is a delay in integration into the group when firms are added, as new investors drawn to the stock take time to rebalance their holdings. This could reduce short term comovement for two reasons. If integration is slow, newly added firms are initially not part of the group portfolio and thus do not co-move. Comovement would increase slowly as investor holdings increase and firms become exposed to the group component of returns over time. The delay could also cause a competing decrease in comovement during the rebalancing period, as the rebalancing itself creates an unconditional increase in demand for the stock over time, which is otherwise uncorrelated to both fundamentals and the group return. [9]

Whatever the reason, it does appear that the treatment effect is different for newly added firms and established members. Further research into the difference, perhaps exploring the role of passive investor flows directly after inclusion, may shed additional light on the difference. However, the interest in this paper is understanding the average effect on member firms, not the dynamics of adjustment that take place to integrate the stock into the group, and the RDD estimation contributes direct evidence that the average S&P 500 firm experiences excess comovement.

The final column of Table 3.5 explores the potential for differences driven by the business cycle. The sample used includes only observations during recessions, as defined by the NBER. The smaller sample size limits power, but point estimates are consistent, if smaller in magnitude than the full sample, indicating no clear sign of

[9]Investors may reduce demand for the index as a whole, but simultaneously increase demand on the stock as they add it to their portfolio.

heterogeneous effects across the business cycle.

3.5.3 Robustness

I next conduct several standard robustness checks to ensure misspecification of my control function is not driving my results. Table 3.6 shows results across several different polynomial spline specifications, with similar results across each. In fact, for the fairly narrow +/- 50 sample, a simple linear spline seems sufficient to capture any minor relationship between the assignment and outcome variables. This is not surprising, given that the narrow window itself should control most differences. The control function will likely have more work to capture differences for wider windows.

Next, I present results for several different control specifications in Table 3.7. Results are similar, showing that the choice of controls does not drive my findings. Finally, I present a standard covariate balance test in Table 3.8, which follows Lee and Lemieux (2010) and repeats the fuzzy RDD using a variety of firm characteristics as the outcome variable. There is no sign of uncontrolled differences in observables for S&P and non-S&P firms which could influence results.

3.6 Conclusion

In this paper, I use a fuzzy regression discontinuity design to estimate the long term causal effect of membership in the S&P 500 on comovement. I find that, in contrast to the previously studied short term effect at inclusion, S&P 500 members experience significant excess comovement with the S&P 500 that cannot be explained by differences in fundamentals.

My findings support the theory of Barberis et al. (2005), which states that membership in a popular group of assets such as the S&P 500 can expose firms to correlated exogenous demands of certain classes of investors which primarily inhabit the index, and introduce a non-fundamental group factor to asset returns. Such a finding points to a violation of informational efficiency, which has important implications for asset pricing theory.

The concept of group based comovement that I study focuses on the S&P 500 but implies the potential for similar frictions in other groups of assets. Further investigations into the comovement effects of other groupings could provide additional insights. The difference in the short- and long-term estimates points to potential delays in integration into the group after inclusion, which suggests another exciting avenue for future research.

Table 3.1: Full Sample Description

This table provides descriptive statistics for the main outcome variables, controls used in later analysis, and additional characteristics of interest for firms included and not-included in the S&P 500 between Jan. 1, 1995 and Dec. 31, 2017. The unit of observation is at the firm-month level. S&P 500 comovement is calculated using regression specification (1) with the influence of the firm in the regression removed from the S&P 500 return. Dimson adjusted comovement sums the S&P 500 betas on a 3-day window centered on the day of interest to account for non-synchronous trading and delayed reaction due to liquidity. Illiquidity is the monthly measure from Amihud (2002). Eff. Tick is the effective spread proxy from Holden (2009). Factors umd, hml, smb, cma, rmw represent firms comovement over a given month with the respective factor portfolio listed. Lagged return variables are holding period returns. FinLiq is the ratio of cash and short term investments to current liabilities. Current is the ratio of current assets to current liabilities. Leverage is the ratio of liabilities to shareholder's equity. Non-ratio financial characteristics are in $millions.

	Non-S&P 500			*S&P 500*		
Variable	N	Mean	Std Dev	N	Mean	Std Dev
SP500 Beta	264057	1.03	0.51	136302	1.01	0.44
Dimson Beta	263746	1.104	0.608	136203	1.053	0.528
Illiquidity	275167	0.219	0.830	137803	0.020	0.543
Eff. Tick	275174	0.187	0.412	137807	0.089	0.185
1mo return	275028	0.009	0.148	137776	0.007	0.12
1yr return	270001	0.092	0.492	135866	0.078	0.481
3yr return	262395	0.286	1.313	133325	0.21	0.806
5yr return	246517	0.556	3.171	130137	0.374	1.472
Firmsize	275242	1550691	1596633	137890	23383822	43335298
Umd	275167	-0.078	2.038	137803	-0.085	1.534
Hml	275167	0.158	2.567	137803	0.084	1.957
Smb	275167	0.978	2.023	137803	0.176	1.427
Cma	275167	-0.028	3.219	137803	0.107	2.44
Rmw	275167	-0.001	0.029	137803	-0.001	0.022
Assets	274965	2828	5706	137785	57207	202123
Acquisitions	244407	61.75	254.79	112863	419.61	1407.9
Capx	251316	85.975	174.134	124897	1157.399	2733.404
Roa	256363	0.029	0.137	127004	0.051	0.084
Roe	256092	-0.022	0.401	126924	2.765	132.816
Roi	256078	0.046	0.344	127004	0.094	0.222
TobinQ	256104	1.878	1.305	126924	2.081	1.326
Sales	256576	1662	2754	127109	18740	38332
R&D	275242	21.143	51.775	137890	401.887	1405.468
FinLiq	230258	1.014	2.05	113376	0.596	0.918
Current	231721	2.557	2.337	113150	1.689	1.116
Leverage	255755	1.24	3.233	126647	476.317	22995.43
Div. Yield	273290	0.954	3.464	136869	12.422	31.845

Table 3.2: Logistic Regression: Probability a firm is in the S&P 500

This table reports the results of the logistic regression outlined in equation (3.6), where the outcome variable is equal to 1 if a firm is in the S&P 500 and 0 if it is not. Above is the indicator variable of interest $\tau 1[S_{it} \geq \overline{S_t}]$ representing whether the firm sizerank is greater than the threshold for inclusion. The coefficient captures the discontinuous change in probability of treatment (S&P 500 membership) around the inclusion threshold. Unreported controls include a cubic polynomial spline on sizerank plus Year, Industry and Year X Industry fixed effects. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

Variable	Estimate	Chi Square
Above	0.1269***	9.9270
1mo Return	-0.4173***	44.3132
1yr Return	-0.5548***	636.4334
3yr Return	-0.4815***	1391.4599
5yr Return	-0.0767***	248.2343
UMD	-0.0628***	240.9485
HML	0.0421***	139.5779
SMB	-0.0722***	276.4700
RMW	1.3203***	20.8034
CMA	0.0445***	271.4598
Firmsize	1.1722***	216.1258
Assets	0.000086***	3795.0173
TobinQ	-0.1144***	321.0367

Table 3.3: Fuzzy RD Effect Estimates across Window Size

This table reports the estimated treatment effect of SP500 membership based on the fuzzy regression discontinuity design (RDD) specification in equation (3.5). P(500) is a firms propensity score estimated from the first stage in equation (3.6) as described in the text. The dependent variable is the estimated beta between the stock and the SP500 index from equation (3.1). Specifications differ in window size around the SP500 size based cutoff point. Controls include 1mo, 1yr, 3yr, 5yr lagged returns, factor exposures to hml, smb, cma, rmw, umd, firmsize, assets, and tobinsq, illiquidity, and effective tick. Two stage least squares standard errors are heteroskedasticity adjusted and clustered at the firm level. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively. P-values are in parentheses.

Specification	(1)	(2)	(3)	(4)
Dep. Var	Beta	Beta	Beta	Beta
Window	25	50	100	Full
Spline	Cubic	Cubic	Cubic	Cubic
P(SP500)	0.149**	0.200***	0.197***	0.129***
	(0.024)	(0.002)	(< 0.001)	(< 0.001)
Controls	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Industry FE	Y	Y	Y	Y
Year × Industry FE	Y	Y	Y	Y
Obs	11983	23687	46623	342915
R-Square	0.47	0.46	0.45	0.40

Table 3.4: Fuzzy RD Effect Estimates across Liquidity Controls

This table reports the estimated treatment effect of SP500 membership based on the fuzzy regression discontinuity design (RDD) specification equation (3.5). P(500) is a firms propensity score estimated from the first stage in equation (3.6) as described in the text. The dependent variable is the estimated beta between the stock and the SP500 index from equation (3.1). The specification differs on dependent variable and liquidity controls. Specification (1) excludes liquidity controls. Specification (2) is the same as reported in table (5) with liquidity controls. Specification (3) uses a dimson adjusted beta without liquidity controls. Specification (4) uses dimson adjusted beta with liquidity controls. Controls include 1mo, 1yr, 3yr, 5yr lagged returns, factor exposures to hml, smb, cma, rmw, umd, firmsize, assets, and tobinsq. Two stage least squares standard errors are heteroskedasticity adjusted and clustered at the firm level. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively. P-values are in parentheses.

Specification	(1)	(2)	(3)	(4)
Dep. Var	Beta	Beta	Dimson	Dimson
Window	50	50	50	50
Spline	Cubic	Cubic	Cubic	Cubic
P(SP500)	0.200***	0.188***	0.214**	0.206**
	(0.002)	(0.005)	(0.011)	(0.018)
Illiquidity	−0.276***		−0.303***	
	(< 0.001)		(0.002)	
Eff. Tick	0.259***		0.282***	
	(< 0.001)		(< 0.001)	
Controls	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Industry FE	Y	Y	Y	Y
Year × Industry FE	Y	Y	Y	Y
Obs	23687	23687	23671	23671
R-Square	0.46	0.45	0.41	0.40

Table 3.5: Heterogenous Effects across Subsamples

This table reports the estimated treatment effect of SP500 membership based on the fuzzy regression discontinuity design (RDD) specification in equation (3.5). P(500) is a firms propensity score estimated from the first stage in equation (3.6) as described in the text. The dependent variable is the estimated beta between the stock and the SP500 index from equation (3.1). The specification differs in sample. Specification (1) is 1995-2017. Specification (2) is 2001-2012. Specification (3) includes observations within 1 year of being included or removed from the SP500 index. Specification (4) includes observations not within 1 year of being included or removed from the SP500 index. Specification (5) only includes observations during an NBER defined recession period. Controls include 1mo, 1yr, 3yr, 5yr lagged returns, factor exposures to hml, smb, cma, rmw, umd, firmsize, assets, and tobinsq, illiquidity, and effective tick. Two stage least squares standard errors are heteroskedasticity adjusted and clustered at the firm level. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively. P-values are in parentheses.

Specification	(1)	(2)	(3)	(4)	(4)
Sample	Full	2001-2012	Within 1 Yr	Out 1 Yr	Recession
Dep. Var	Beta	Beta	Beta	Beta	Beta
Window	50	50	50	50	50
Spline	Cubic	Cubic	Cubic	Cubic	Cubic
P(SP500)	0.200***	0.227***	0.015	0.205***	0.101
	(0.002)	(0.006)	(0.909)	(0.001)	(0.367)
Controls	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y
Industry FE	Y	Y	Y	Y	Y
Year × Industry FE	Y	Y	Y	Y	Y
Obs	23687	12405	1334	22367	2207
R-Square	0.46	0.42	0.68	0.46	0.55

Table 3.6: Fuzzy RD Effect Estimate across Spline

This table reports the estimated treatment effect of SP500 membership based on the fuzzy regression discontinuity design (RDD) specification in equation (3.5). P(500) is a firms propensity score estimated from the first stage in equation (3.6) as described in the text. The dependent variable is the estimated beta between the stock and the SP500 index from equation (3.1). The specifications differ by the order of the polynomial spline used as a control function. Controls include 1mo, 1yr, 3yr, 5yr lagged returns, factor exposures to hml, smb, cma, rmw, umd, firmsize, assets, and tobinsq, illiquidity, and effective tick. Two stage least squares standard errors are heteroskedasticity adjusted and clustered at the firm level. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively. P-values are in parentheses.

Specification	(1)	(2)	(3)	(4)
Dep. Var	Beta	Beta	Beta	Beta
Window	50	50	50	50
Spline	Linear	Quadratic	Cubic	Quartic
P(SP500)	0.198***	0.200***	0.200***	0.200***
	(0.002)	(0.002)	(0.002)	(0.002)
Controls	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Industry FE	Y	Y	Y	Y
Year $\times$ Industry FE	Y	Y	Y	Y
Obs	23687	23687	23687	23687
R-Square	0.46	0.46	0.46	0.46

Table 3.7: Fuzzy RD Effect Estimates over Control Specification

This table reports the estimated treatment effect of SP500 membership based on the fuzzy regression discontinuity design (RDD) specification in equation (3.5). P(500) is a firms propensity score estimated from the first stage in equation (3.6) as described in the text. The dependent variable is the estimated beta between the stock and the SP500 index from equation (3.1). The specifications differ in included turnover controls. Controls include 1mo, 1yr, 3yr, 5yr lagged returns, factor exposures to hml, smb, cma, rmw, umd, firmsize, assets, and tobinsq, illiquidity, and effective tick. Two stage least squares standard errors are heteroskedasticity adjusted and clustered at the firm level. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively. P-values are in parentheses.

Specification	(1)	(2)	(3)	(4)
Dep. Var	Beta	Beta	Beta	Beta
Window	50	50	50	50
Spline	Cubic	Cubic	Cubic	Cubic
P(SP500)	0.200***	0.412***	0.250***	0.160***
	(0.002)	(0.009)	(< 0.001)	(0.005)
Controls	Y	N	Y	Ret & Size
Year FE	Y	Y	Y	Y
Industry FE	Y	Y	Y	Y
Year × Industry FE	Y	Y	N	Y
Obs	23687	27227	23687	25637
R-Square	0.46	0.41	0.31	0.44

Table 3.8: Fuzzy RD Covariate Balance Test

This table reports the results of a covariate balance test which includes various firm characteristics as the outcome variable in the same fuzzy RD specification used to generate main results from equation (3.5). Estimates represent the difference in the listed variable for S&P 500 and non-S&P 500 firms in the +/- 50 discontinuity sample. The polynomial spline is cubic and controls include 1mo, 1yr, 3yr, 5yr lagged returns, factor exposures to hml, smb, cma, rmw, umd, firmsize, assets, and tobinsq, illiquidity, and effective tick. Two stage least squares standard errors are heteroskedasticity adjusted and clustered at the firm level. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively. P-values are in parentheses.

Variable	Estimate	p-value
ROE	0.020	(0.550)
ROI	-0.019	(0.509)
Sales	244.258	(0.958)
Capx	40.496	(0.871)
R&D	30.893	(0.322)
Acquisitions	269.820	(0.125)
Dividend Yield	2.347	(0.264)
Fin. Liquidity	-0.072	(0.768)
Current	-0.299	(0.312)
Leverage	1.087	(0.255)

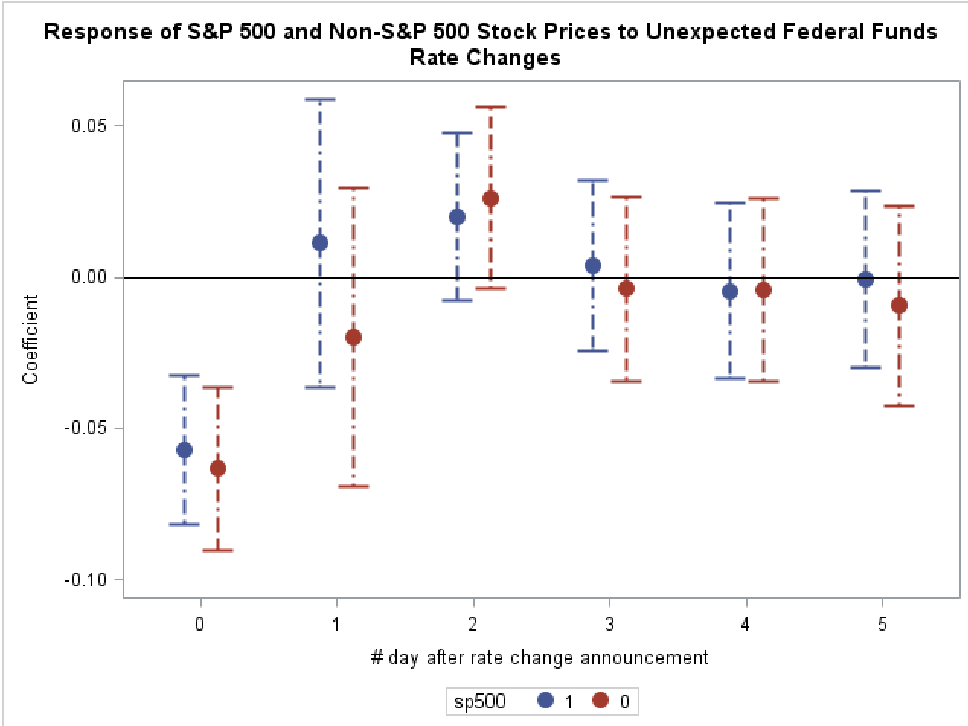

(a)

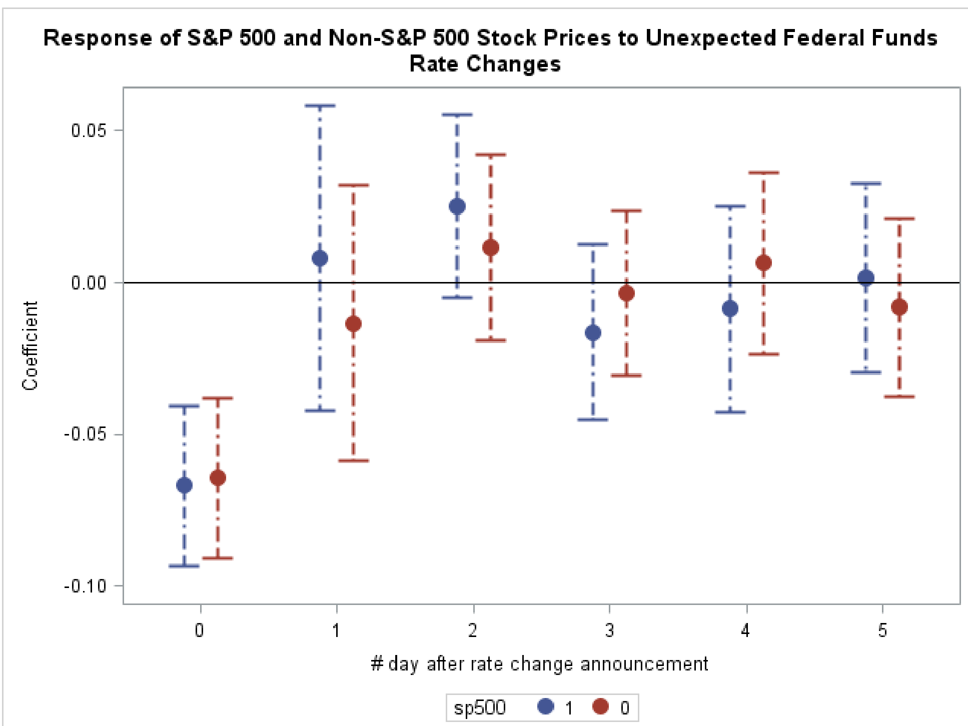

(b)

Figure 3.1: Stock reaction to FOMC rate changes

Displays point estimates and 95% confidence intervals for the stock price reaction of S&P 500 and non-S&P 500 stocks to surprise FOMC announcements about federal funds rate changes on the announcement date and over the post-announcement window. Index returns are calculated as a value weighted average of (a) the full sample of S&P 1500 firms (b) the +/- 50 sample of firms around the S&P 500 size rank cutoff.

Chapter 4

Excess Coskewness and Financialization

4.1 Introduction

Increased investment flows into passively managed investment vehicles, such as ETFs, have raised questions over how well the market can price securities when so many participants are not investing based on a fundamental evaluation of the investment characteristics of the assets they are buying. Generally, ETFs and other passive investments track a large basket of securities, and the demand for the assets within the basket becomes a function of the demand for the basket as a whole, regardless of its specific investment qualities. This leads to the obvious question of whether such flows can add non-fundamental characteristics to the assets' return process and distort its price.

A few signs of this potential distortion include the documented increase in return when firms enter the S&P 500 (Jain, 1986; Shleifer, 1987), excess comovement of stocks within an index with each other (Barberis, Shleifer, and Wurgler, 2005), and the trend of increased comovement of commodities with broader financial markets over time, particularly when those commodities are added to commodity indices (Tang and Xiong, 2012). In this paper, I move away from the previous focus on expected returns and comovement and, for the first time, look at how financialization affects asset skewness and coskewness.

Coskewness, or systematic skewness, refers to the portion of an asset's return skewness that cannot be diversified away. Assets with positive coskewness increase the skewness of a portfolio when added, resulting in a higher likelihood that assets move together during upswings in their price and a lower likelihood that assets move together during downswings. This means that, with relatively higher coskewness, a portfolio will be better diversified against downturns without diversifying away upside. For this reason, positive coskewness is generally viewed as beneficial to investors, while negative coskewness is viewed as detrimental.

While my primary goal is exploration rather than isolation of the exact mechanism by which financialization might induce excess coskewness, there are some potential explanations. If excess coskewness arises due to large passive investment flows, it is plausible that it is more difficult during fast moving market downturns for active market participants to identify and trade away mispricing. If so, we would expect more correlated losses during market downturns and less correlated gains during upturns.

Loss aversion is another explanation, with a relatively stronger correlated response to losses than to gains of a similar size amongst less sophisticated index or habitat traders that trade in a particular group of securities. Even if the trading is handled by a more sophisticated investor, like a managed mutual fund that trades a certain group of assets, they ultimately represent the wealth of less sophisticated individuals. If those individuals are loss averse, they can put pressure through redemptions on mutual funds, causing them to liquidate a broad basket of assets within the group, regardless of fundamentals.

I utilize two methodologies to explore how financialization affects asset coskewness with a given market. The first is an event study methodology that evaluates the change in an asset's coskewness with the S&P 500 market before and after inclusion in the index. This method primarily captures relatively short-term dynamics of S&P 500 inclusion. I find that, over the short run, there is excess positive coskewness. This could partially be explained by changes in fundamentals or short-run transitory pressures that arise from inclusion. [1] One example could be the documented increase in momentum upon index inclusion in Chen, Singal, and Whitelaw (2016), which could cause stocks to experience generally positive returns after inclusion, naturally leading to positive coskewness. This example does not preclude others and simply illustrates one of the weaknesses of this type of analysis.

[1] Indeed, significant positive coskewness is weak if we remove the first year after inclusion, suggesting that it is driven by short-run dynamics.

To more accurately explore the effects of financialization, I estimate the average effect of S&P 500 membership on long-term comovement. Membership in the S&P 500 is largely based on the ranked size of a firm. I exploit this by using a fuzzy regression discontinuity design to isolate exogenous variation in the probability of S&P 500 membership which arises between firms that just meet the size criteria and those that just miss the size criteria. Since firm size has an idiosyncratic component, firms which are very close in size should not be systematically different from each other, yet meeting the membership criteria leads to a discrete increase in the probability of being selected into the S&P 500. This can be used to measure the effect of membership itself.

I find that, relative to a firm excluded from the S&P 500 due to a minor, random size difference, firms included in the S&P 500 have more undesirable, negative coskewness with the S&P 500. This has several interesting implications. First, due to the fuzzy RDD control scheme, the differences in fundamentals across firms are randomized. This implies that the excess negative coskewness is a result of a non-fundamental factor driving asset returns and that this factor is directly tied to exposure to the attentions of S&P 500 investors, supporting the idea that grouping assets together into widely traded habitats, or baskets, can attract flows that ignore individual asset level fundamentals and induce mispricing.

Second, since negative coskewness is generally viewed as undesirable, it could have a feedback effect on the expected returns of S&P 500 firms, increasing expected returns and the cost of capital for S&P 500 firms. Such non-fundamental increases in the cost of capital could have implications for allocative efficiency.

Finally, negative coskewness implies that the diversification benefit of passively holding a market portfolio, like the S&P 500, which is common among unsophisticated investors, may be reduced exactly when diversification is most needed. This could have widespread welfare implications for those flocking into such passive strategies, given the growth of passive investing.

To ensure the robustness of the results, I estimate a plethora of specifications across multiple window sizes and control specifications. I show that estimates of a firm's probability of S&P 500 membership are unrelated to observable firm characteristics and specifically rule out liquidity differences between indexes as a confounding effect.

The remainder of the paper is organized as follows: Section 2 reviews the literature and provides more context for my contributions, while section 3 describes my data. Section 4 provides details and short-run results for my event studies, while section 5 provides details and long run results for my regression discontinuity analysis. Section 6 discusses the robustness of my regression discontinuity results. Section 7 concludes.

4.2 Literature Review

4.2.1 Financialization

One of the defining trends within financial markets in the last two decades has been the increasing popularity of passive investments across a wide range of assets. This has the well-known benefit of increasing trade and liquidity in these assets, theoretically reducing transaction costs for financial market participants, and increasing the flow of speculative funds to meet hedging demand, particularly in commodity markets. However, there are some interesting documented side effects. Amongst equities, several studies show that stocks added to popular indices experience higher returns, [2] while Barberis, Shleifer, and Wurgler (2005) show that inclusion in the S&P 500

[2] Studies include Shleifer (1986) and Wurgler and Zhuravskaya (2002) for S&P 500 inclusions, Kaul, Mehrotra, and Morck (2000) for the TSE 300, and Greenwood (2008) for the Nikkei 225.

leads to excess comovement amongst stocks. Additionally, Harford and Kaul (2005) are able to trace correlated order flow to strong common effects in the returns of S&P 500 stocks. Amongst commodities, Tang and Xiong (2012) show that commodities included in popular broad-based commodity indices experience increased correlations between themselves and wider financial markets.

There is a contentious literature going back to 2008 debating whether passive investment was to blame for a dramatic bubble in commodity prices, which is summarized neatly in Cheng and Xiong (2014). Meanwhile, Henderson, Pearson, and Wang (2016) show that trading in futures contracts unrelated to asset fundamentals impacts both futures and spot prices of underlying commodities, and chapter 2 of this thesis links rising financialization to speculation driving distortions in commodity futures and spot prices.

The literature so far has focused on finding the direct effects of financialization on firms' returns or their correlation and comovement with each other. The effect of financialization on asset skewness is an important, but largely ignored question, and is the focus of my study.

4.2.2 Investor Skewness Preferences

Preference for positive skewness is an accepted characteristic of rational investors, alongside preference for higher returns and aversion to variance within an expected utility framework. The theoretical underpinnings for such a preference simply rely on investors' utility functions having decreasing marginal utility of wealth, and non-increasing absolute risk aversion, which are two of the desirable properties outlined in Arrow (1971). Further, asset returns are generally non-normal and often non-symmetric, so skewness preferences are a relevant feature in asset pricing. These characteristics motivate the work of Kraus and Litzenberger (1976) and Harvey and

Siddique (2000), both of which include coskewness in a three-moment CAPM model. It is this strand of literature regarding coskewness that I build upon by evaluating the impact that financialization can have on firms' coskewness.

4.2.3 Benefits of Diversification

My research is also potentially related to observed asymmetries in the benefits of diversification. Several studies have shown that correlations between assets seem to increase during market downturns, reducing the effectiveness of diversification to reduce risk, precisely when needed most. The most famous example is the well documented increase in correlation between international equity markets during bear markets. (Longin and Solnik, 2001) However, increased correlations are not limited to international markets, as Ang and Chen (2002) show that correlations between US stocks and the aggregate US market increase during market downturns. This could be explained if excess coskewness arises due to asymmetric effects of financialization. For example, if comovement increases during economic downswings when limits to arbitrage may be binding, or if the correlated flows of passive investors are more extreme in market downturns, this could manifest as both negative coskewness in asset returns and a reduction in benefits of diversification during downturns.

4.3 Data

To conduct my analysis, I follow the same procedure as chapter 3, and utilize monthly index constituency data and a sample of inclusion dates from the Compustat S&P index file. I limit the sample to S&P 1500 stocks, as they form the pool of stocks that meet the specific minimum liquidity and ownership characteristics required for S&P 500 inclusion. Since the S&P 1500 was created with the introduction of the S&P 600 small cap index in 1994, the study period is January 1995 to Dec 2017.

Security level data and industries are merged from the monthly CRSP file to make a sample of firm-month observations. Coskewness calculations rely on index return data from CRSP daily index files, and I obtain daily stock returns and volume data from the CRSP daily and monthly stock files. To measure monthly exposures of stocks to Fama-French factors (Fama and French, 2015) and Carhart's momentum factor (Carhart, 1997), I use data from Ken French's data library and WRDS, respectively, and classify industries according to the standard Fama-French 12 industry definitions available from Ken French's data library.

The measure of coskewness I use in the fuzzy RDD follows the standard systematic skewness calculation introduced by Kraus and Litzenberger (1976):

$$Return_{id} = \omega + \beta_{im} Index_{kd} + \gamma_{im} Index_{kd}^2 + \epsilon_{id} \tag{4.1}$$

Which is the quadratic regression of excess daily returns of stock i on the daily excess returns for index k for each month m, estimated monthly using a rolling 12 month regression period. Coskewness is captured as the parameter γ_{im}, and this γ_{im} is the month m gamma for firm i with index k, which I use as the monthly coskewness outcome variable. The index of interest is the S&P 500 due to its unique exposure to large classes of investors. In all cases, I remove the contribution of the individual stock being regressed to the index return to avoid spurious results, though given the size of the firms around the discontinuity and the fact that the index is value-weighted, this makes no practical difference to the skewness measure or later results. The short-run event study results use the same measure of coskewness but are calculated differently, as outlined in section 3.

Table 4.1 shows descriptive statistics for all variables over the full sample period.

Notably, coskewness is negative for both groups of stocks, but more so for non-S&P 500 firms. These simple means capture coskewness from all sources, including coskewness that arises naturally from firm fundamentals. Unsurprisingly, coskewness is also relatively noisy compared to comovement.

4.4 Short Term S&P 500 Inclusion Effect

In line with the traditional literature measuring inclusion effects on the moments of the CAPM, I start by analyzing the change in coskewness for firms after they are added to the index. I run a simple, univariate test of the return of stock i against the return on the S&P 500 index, excluding the influence of stock i from the index whenever it is a member:

$$R_{i,t} = \alpha_i + \beta_{i1} POST_t + \beta_{i2} SP500_t + \beta_{i3} POST_t \times SP500_t + \gamma_{i1} SP500_t^2 + \gamma_{i2} POST_t \times SP500_t^2 + \epsilon_{it} \tag{4.2}$$

The coefficient γ_{i1} captures the baseline coskewness of stock i during the 12 month estimation period prior to the start of the month in which stock i is added to the index. The coefficient γ_{i2} captures the additional coskewness estimate for the 12 month post-estimation period starting at the beginning of the month following stock i's inclusion, since $POST$ is a dummy variable representing whether stock i is in the S&P 500 index at time t. I can then interpret γ_{i2} as the change in coskewness for stock i when it is added to the S&P 500 index. This estimate of coskewness is consistent with Kraus and Litzenburger (1976), and is noted by Harvey and Siddique (2001) as a valid measure of a stock's coskewness with an index.

Table 4.2 provides the average OLS coefficients across all stocks in the sample. Consistent with similar tests in Barberis, Shleifer, and Wurgler (2005), there is a significant increase in comovement upon inclusion. Surprisingly, there is a significant increase in the coskewness of a stock with the S&P 500 in the year following inclusion. This implies that firms added to the index are more correlated with large positive moves of the index, but less closely correlated with large negative moves.

There are at least two possible explanations for this increase in coskewness. First, I could be capturing a permanent increase in coskewness related to membership in the S&P 500. However, later long term regression discontinuity results do not corroborate this. Second, there could be a violation of the assumption that firm fundamentals remain constant after inclusion. For example, momentum changes on inclusion could lead to positive returns in the short-run afterward, creating a mechanical correlation with the market when it is going up, and no correlation when it is going down. Chen et al. (2016) find effects from momentum partially driving increased comovement when firms are added to the index, pointing to a violation of the event study's assumptions. In this case, the captured increase in coskewness is real, but is not representative of the permanent change we are interested in capturing.[3]

Next, I conduct a bivariate regression to compare the coskewness response of a stock to both the S&P 500 index, which it enters, and the non-S&P 500 group of stocks, which it leaves:

$$R_{i,t} = \alpha_i + \beta_{i1} POST + \beta_{i2} SP500 + \beta_{i3} POST \times SP500 + \beta_{i4} NonSP500 + \beta_{i5} POST \times NonSP500 + \gamma_{i1} SP500^2 + \gamma_{i2} POST \times SP500^2 + \gamma_{i3} NonSP500^2 + \gamma_{i4} POST \times NonSP500^2 + \epsilon_{it} \tag{4.3}$$

In the absence of any frictions or distortions to the stocks return process or any change to fundamentals around inclusion, we would expect coskewness to be un-

[3] I find positive coskewness even when the first full month after inclusion is dropped. However, when I drop the first year, positive coskewness largely disappears. This may suggest that positive coskewness is a short-run phenomenon and is not permanent.

changed with both groups. If there is a distortion, or a change in fundamentals, then we would expect additions to the S&P 500 to result in significant changes in coskewness with the new index, and potentially an opposite effect on the old non-S&P 500 index. Table 4.3 shows that this is roughly the case; there is a positive and significant increase upon inclusion with the S&P 500 and a negative, albeit insignificant effect on coskewness with the old non-S&P 500 group of stocks. The positive sign on inclusion remains consistent with the results of the univariate test.

Chen et al. (2016) argue that this bivariate regression is highly sensitive to small changes in parameters, and instead recommend using separate univariate regressions to estimate the change in coskewness with each group before and after the event. I, therefore, repeat the univariate test for non-SP500 returns. Results in panel A of Table 4.4 are slightly changed. There is still a positive and significant increase in coskewness of returns with the SP500 upon inclusion, but the change in coskewness with the non-SP500 group changes sign to positive as well, albeit still insignificant. I then test the difference and find that the increase in coskewness with the SP500 is still significantly more than the increase in coskewness with the non-SP500 group, which is consistent with excess positive coskewness upon inclusion. Results comparing two bivariate regressions run separately before and after inclusion, presented in Panel B of Table 4.4, are materially similar to the full bivariate regression above.

4.4.1 Subperiod Results

I address whether coskewness effects of S&P 500 inclusion have changed over time by splitting the sample before and after the start of 2003 and repeating all of the tests. The post 2002 period is generally associated with a rapid rise in passive investing, during which distortive asset class effects driven by passive investment should be stronger. However, I find mixed results for differences across subperiods in the tests.

The univariate test of Table 4.2 yields a more positive coskewness effect upon inclusion for the pre-2003 period relative to the post-2003 period, while the bivariate results in Table 4.3 are the exact opposite. Given the critique of Chen et al. (2016), this may point to sensitivity in the bivariate results. When I compare univariate estimates between the non-SP500 group and the SP500 index, I find that changes to coskewness with the SP500 upon inclusion are more positive in the earlier period, and the same is true of coskewness with the non-SP500 group. When I evaluate the difference in the change in coskewness between the two groups for each time period, I find that the increase in coskewness with the SP500 upon index inclusion relative to the increase in coskewness with the non-SP500 group, is virtually unchanged.

Once again, these mixed results may point to the event study methodology capturing short term, transitory, confounding effects upon inclusion. For example, short term momentum after inclusion like that documented in Chen et al. (2016), leading the stock to rise regardless of market conditions, and mechanically increasing coskewness with both SP500 and non-SP500 markets. Even though passive investment has risen in the more recent period, as Barberis, Shleifer, and Wurgler (2005) point out, the S&P 500 has long been a popular investment habitat, and even without passive investment, addition to the S&P 500 attracts significant attention and access to a large group of investors who populate that specific group of stocks.[4]

4.5 Regression Discontinuity Analysis

While the above event study methodology captures the short-term response of coskewness upon inclusion, potentially confounded by transitory effects, our actual interest

[4] Appendix A, Tables A.6-A.7, contains an analysis of deletions and supports a generally reversed, but much weaker effect of deletion on comovement and coskewness. This is consistent with past evidence of weaker deletion effects for S&P 500 stocks, and may point to residual attention that a stock maintains after S&P 500 membership.

lies in estimating the effect of S&P 500 membership, not addition to the S&P 500, on coskewness. This is a subtle but important difference, as what we are truly interested in is isolating the persistent, long term effect of S&P 500 membership in driving excess coskewness and distorting asset returns.

To estimate this long term treatment effect of S&P 500 membership I use a fuzzy regression discontinuity framework. Unlike the event study, this method does not rely on a discontinuous effect before and after the event time period, or an assumption that the firm is exactly the same before and after that time period, to achieve identification. Instead, the regression discontinuity relies on a sample being similar on either side of a treatment threshold based on treatment criteria. In my case, this treatment criteria is the ranked size of a firm, and I isolate variation in outcomes due to a discontinuous probability of being included in the S&P 500, arising from meeting the S&P 500's size criteria. Since firm size has a random component, very small differences in firm size and size rank are likely to be random, and not due to systematic differences, making treatment effectively random within a small window around the inclusion criteria threshold.

I use a fuzzy, instead of a sharp, RDD, because meeting the size rank criteria alone does not guarantee immediate membership in the S&P 500, which would be required to use a sharp RDD. In addition to choosing the largest firms for membership, S&P also has a stated goal of minimizing excessive turnover in the index, and promoting industry representativeness. This can result in a lag before firms are added or deleted after meeting the inclusion criteria, and firms in some industries may be more likely to be added than firms in others. Therefore, instead of isolating variation in S&P 500 membership due to a guaranteed treatment, I isolate variation in the probability of being treated, making this a fuzzy RDD.

I implement the fuzzy RDD in two stages. In the first stage I estimate the logistic regression to isolate the discontinuous probability of being included in the index based on size rank:

$$SP500_{it} = \alpha + \lambda f(S_{it}) + \tau 1[S_{it} \geq \overline{S_t}] + \delta TO_{it} + \psi_i + \phi_Y + \psi_i * \phi_Y \epsilon_{it} \tag{4.4}$$

Here $SP500_{it}$ is an indicator variable equal to 1 if a firm is in the S&P 500, $\overline{S_t}$ is an indicator variable equal to 1 if a firm meets the stated size criteria of the S&P 500, $f(S_{it})$ is a polynomial spline controlling for continuous differences in probability of inclusion due to sizerank. The coefficient τ then captures the discontinuous difference in probability of inclusion due to meeting the inclusion criteria. I also include ψ_i and ϕ_Y as industry and year fixed effects, alongside turnover controls TO_{it} such as lagged returns, firmsize, tobinsq, and assets. These controls are included to increase the efficiency of estimates by controlling other sources of variation in treatment probability across firms, in line with Lee and Lemieux (2010), Cellini, Ferreria, and Rothstein (2010) and Cunat, Gine, and Guadalupe (2012). Estimation of this logistic regression in Table 4.5 reveals a clear discontinuity, with firms above the sizerank threshold more likely to be S&P 500 members.

The fitted values of $SP500_{it}$ represent firm i's estimated probability of being included in the S&P 500 conditional on sizerank and controls. I then run the following second stage regression, using the fitted values from the first stage $P(SP500_{it})$:

$$Y_{it} = \alpha + \tau P(SP500_{it}) + \lambda f(S_{it}) + \delta TO_{it} + \psi_I + \phi_Y + \psi_I * \phi_Y + \epsilon_{it} \tag{4.5}$$

Here, Y_{it} is the outcome variable of interest, which is the coskewness measure from equation (4.1). I also include all control variables from the first stage regression, including the polynomial spline $f(S_{it})$. The primary regressor of interest is $P(SP500_{it})$, and the coefficient τ in this case captures the change in coskewness if the probability of inclusion in the index increases by 1. It has the interpretation of being the estimated treatment effect on coskewness, of being in the S&P 500 vs not being in the S&P 500. When this regression is applied to a narrow window around the size rank threshold, $P(SP500_{it})$ will be unrelated to any variables outside of the model, since it is obtained from a narrow sample where the differences in size rank are effectively random. While technically the fitted values $P(SP500_{it})$ may vary because they capture variation from the controls included in the first stage, this variation is controlled since I include those controls in the second stage, and effectively isolate variation in the outcome variable, coskewness, related to the random, discontinuous probability of meeting the size rank criteria.

Since I use a non-linear logit first stage regression with a linear OLS second stage regression, I also run an intermediate OLS regression where $SP500_{it}$ is regressed against the fitted values from the first stage $P(SP500_{it})$ plus controls, and use the fitted probability values from this intermediate stage regression in the second stage. This ensures consistent estimates as outlined in Wooldridge (2002) and Adams, Almeida, and Ferreira (2009). The 2SLS standard errors are heteroskedasticity adjusted and clustered at the firm level.

4.5.1 Main Results

I start by analyzing the fuzzy RDD in Table 4.6, and varying the window size for the sample between +/- 25 to +/- 100 on either size of the S&P 500 size rank inclusion threshold. I find a negative and significant effect of S&P 500 membership on coskewness with the S&P 500 index, which is robust across windows. This indicates that firms in the S&P 500 experience undesirable excess negative coskewness as a result of being in the index. This is counter to the previous event study results, but it is important to note that these samples are not limited to observations for firms within a short period after inclusion, instead, they are measuring differences amongst firms around the inclusion threshold regardless of their tenure within the index. They therefore capture long term effects of membership instead of simple short-term, potentially transitory effects like an event study. The difference in results indicates the importance of using alternative methods to measure index membership effects.

The existence of excess negative coskewness has several important implications. The first is that it provides further evidence of financialization impacting and distorting the return process of firms, particularly those included in popular indices. If there is no non-fundamental component introduced into the return process of S&P 500 firms, then we should see no excess coskewness. The fact that we do adds to growing evidence of such non-fundamental effects. This is in line with observed excess returns and comovement of firms in the S&P 500. It is also consistent with observations from other markets, such as increasing correlations amongst commodities and increasing correlations of commodities with other financial markets.

The next implication is that, if coskewness is priced, excess coskewness in a stock may represent a distortion to a stock's risk characteristics, which could feedback into expected returns and its cost of capital. While I do not directly analyze expected returns, such an indirect risk effect on expected returns could also help explain previously observed excess returns for S&P 500 firms, even without a first order effect of membership on expected returns directly.

Finally, there are implications for diversification. Results point to increased financialization as a driver of excess coskewness, and excess coskewness implies that assets decrease the coskewness of portfolios they are added to. Broadly speaking, as financialization decreases the coskewness of portfolios, those portfolios will offer less diversification during downswings, which is exactly when diversification is most needed. There is a long literature, particularly amongst international markets, looking into increased asset correlations during market downturns. The implication of excess negative coskewness, is that this trend of increased asset correlations during market downturns will worsen as financialization increases, and it also points to financialization being one of the reasons downside correlation has decreased. I do not evaluate this in the current paper, but it is an interesting area for further exploration.

4.6 Robustness of regression discontinuity

4.6.1 Liquidity

While I focus on a potential and unintended negative consequence of financialization on asset returns, it is important to remember that inclusion in a popular index and exposure to the investment flows of its investor can also bring beneficial increases in liquidity. To rule out results being driven by beneficial increases in liquidity, I repeat the analysis with several specifications designed to control for liquidity, and present the results in Table 4.7.

First, I introduce additional liquidity controls, specifically the Amihud (2002) illiquidity measure and the effective tick measure of Goyenko, Holden, Trzcinka (2009), which is a monthly proxy for effective spread. Results when including these liquidity measures are virtually unchanged from the main results.

Next, I use dimson adjusted coefficients (Dimson, 1979) for the coskewness measure to capture the impact of liquidity on the differential speed of information diffusion for firms within the S&P500, and once against the results are unchanged. In Figure 4.1, I also present the results of an event study on unexpected federal funds rate changes, tracking the return response of S&P 500 and non-S&P 500 firms during and after an announcement. I find the daily response is almost identical, ruling out significant effects of inclusion on the speed of information diffusion which could impact coskewness results.[5]

Finally, I combine both control methods with dimson adjusted coefficients for coskewness, plus liquidity controls, and find that results are robust. I conclude that results are not due to liquidity differences for S&P 500 firms.

4.6.2 Control Specification

I also evaluate the sensitivity of results to the specification of the regressions equations in the fuzzy RDD, and check to see if it successfully eliminates observable differences in firms on either size of the treatment threshold.

Results in Table 4.8 show that results are insensitive to the choice of the order of polynomial used in the polynomial spline I use to control any size rank differences between firms. I utilize 1st to 4th order polynomials and find unchanged coefficients. I ignore higher order polynomials as they have been criticized as being inappropriate (Gelman and Imbens, 2019). The insensitivity of these estimates point to the narrow window being fairly effective at controlling important differences in firms related to sizerank.

In Table 4.9, I present different control specifications to show how control choice affects the results. Results are generally the same, except for a regression removing all turnover controls, in which efficiency drops drastically, and the coefficient becomes insignificant from zero. If I include lagged return and size controls, efficiency improves somewhat. This points to the importance of controlling other sources of variation in firms around the threshold to improve efficiency of estimation, particularly when using coskewness, which is quite noisy relative to comovement, as an outcome variable.

[5] This is in line with research by Zebedee, Bentzen, Hansen, and Lunde (2008) and by Jubinkski and Tomljanovich (2013), which find that S&P 500 stocks, and CRSP stocks generally, seem to fully incorporate macro news within 15 minutes.

Finally, Table 4.10 shows the results of a covariate balance test, which tests for differences in the treated vs. untreated firms on either side of the treatment threshold. If, as intended, the regression discontinuity approach is successful and firms are not systematically different on either side of the threshold in ways that could drive results, then observable characteristics of the firms should be similar. I find that to be the case across a variety of firm characteristics. Thus, the firms analyzed do not seem to be systematically different, and my analysis captures the effect of being randomly treated.

4.6.3 Self-selection

A common concern in regression discontinuity designs is whether the subject being treated can control whether they are treated or not by manipulating the assignment variable, and therefore self-select into the treated or untreated groups leading to differences between the two. This is a concern if subjects can precisely control the variable used for assignment. A typical example of this, as outlined in McCrary (2008) is the study of outcomes of income-based support programs, where subjects can easily manipulate their earnings to qualify for treatment.

There are a few reasons this is not a concern in my setting. First, the assignment variable is size rank, and size rank is generated from market capitalization, which is determined by an evaluation of the firm by the market, which the firm cannot precisely control. Second, the explicit goal of all firms is to maximize their market capitalization, so to the extent that firms could manipulate their size rank, they should all be trying to maximize it, meaning a firms size rank is not only controlled by their actions and the evaluation of the market, but by other firms and the markets evaluation of them. Finally, treatment in the model is fuzzy, so simply meeting the size criteria does not guarantee treatment and vice versa, making it highly unlikely that firm managers would be motivated to, let alone be able to, precisely manipulate their size rank in a small window around the threshold for purposes of seeking or avoiding treatment.[6]

4.7 Conclusion

I present the first results showing the impact of S&P 500 membership on firm coskewness and find that the short-run transitory effect of being added to the S&P 500 is very different from the long-term effect experienced by member firms. In the short-run after inclusion, there appears to be a significant positive excess coskewness between the firm and the S&P 500. This is not surprising if we consider transitory pressures, like momentum, surrounding firms when added to the index. In the long run, S&P 500 member firms experience undesirable excess negative coskewness with the S&P 500, indicating that they move together more closely during market downturns than during market upswings. This negative excess coskewness points to price distortions caused by index membership and exposure to the large flows of passive investment, and habitat traders, that membership entails.

Additional implications of this distortion could include feedback into expected returns, firms' cost of capital, and reduced benefits to diversification. These implications point to avenues for future research. One is to look at how much inclusion directly affects S&P 500 returns versus any indirect effects it may have through changing risk characteristics. Another avenue to explore is whether financialization has contributed to changing downside correlations across assets, and therefore impacted benefits to diversification.

[6] Commonly used density tests such as McCrary (2008) are not useful given the design, as the size rank assignment variable is uniformly distributed.

Table 4.1: Full Sample Description

This table provides descriptive statistics for the main outcome variables, controls used in later analysis, and additional characteristics of interest for firms included and not-included in the S&P 500 between Jan. 1, 1995 and Dec. 31, 2017. The unit of observation is at the firm-month level. S&P 500 comovement is calculated using regression specification (1) with the influence of the firm in the regression removed from the S&P 500 return. Dimson adjusted coefficients sums the S&P 500 betas and gammas on a 3-day window centered on the day of interest to account for non-synchronous trading and delayed reaction due to liquidity. Illiquidity is the monthly measure from Amihud (2002). Eff. Tick is the effective spread proxy from Goyenko, Holden, and Trzcinka (2009). Factors umd, hml, smb, cma, rmw represent firms comovement over a given month with the respective factor portfolio listed. Lagged return variables are holding period returns. FinLiq is the ratio of cash and short term investments to current liabilities. Current is the ratio of current assets to current liabilities. Leverage is the ratio of liabilities to shareholder's equity. Non-ratio financial characteristics are in $millions.

	Non-S&P 500			*S&P 500*		
Variable	N	Mean	Std Dev	N	Mean	Std Dev
SP500 Beta	264057	1.03	0.51	136302	1.01	0.44
Dimson Beta	263746	1.104	0.608	136203	1.053	0.528
Coskew	264073	-2.092	15.179	136245	-0.677	10.659
Dimson Coskew	262921	-3.331	24.126	136169	-1.226	17.328
Illiquidity	275167	0.219	0.830	137803	0.020	0.543
Eff. Tick	275174	0.187	0.412	137807	0.089	0.185
1mo return	275028	0.009	0.148	137776	0.007	0.12
1yr return	270001	0.092	0.492	135866	0.078	0.481
3yr return	262395	0.286	1.313	133325	0.21	0.806
5yr return	246517	0.556	3.171	130137	0.374	1.472
Firmsize	275242	1550691	1596633	137890	23383822	43335298
Umd	275167	-0.078	2.038	137803	-0.085	1.534
Hml	275167	0.158	2.567	137803	0.084	1.957
Smb	275167	0.978	2.023	137803	0.176	1.427
Cma	275167	-0.028	3.219	137803	0.107	2.44
Rmw	275167	-0.001	0.029	137803	-0.001	0.022
Assets	274965	2828	5706	137785	57207	202123
Acquisitions	244407	61.75	254.79	112863	419.61	1407.9
Capx	251316	85.975	174.134	124897	1157.399	2733.404
Roa	256363	0.029	0.137	127004	0.051	0.084
Roe	256092	-0.022	0.401	126924	2.765	132.816
Roi	256078	0.046	0.344	127004	0.094	0.222
TobinQ	256104	1.878	1.305	126924	2.081	1.326
Sales	256576	1662	2754	127109	18740	38332
R&D	275242	21.143	51.775	137890	401.887	1405.468
FinLiq	230258	1.014	2.05	113376	0.596	0.918
Current	231721	2.557	2.337	113150	1.689	1.116
Leverage	255755	1.24	3.233	126647	476.317	22995.43
Div. Yield	273290	0.954	3.464	136869	12.422	31.845

Table 4.2: Average Coefficients for Univariate Inclusion Event Study

This table reports the average coefficients and the standard error of the coefficient estimates across regressions for stocks i

$$R_{i,t} = \alpha_i + \beta_{i1} POST_t + \beta_{i2} SP500_t + \beta_{i3} POST_t \times SP500_t + \gamma_{i1} SP500_t^2 + \gamma_{i2} POST_t \times SP500_t^2 + \epsilon_{it}$$

POST is a dummy variable equal to 1 if a return is in a 12 month window after stock i is added to the S&P 500 index, and 0 if the return is in the 12 months window before stock i is added to the index.SP500 is the return on the S&P 500 excluding the influence of stock i during the period it is within the index. Standard errors are in parentheses.

	Univariate Regression		
Sample Period	(1995-2017)	(1995-2002)	(2003-2017)
Constant	0.0014	0.0017	0.0012
	(0.0001)	(0.0002)	(0.0001)
POST	−0.0015	−0.0020	−0.0012
	(0.0001)	(0.0003)	(0.0001)
$SP500$	1.0613	1.0387	1.0784
	(0.0236)	(0.0451)	(0.0238)
$SP500^2$	−2.2277	−1.5355	−2.7485
	(0.5784)	(0.9579)	(0.7126)
$POST \times SP500$	0.1227	0.1926	0.0702
	(0.0193)	(0.0357)	(0.0201)
$POST \times SP500^2$	3.0259	4.0186	2.2790
	(0.7406)	(1.2709)	(0.8770)
Obs	531	228	303
R-Square	0.27	0.21	0.32
	(0.0064)	(0.0076)	(0.0085)

Table 4.3: Average Coefficients for Bivariate Inclusion Event Study

This table reports the average coefficients and the standard error of the coefficient estimates across regressions for stocks i

$$R_{i,t} = \alpha_i + \beta_{i1} POST + \beta_{i2} SP500 + \beta_{i3} POST \times SP500 + \beta_{i4} NonSP500 + \beta_{i5} POST \times NonSP500$$

$$+\gamma_{i1} SP500^2 + \gamma_{i2} POST \times SP500^2 + \gamma_{i3} NonSP500^2 + \gamma_{i4} POST \times NonSP500^2 + \epsilon_{it}$$

POST is a dummy variable equal to 1 if a return is in a 12 month window after stock i is added to the S&P 500 index, and 0 if the return is in the 12 months window before stock i is added to the index. SP500 and NonSP500 are the returns on the S&P 500 and those non-SP500 firms in the SP1500 excluding the influence of stock i during the period it is within the respective index at time t. Standard errors are in parentheses.

	Bivariate Regression		
Sample Period	(1995-2017)	(1995-2002)	(2003-2017)
Constant	0.0006	0.0003	0.0008
	(0.0001)	(0.0002)	(0.0001)
POST	−0.0012	−0.0016	−0.0009
	(0.0001)	(0.0003)	(0.0001)
$SP500$	0.1742	0.1380	0.2015
	(0.0259)	(0.0363)	(0.0362)
$SP500^2$	0.1425	1.7103	−1.0373
	(0.9696)	(1.1590)	(1.4563)
$POST \times SP500$	0.4505	0.4839	0.4254
	(0.0309)	(0.0457)	(0.0418)
$POST \times SP500^2$	3.3155	1.7613	4.4850
	(1.3878)	(1.5353)	(2.1399)
$NonSP500$	0.9163	1.0926	0.7836
	(0.0337)	(0.0571)	(0.0389)
$NonSP500^2$	−1.0601	−0.9663	−1.1307
	(0.7625)	(1.1088)	(1.0454)
$POST \times NonSP500$	−0.3488	−0.4006	−0.3098
	(0.0280)	(0.0443)	(0.0359)
$POST \times NonSP500^2$	−1.2417	0.7935	−2.7732
	(1.0470)	(1.5051)	(1.4394)
Obs	531	228	303
R-Square	0.30	0.25	0.34
	(0.0064)	(0.0078)	(0.0088)

Table 4.4: Pre vs Post S&P 500 inclusion coskewness estimates

This table reports the average coskewness coefficients and across the univariate regressions of each individual stock against the returns of S&P 500 and Non-S&P 500 stocks in Panel A:

$$R_{i,t} = \alpha_i + \beta_{i1} index_{jt} + \gamma_{i1} index_{jt}^2 + \epsilon_{it}$$

and the bivariate regressions in Panel B:

$$R_{i,t} = \alpha_i + \beta_{i1} SP500_t + \beta_{i2} NonSP500_t + \gamma_{i1} SP500_t^2 + +\gamma_{i2} NonSP500_t^2 + \epsilon_{it}$$

The regressions are run seperately on the pre-period including the 12 full calendar months before the month of inclusion, and the post-period including the full 12 calendar months after the month of inclusion. Significance of the difference estimates at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively. Standard errors are in parenthesis.

Univariate Regression								
		S&P500			*Non-S&P500*			*Diff*
Sample Period	*Obs*	$\gamma_{1,pre}$	$\gamma_{1,post}$	$\Delta\gamma_{1,sp}$	$\gamma_{1,pre}$	$\gamma_{1,post}$	$\Delta\gamma_{1,nsp}$	$\Delta\gamma_{1,sp}-\Delta\gamma_{1,nsp}$
1995-2017	531	−2.2276	0.7983	3.0259***	−1.3275	−0.3474	0.9801	2.0458***
		(0.5784)	(0.4840)	(0.7406)	(0.5299)	(0.5920)	(0.6435)	(0.5392)
1995-2002	228	−1.5354	2.4831	4.0185***	−0.0725	1.5447	1.6172	2.4013***
		(0.9579)	(0.8593)	(1.2709)	(0.8479)	(0.8648)	(1.1325)	(0.8910)
2003-2017	303	−2.7485	−0.4695	2.279*	−2.2720	−1.7950	0.4770	2.2313**
		(0.7126)	(0.5389)	(0.8770)	(0.5299)	(0.5920)	(0.7342)	(0.6654)

Bivariate Regression								
		S&P500			*Non-S&P500*			*Diff*
Sample Period	*Obs*	$\gamma_{1,pre}$	$\gamma_{1,post}$	$\Delta\gamma_1$	$\gamma_{2,pre}$	$\gamma_{2,post}$	$\Delta\gamma_2$	$\Delta\gamma_1-\Delta\gamma_2$
1995-2017	531	0.1425	3.7274**	3.5849	−1.0601	−2.2826	−1.2225	4.8074*
		(0.9696)	(0.7714)	(1.401)	(0.7625)	(0.7714)	(1.0569)	(2.3133)
1995-2002	228	1.7103	3.4717	1.7614	−0.9663	−0.1728	0.7935	0.9679
		(1.1590)	(1.0918)	(1.5353)	(1.1087)	(1.0469)	(1.5051)	(2.7449)
2003-2017	303	−1.0373	3.9231	4.9604**	−1.1307	−3.8968	−2.7661	7.7265**
		(1.4563)	(1.6624)	(2.1756)	(1.0454)	(1.0934)	(1.4635)	(3.4973)

Table 4.5: Logistic Regression: Probability a firm is in the S&P 500

This table reports the results of the logistic regression outlined in equation (4.4), where the outcome variable is equal to 1 if a firm is in the S&P 500 and 0 if it is not. Above is the indicator variable of interest $\tau 1[S_{it} \geq \overline{S_t}]$ representing whether the firm sizerank is greater than the threshold for inclusion. The coefficient captures the discontinuous change in probability of treatment (S&P 500 membership) around the inclusion threshold. Unreported controls include a cubic polynomial spline on sizerank plus Year, Industry and Year X Industry fixed effects. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

Variable	Estimate	Chi Square
Above	0.1152***	8.1658
1mo Return	-0.4122***	43.1455
1yr Return	-0.5421***	608.3394
3yr Return	-0.4710***	1327.4953
5yr Return	-0.0746***	232.6802
UMD	-0.0609***	227.0554
HML	0.0427***	144.0431
SMB	-0.0718***	275.5307
RMW	1.2986***	20.2192
CMA	0.0441***	266.9084
Firmsize	1.3293***	274.6296
Assets	0.000085***	3721.1112
TobinQ	-0.1140***	317.3258
Illiquidity	-6.1599***	69.8409
Eff.Tick	81.5374***	654.7065

Table 4.6: Fuzzy RD Effect Estimates across Window Size

This table reports the estimated treatment effect of SP500 membership based on the fuzzy regression discontinuity design (RDD) specification in equation (4.5). P(500) is a firms propensity score estimated from the first stage in equation (4.4) as described in the text. The dependent variable is the estimated coskewness between the stock and the SP500 index from equation (4.1). Specifications differ in window size around the SP500 size based cutoff point. Controls include 1mo, 1yr, 3yr, 5yr lagged returns, factor exposures to hml, smb, cma, rmw, umd, firmsize, assets, and tobinsq, illiquidity, and effective tick. Two stage least squares standard errors are heteroskedasticity adjusted and clustered at the firm level. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively. P-values are in parentheses.

Specification	(1)	(2)	(3)
Dep. Var	Coskew.	Coskew.	Coskew.
Window	25	50	100
Spline	Cubic	Cubic	Cubic
P(SP500)	-2.846^{**}	-2.361^{**}	-2.720^{***}
	(0.027)	(0.026)	(0.004)
Controls	Y	Y	Y
Year FE	Y	Y	Y
Industry FE	Y	Y	Y
Year $\times$ Industry FE	Y	Y	Y
Obs	11980	23681	46610
R-Square	0.15	0.13	0.13

Table 4.7: Fuzzy RD Effect Estimates across Liquidity Controls

This table reports the estimated treatment effect of SP500 membership based on the fuzzy regression discontinuity design (RDD) specification in equation (4.5). P(500) is a firms propensity score estimated from the first stage in equation (4.4) as described in the text. The dependent variable is the estimated coskewness between the stock and the SP500 index from equation (4.1). The specification differs on dependent variable and liquidity controls. Specification (1) excludes liquidity controls. Specification (2) is the same as reported in table (5) with liquidity controls. Specification (3) uses a dimson adjusted skewness without liquidity controls. Specification (4) uses dimson adjusted skewness with liquidity controls. Controls include 1mo, 1yr, 3yr, 5yr lagged returns, factor exposures to hml, smb, cma, rmw, mom, firmsize, assets, and tobinsq. Two stage least squares standard errors are heteroskedasticity adjusted and clustered at the firm level. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively. P-values are in parentheses.

Specification	(1)	(2)	(3)	(4)
Dep. Var	Coskew.	Coskew.	Dimson	Dimson
Window	50	50	50	50
Spline	Cubic	Cubic	Cubic	Cubic
P(SP500)	-2.267^{**}	-2.992^{***}	-3.078^{*}	-3.702^{**}
	(0.026)	(0.005)	(0.065)	(0.044)
Illiquidity	3.641		-0.559	
	(0.358)		(0.307)	
Eff. Tick	-286.514		-275.517	
	(0.143)		(0.307)	
Controls	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Industry FE	Y	Y	Y	Y
Year × Industry FE	Y	Y	Y	Y
Obs	23681	23681	23664	23664
R-Square	0.13	0.13	0.11	0.11

Table 4.8: Fuzzy RD Effect Estimate across Spline

This table reports the estimated treatment effect of SP500 membership based on the fuzzy regression discontinuity design (RDD) specification in equation (4.5). P(500) is a firms propensity score estimated from the first stage in equation (4.4) as described in the text. The dependent variable is the estimated coskewness between the stock and the SP500 index from equation (4.1). The specifications differ by the order of the polynomial spline used as a control function. Controls include 1mo, 1yr, 3yr, 5yr lagged returns, factor exposures to hml, smb, cma, rmw, umd, firmsize, assets, and tobinsq, illiquidity, and effective tick. Two stage least squares standard errors are heteroskedasticity adjusted and clustered at the firm level. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively. P-values are in parentheses.

Specification	(1)	(2)	(3)	(4)
Dep. Var	Coskew.	Coskew.	Coskew.	Coskew.
Window	50	50	50	50
Spline	Linear	Quadratic	Cubic	Quartic
P(SP500)	-2.325^{**}	-2.370^{**}	-2.361^{**}	-2.356^{**}
	(0.028)	(0.025)	(0.026)	(0.026)
Controls	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Industry FE	Y	Y	Y	Y
Year $\times$ Industry FE	Y	Y	Y	Y
Obs	23687	23681	23687	23687
R-Square	0.13	0.13	0.13	0.13

Table 4.9: Fuzzy RD Effect Estimates over Control Specification

This table reports the estimated treatment effect of SP500 membership based on the fuzzy regression discontinuity design (RDD) specification in equation (4.5). P(500) is a firms propensity score estimated from the first stage in equation (4.4) as described in the text. The dependent variable is the estimated coskewness between the stock and the SP500 index from equation (4.1). The specifications differ in included turnover controls. Controls include 1mo, 1yr, 3yr, 5yr lagged returns, factor exposures to hml, smb, cma, rmw, umd, firmsize, assets, and tobinsq, illiquidity, and effective tick. Two stage least squares standard errors are heteroskedasticity adjusted and clustered at the firm level. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively. P-values are in parentheses.

Specification	(1)	(2)	(3)	(4)
Dep. Var	Coskew.	Coskew.	Coskew.	Coskew.
Window	50	50	50	50
Spline	Cubic	Cubic	Cubic	Cubic
P(SP500)	-2.361^{**}	0.655	-2.850^{***}	-1.768
	(0.026)	(0.925)	(0.021)	(0.275)
Controls	Y	N	Y	Ret & Size
Year FE	Y	Y	Y	Y
Industry FE	Y	Y	Y	Y
Year $\times$ Industry FE	Y	Y	N	Y
Obs	23681	27229	23681	25631
R-Square	0.13	0.12	0.05	0.13

Table 4.10: Fuzzy RD Covariate Balance Test

This table reports the results of a covariate balance test which includes various firm characteristics as the outcome variable in the same fuzzy RD specification used to generate main results from equation (4.5). Estimates represent the difference in the listed variable for S&P 500 and non-S&P 500 firms in the +/- 50 discontinuity sample. The polynomial spline is cubic and controls include 1mo, 1yr, 3yr, 5yr lagged returns, factor exposures to hml, smb, cma, rmw, umd, firmsize, assets, and tobinsq, illiquidity, and effective tick. Two stage least squares standard errors are heteroskedasticity adjusted and clustered at the firm level. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

Variable	Estimate	P-Value
ROE	0.024	(0.497)
ROI	-0.026	(0.359)
Sales	438.675	(0.928)
Capx	56.488	(0.828)
R&D	24.107	(0.460)
Acquisitions	285.957	(0.121)
Dividend Yield	2.511	(0.243)
Fin. Liquidity	-0.099	(0.695)
Current	-0.327	(0.287)
Leverage	1.209	(0.209)

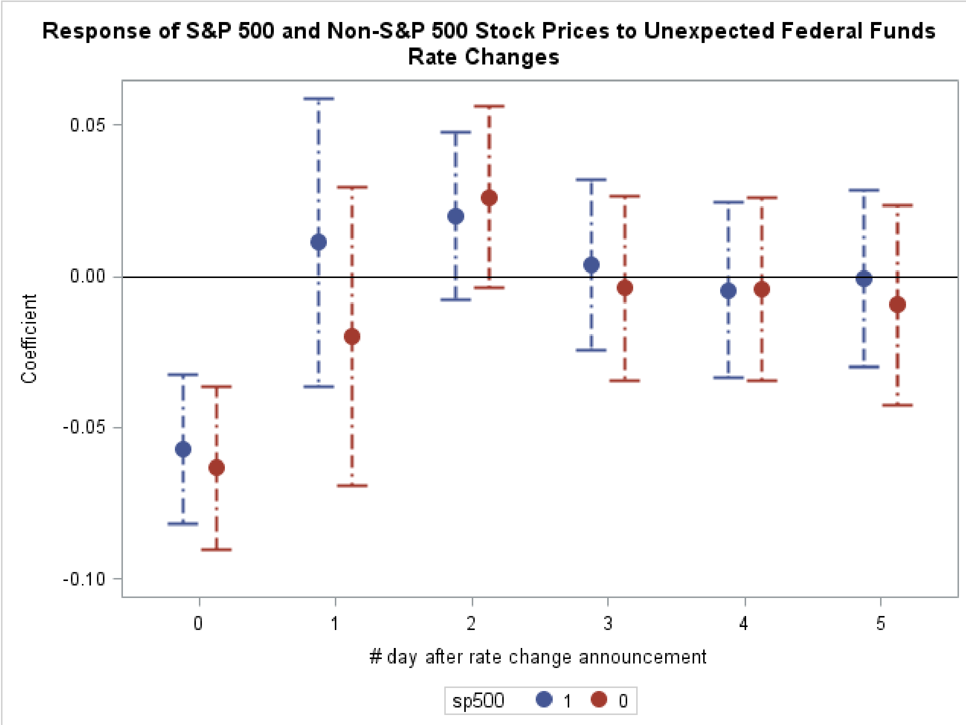

(a)

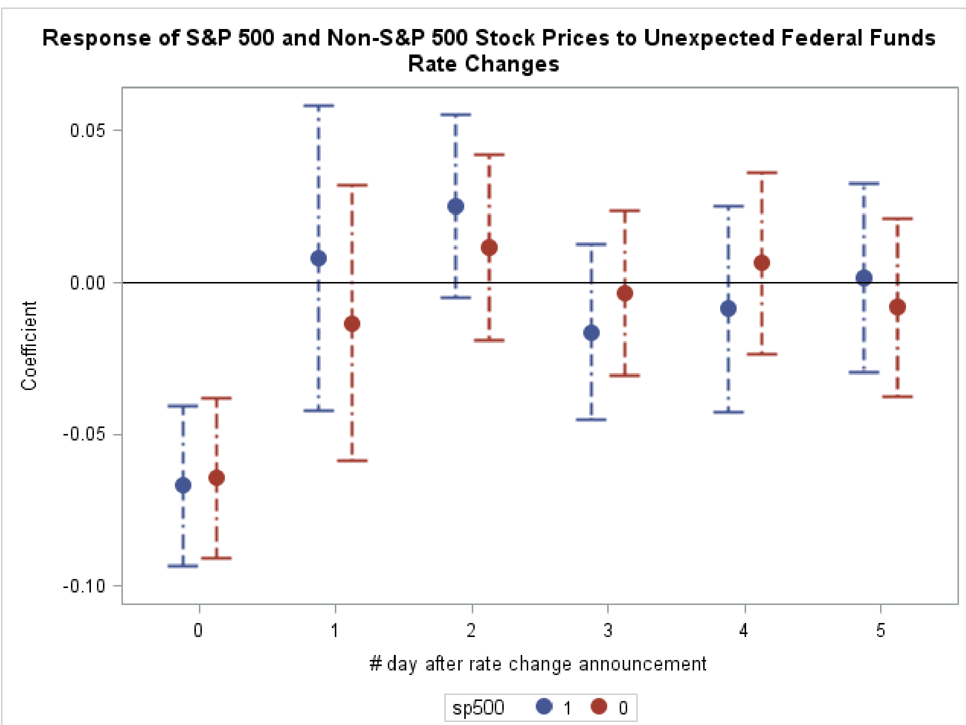

(b)

Figure 4.1: Stock price reaction to FOMC announcements

Displays point estimates and 95% confidence intervals for the stock price reaction of S&P 500 and non-S&P 500 stocks to surprise FOMC announcements about federal funds rate changes on the announcement date and over the post-announcement window. Index returns are calculated as a value weighted average of (a) the full sample of S&P 1500 firms (b) the +/- 50 sample of firms around the S&P 500 size rank cutoff.

REFERENCES

Acharya, V. V., Lochstoer, L. A., Ramadorai, T., 2013. Limits to arbitrage and hedging: Evidence from commodity markets. Journal of Financial Economics 109, 441–465.

Adams, R., Almeida, H., Ferreira, D., 2009. Understanding the relationship between founder-ceos and firm performance. Journal of Empirical Finance 16(1), 136–150.

Alquist, R., Kilian, L., 2010. What do we learn from the price of crude oil futures? Journal of Applied Econometrics 25, 539–573.

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

Ang, A., Chen, J., 2002. Asymmetric correlations of equity portfolios. Journal of Financial Economics 63, 443–494.

Antolin-Diaz, J., Rubio-Ramirez, J., 2018. Narrative sign restrictions for svars. American Economic Review 108, 2802–29.

Appel, I., Gormley, T., Keim, D., 2019. Identification using russell 1000/2000 index assignments: A discussion of methodologies. Available at SSRN 2641548 .

Arrow, K. J., 1971. Essays in the theory of risk-bearing. Markham Publishing Company, Chicago.

Barberis, N., Shleifer, A., Wurgler, J., 2005. Comovement. Journal of Financial Economics 75(2), 283–317.

Bernanke, B. S., Kuttner, K. N., 2005. What explains the stock market's reaction to federal reserve policy? The Journal of Finance 60(3), 1221–1257.

Bond, P., Edmans, A., Goldstein, I., 2012. The real effects of financial markets. Annual Review of Financial Economics 4, 339–360.

Bouchouev, I., 2020. From risk bearing to propheteering. Quantitative Finance 20(6), 887–894.

Brennan, M., 1991. The cost of convenience and the pricing of commodity contingent claims. In: Lund, D. Oksendal, B. (ed.), *Stochastic models and option values : applications to resources, environment, and investment problems*, Elsevier Science Pub., North Holland, NY.

Brogaard, J., Ringgenberg, M. C., Sovich, D., 2019. The economic impact of index investing. The Review of Financial Studies 32, 3461–3499.

Cabellero, R., Farhi, E., Gourinchas, P., 2008. Financial crash, commodity prices, and global imbalances. Brookings Papers on Economic Activity pp. 1–55.

Carhart, M., 1997. On persistence in mutual fund performance. The Journal of Finance 52(1), 57–82.

Cellini, S. R., Ferreira, F., Rothstein, J., 2010. The value of school facility investments: Evidence from a dynamic regression discontinuity design. The Quarterly Journal of Economics 125(1), 215–261.

Chen, H., Noronha, G., Singal, V., 2004. The price response to s&p 500 index additions and deletions: Evidence of asymmetry and a new explanation. The Journal of Finance 59(4), 1901–1930.

Chen, H., Singal, V., Whitelaw, R. F., 2016. Comovement revisited. Journal of Financial Economics 121(3), 624–644.

Cheng, H., Xiong, W., 2014. Financialization of commodity markets. Annual Review of Financial Economics 6, 419–41.

Chordia, T., Goyal, A., Sadka, G., Sadka, R., Shivakumar, L., 2009. Liquidity and the post-earnings-announcement drift. Financial Analysts Journal 65(4), 18–32.

Cunat, V., Gine, M., Guadalupe, M., 2012. The vote is cast: The effect of corporate governance on shareholder value. The Journal of Finance 67(5), 1943–1977.

Dessaint, O., Foucault, T., Frésard, L., Matray, A., 2019. Noisy stock prices and corporate investment. The Review of Financial Studies 32, 2625–2672.

Dimson, E., 1979. Risk measurement when shares are subject to infrequent trading. Journal of Financial Economics 7(2), 197–226.

Dolatabadi, S., Nielsen, M. Ø., Xu, K., 2015. A fractionally cointegrated var analysis of price discovery in commodity futures markets. Journal of Futures Markets 35, 339–356.

Fama, E. F., French, K. R., 1988. Business cycles and the behavior of metals prices. The Journal of Finance 43, 1075–1093.

Fama, E. F., French, K. R., 2015. A five-factor asset pricing model. Journal of Financial Economics 116(1), 1–22.

Fattouh, B., Kilian, L., Mahadeva, L., 2013. The role of speculation in oil markets: What have we learned so far? Energy Journal 34, 7–33.

Figuerola-Ferretti, I., Gonzalo, J., 2010. Modelling and measuring price discovery in commodity markets. Journal of Econometrics 158, 95–107.

Froot, K. A., Dabora, E. M., 1999. How are stock prices affected by the location of trade? Journal of Financial Economics 53(2), 189–216.

Gelman, A., Imbens, G., 2019. Why high-order polynomials should not be used in regression discontinuity designs. Journal of Business Economic Statistics 37(3), 447–456.

Gorton, G. B., Hayashi, F., Rouwenhorst, K. G., 2013. The fundamentals of commodity futures returns. Review of Finance 17, 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.

Greenwood, R., 2008. Excess comovement of stock returns: Evidence from cross-sectional variation in nikkei 225 weights. The Review of Financial Studies 21, 1153–1186.

Hamilton, J. D., 2009. Causes and Consequences of the Oil Shock of 2007-08 (No. w15002). National Bureau of Economic Research.

Hamilton, J. D., Wu, J. C., 2014. Risk premia in crude oil futures prices. Journal of International Money and Finance 42, 9–37.

Hamilton, J. D., Wu, J. C., 2015. Effects of index-fund investing on commodity futures prices. International Economic Review 56, 187–205.

Harford, J., Kaul, A., 2005. Correlated order flow: Pervasiveness, sources, and pricing effects. Journal of Financial and Quantitative Analysis 40, 29–55.

Harvey, C. R., Siddique, A., 2000. Conditional skewness in asset pricing tests. The Journal of Finance 55, 1263–1295.

Henderson, B. J., Pearson, N. D., Wang, L., 2015. New evidence on the financialization of commodity markets. Review of Financial Studies 28, 1285–1311.

Inoue, A., Kilian, L., 2013. Inference on impulse response functions in structural var models. Journal of Econometrics 177, 1–13.

Irwin, S. H., Sanders, D. R., 2012. Testing the masters hypothesis in commodity futures markets. Energy Economics 34(1), 256–269.

Irwin, S. H., Sanders, D. R., Merrin, R. P., 2009. Devil or angel? the role of speculation in the recent commodity price boom (and bust). Journal of Agricultural and Applied Economics 41, 377–391.

Jain, P., 1987. The effect on stock price of inclusion in or exclusion from the s&p 500. Financial Analysts Journal 43(1), 58–65.

Jordà, O., 2005. Estimation and inference of impulse responses by local projections. American Economic Review 95, 161–182.

Jubinski, D., Tomljanovich, M., 2013. Do fomc minutes matter to markets? an intraday analysis of fomc minutes releases on individual equity volatility and returns. Review of Financial Economics 22(3), 86–97.

Kaufmann, R. K., Ullman, B., 2009. Oil prices, speculation, and fundamentals: Interpreting causal relations among spot and futures prices. Energy Economics 31, 550–558.

Kaul, A., Mehrotra, V., Morck, R., 2000. Demand curves for stocks do slope down: New evidence from an index weights adjustment. The Journal of Finance 55(2), 893–912.

Kilian, L., 2009. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. American Economic Review 99, 1053–69.

Kilian, L., Murphy, D. P., 2014. The role of inventories and speculative trading in the global market for crude oil. Journal of Applied Econometrics 29(3), 454–478.

Kilian, L., Zhou, X., 2018. Modeling fluctuations in the global demand for commodities. Journal of International Money and Finance 88, 54–78.

Kilian, L., Zhou, X., 2020. The econometrics of oil market var models. Working paper, FRB of Dallas, available at SSRN `https://ssrn.com/abstract=3592429`.

Knittel, C. R., Pindyck, R. S., 2016. The simple economics of commodity price speculation. American Economic Journal: Macroeconomics 8, 85–110.

Kraus, A. Litzenberger, R., 1976. Skewness preference and the valuation of risk assets. The Journal of Finance 31(4), 1085–1100.

Lee, D. S., Lemieux, T., 2010. Regression discontinuity designs in economics. Journal of Economic Literature 48(2), 281–355.

Lombardi, M. J., Van Robays, I., 2011. Do financial investors destabilize the oil price? Working paper series 1346, ECB, available at SSRN `https://ssrn.com/abstract=1847503`.

Longin, F., Solnik, B., 2001. Extreme correlation of international equity markets. The journal of Finance 56, 649–676.

Madhavan, A., 2003. The russell reconstitution effect. Financial Analysts Journal 59(4), 51–64.

Masters, M. W., 2008. Testimony of michael w. masters, managing member/portfolio manager, masters capital management, llc,. before the Committee on Homeland Security and Governmental Affairs, United States Senate 110th Cong., 2nd sess.

Ng, V. K., Pirrong, S. C., 1994. Fundamentals and volatility: Storage, spreads, and the dynamics of metals prices. Journal of Business pp. 203–230.

Peri, M., Baldi, L., Vandone, D., 2013. Price discovery in commodity markets. Applied Economics Letters 20, 397–403.

Pindyck, R. S., 1994. Inventories and the short-run dynamics of commodity prices. RAND Journal of Economics 25, 141–159.

Pindyck, R. S., Rotemberg, J. J., 1990. The excess co-movement of commodity prices. Economic Journal (London) 100(403), 1173–1189.

Pindyck, R. S., Rotemberg, J. J., 1993. The comovement of stock prices. The Quarterly Journal of Economics 108(4), 1073–1104.

Rubio-Ramirez, J. F., Waggoner, D. F., Zha, T., 2010. Structural vector autoregressions: Theory of identification and algorithms for inference. The Review of Economic Studies 77, 665–696.

Sanders, D. R., Irwin, S. H., 2010. A speculative bubble in commodity futures prices? cross-sectional evidence. Agricultural Economics 41(1), 25–32.

Schwarz, T. V., Szakmary, A. C., 1994. Price discovery in petroleum markets: Arbitrage, cointegration, and the time interval of analysis. Journal of Futures Markets 14.

Shleifer, A., 1986. Do demand curves for stocks slope down? The Journal of Finance 41(3), 579–590.

Silvapulle, P., Moosa, I. A., 1999. The relationship between spot and futures prices: evidence from the crude oil market. Journal of Futures Markets 19, 175–193.

Singleton, K. J., 2014. Investor flows and the 2008 boom/bust in oil prices. Management Science 60, 300–318.

Sockin, M., Xiong, W., 2015. Informational frictions and commodity markets. The Journal of Finance 70(5), 2063–2098.

Stoll, H. R., Whaley, R. E., 2010. Commodity index investing and commodity futures prices. Journal of Applied Finance 20, 7–47.

Tang, K., Xiong, W., 2012. Index investment and the financialization of commodities. Financial Analysts Journal 68(6), 54–74.

Wooldridge, J., 2002. Econometric analysis of cross section and panel data, MIT Press, Cambridge, MA, p. 623.

Wurgler, J., Zhuravskaya, E., 2002. Does arbitrage flatten demand curves for stocks? The Journal of Business 75, 583–608.

Zebedee, A. A., Bentzen, E., Hansen, P. R., Lunde, A., 2008. The greenspan years: an analysis of the magnitude and speed of the equity market response to fomc announcements. Financial Markets and Portfolio Management 22(1), 3–20.

Appendix A

Additional Tables and Figures

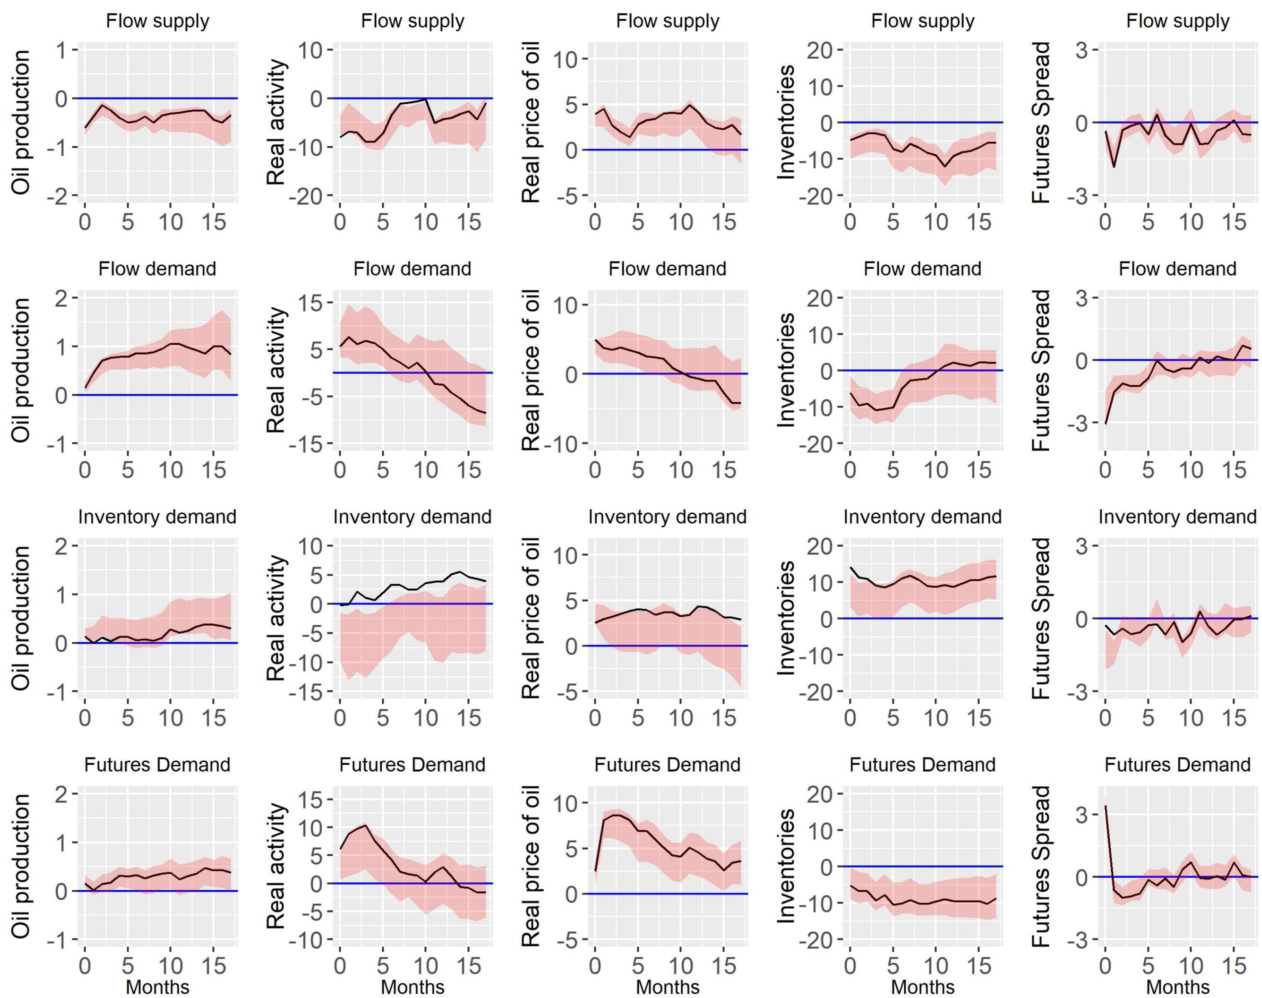

Figure A.1: Structural IRFs: No narrative restrictions

Structural Impulse Response Functions showing the response of each variable to a one standard deviation innovation to each structural shock. Responses are the cumulative % change for production, real activity, and the spot price, and cumulative level change for inventories. The Spread response is the difference in the futures and spot responses. The red band illustrates the 68% error band from the posterior distribution of the IRF's. Obtained as described in section 3, and in Appendix B. Narrative restrictions as described in section 3 are relaxed.

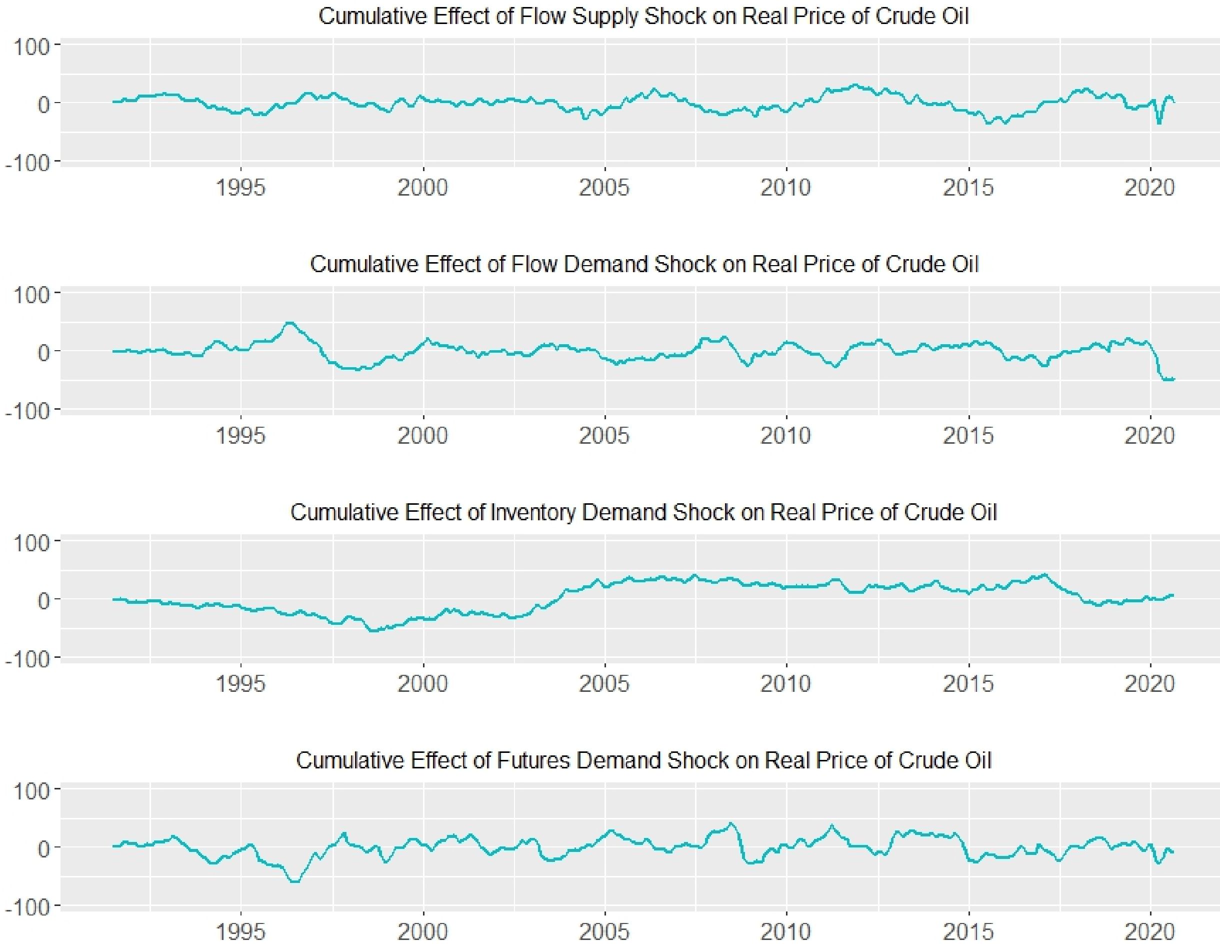

Figure A.2: Historical Decomposition: No narrative restrictions

Historical decomposition of the real spot price of brent oil from July 1991 to September 2020 showing the cumulative percentage change in spot price due to flow supply, flow demand, inventory demand, and futures demand shocks, respectively. Narrative restrictions as described in section 3 are relaxed.

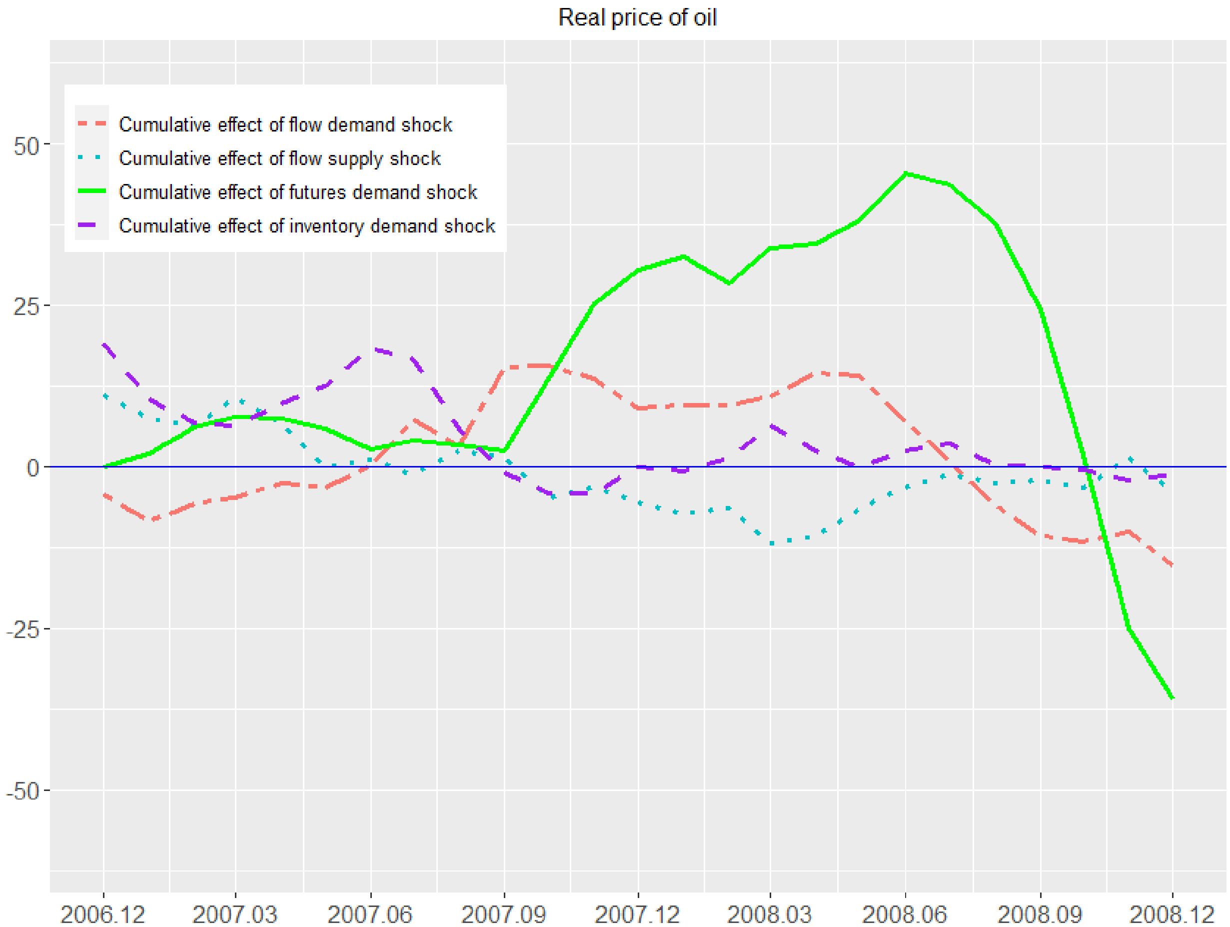

Figure A.3: Historical Decomposition-2008 Financial Crisis: No narrative restrictions

Historical decomposition of the real spot price of brent oil showing cumulative percentage change in spot price due to flow demand and futures demand shocks, respectively, between Jan. 2007 and Jan. 2009. Narrative restrictions as described in section 3 are relaxed.

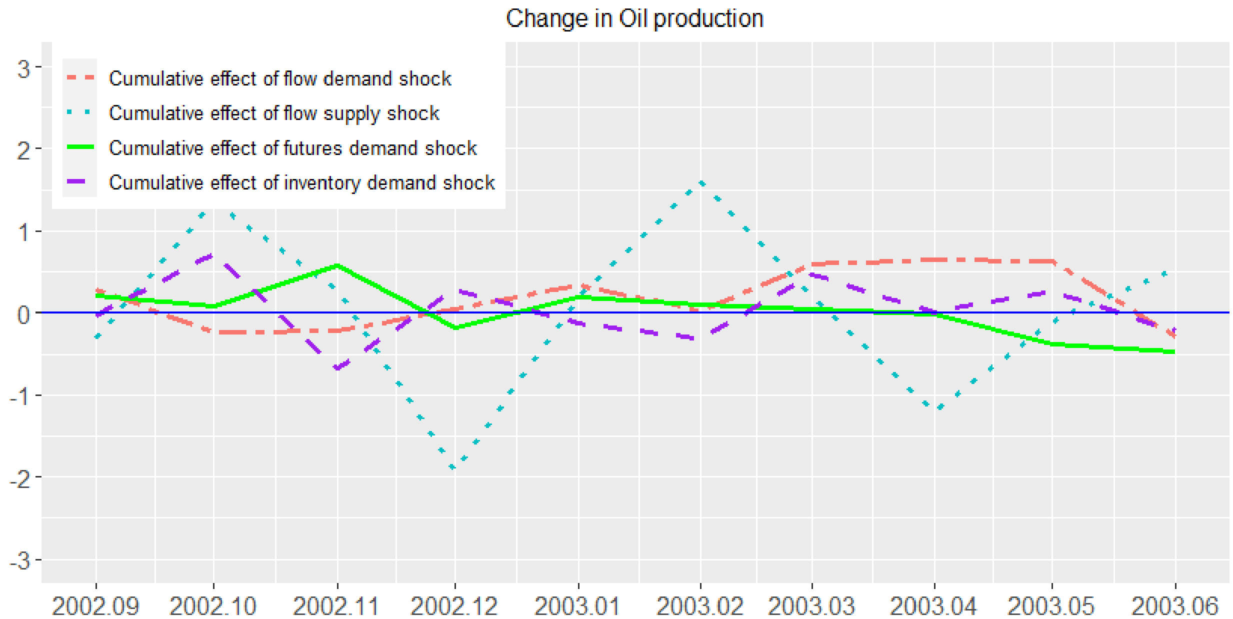

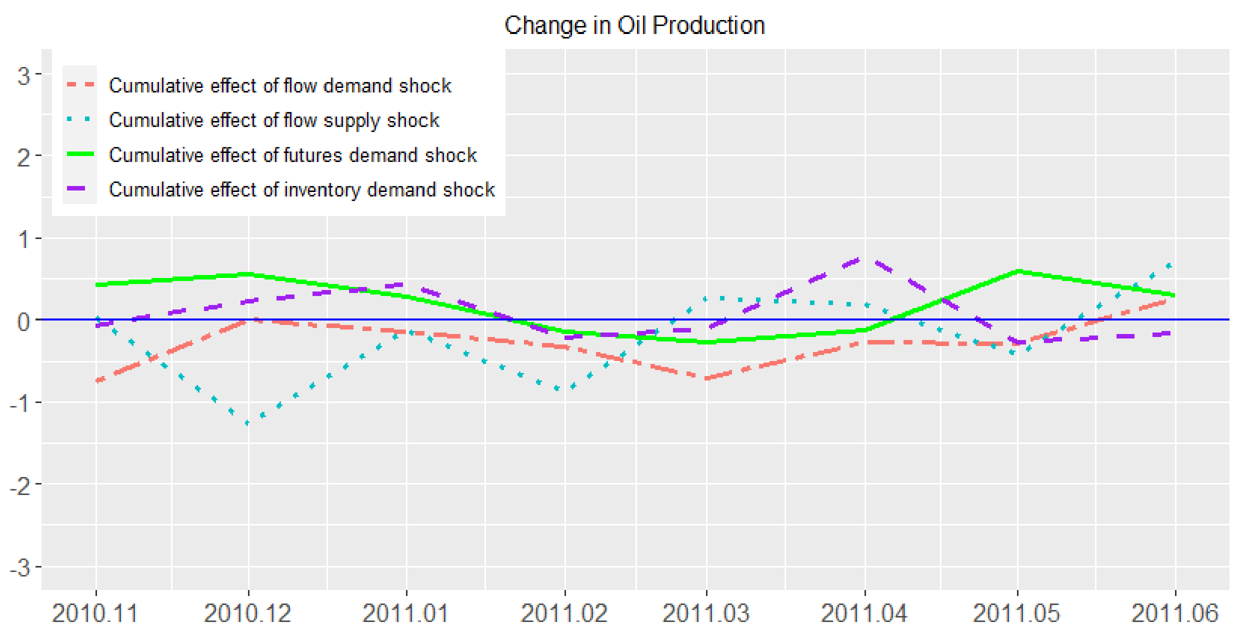

Figure A.4: Historical decomposition of supply shocks: No narrative restrictions

Historical decomposition of global oil production, showing cumulative percentage change in global production due to each shock during the period of the Venezualan Oil Strike (Dec. 2002), Invasion of Iraq (Mar.-Apr. 2003), and Libyan Civil War (Feb. 2011). Narrative restrictions as described in section 3 are relaxed.

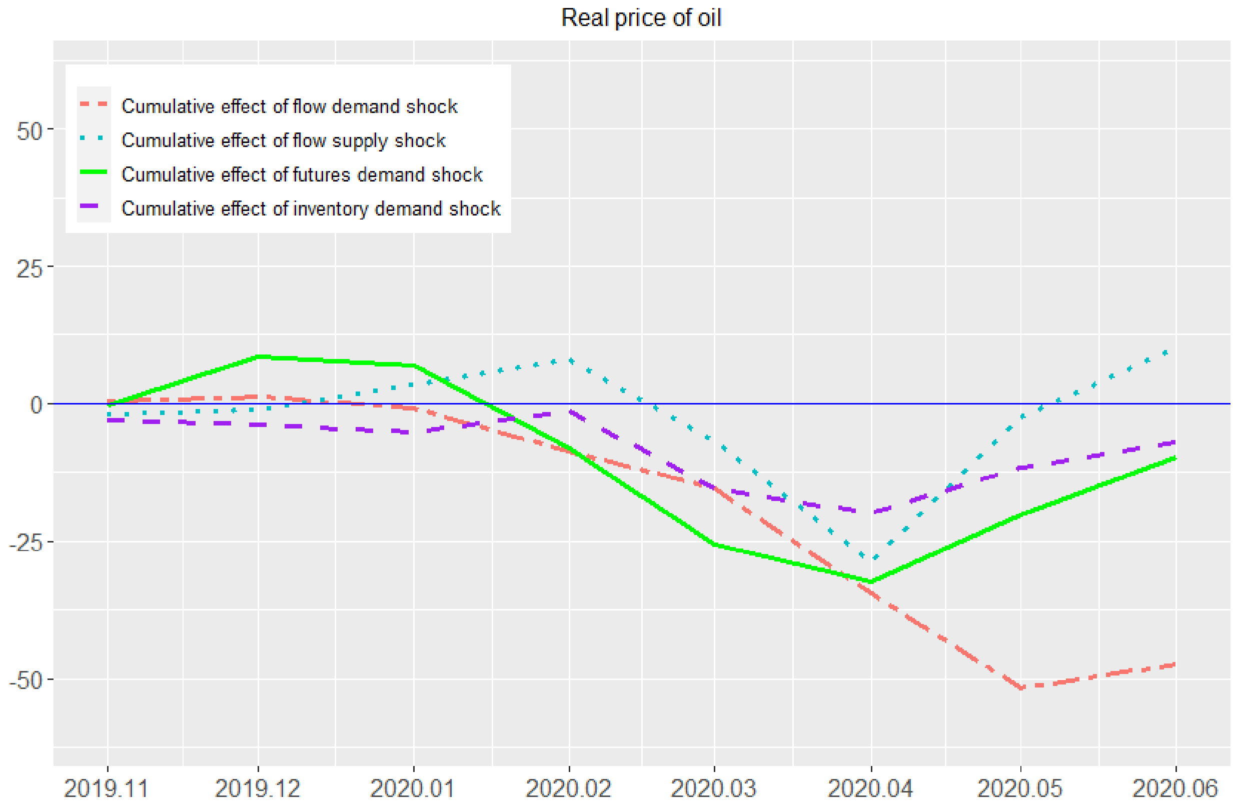

Figure A.5: Historical decomposition Covid-19 Outbreak: No narrative restrictions

Historical decomposition of the real spot price of brent oil showing cumulative percentage change in spot price due to each shock during the onset of the Global Covid-19 Pandemic and the Saudi-Russian Oil price war. Narrative restrictions as described in section 3 are relaxed.

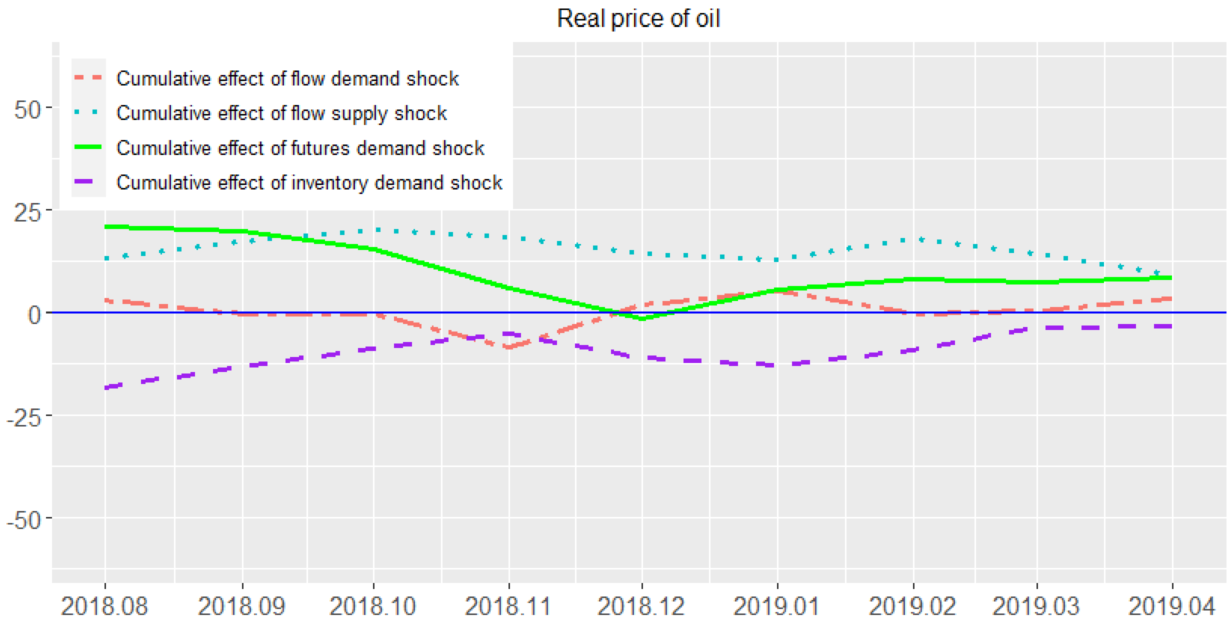

Figure A.6: Historical decomposition Aug. 2018-Apr. 2019

Historical decomposition of the real spot price of brent oil showing cumulative percentage change in spot price due to each shock around the assassination of Jamal Khashoggi on Oct 2nd, 2018.

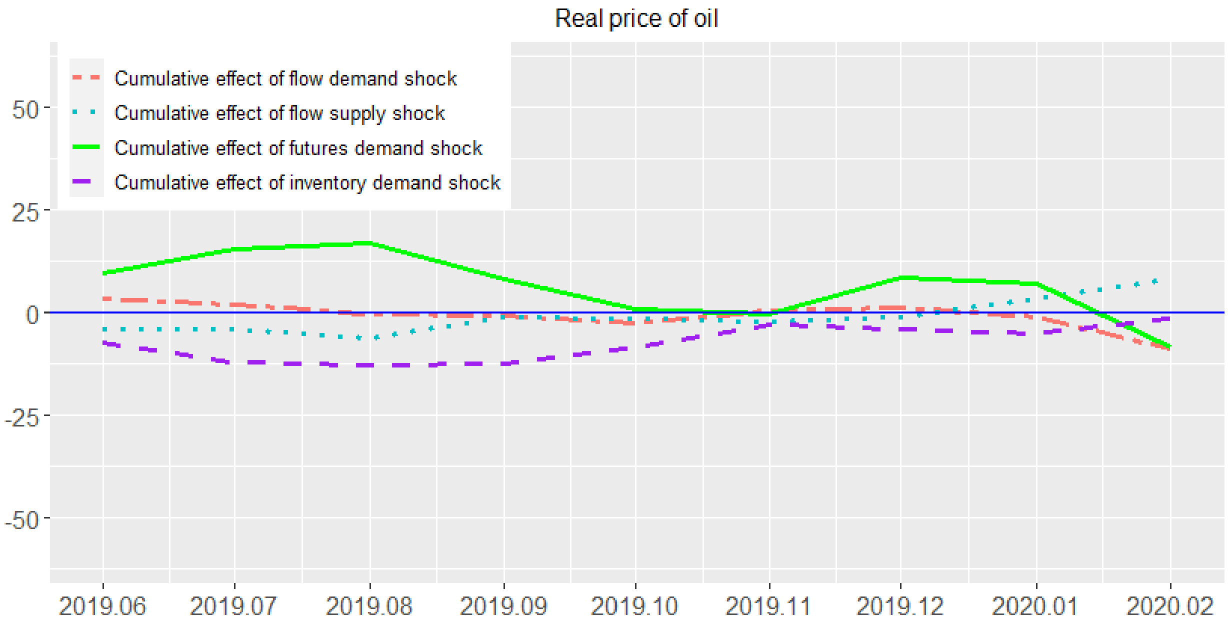

Figure A.7: Historical decomposition Jun. 2019-Feb. 2020

Historical decomposition of the global oil production, showing the cumulative percentage change in global production due to each shock around the Sept. 14 2019 Abqaiq–Khurais drone attacks on Saudi oil facilities.

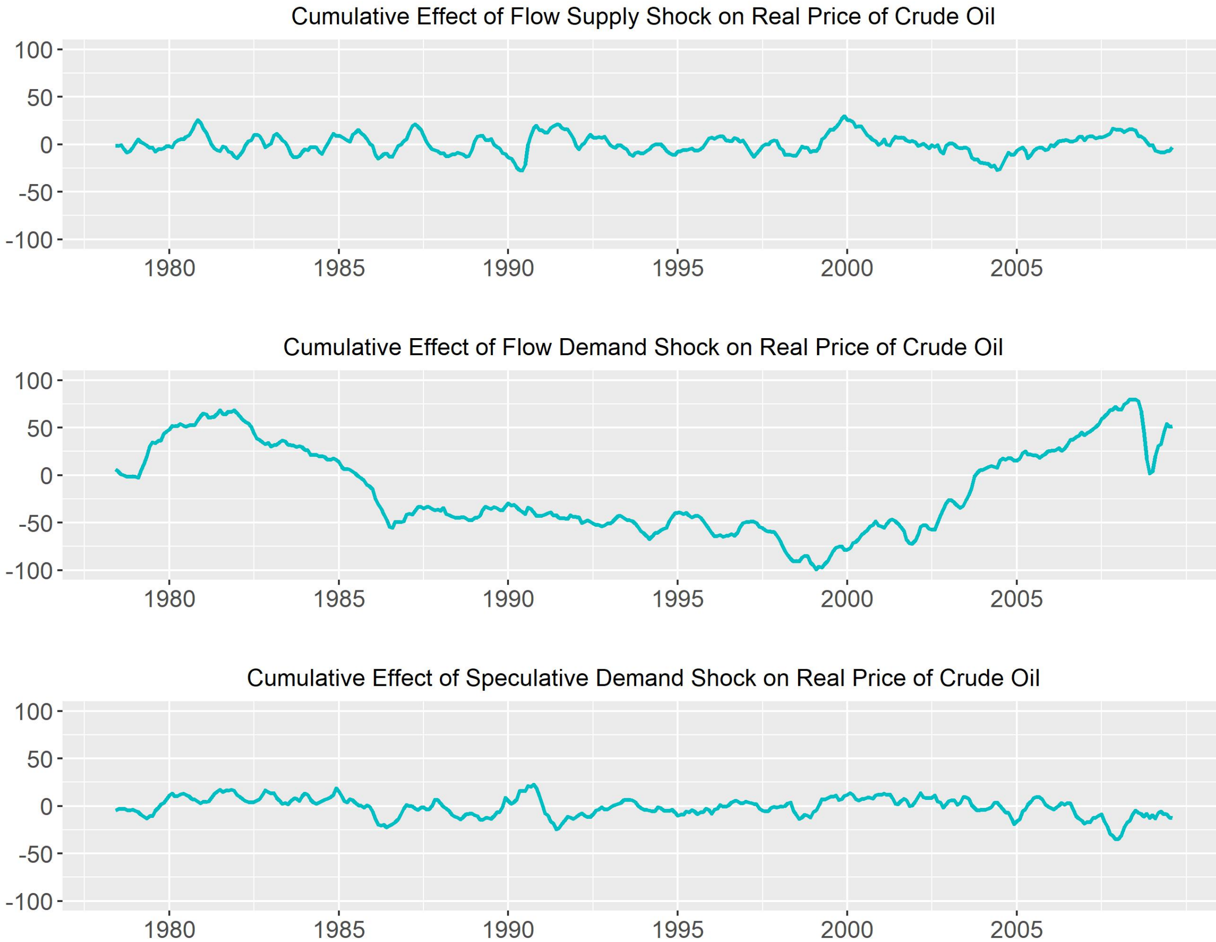

Figure A.8: Kilian and Murphy Replication

Historical decomposition of the real spot price of brent oil from February 1978 to August 2009 showing the cumulative percentage change in spot price due to flow supply, flow demand shocks, and inventory demand shocks respectively. Estimated by replicating Kilian and Murphy (2014) using their original variables, sample period, and restrictions.

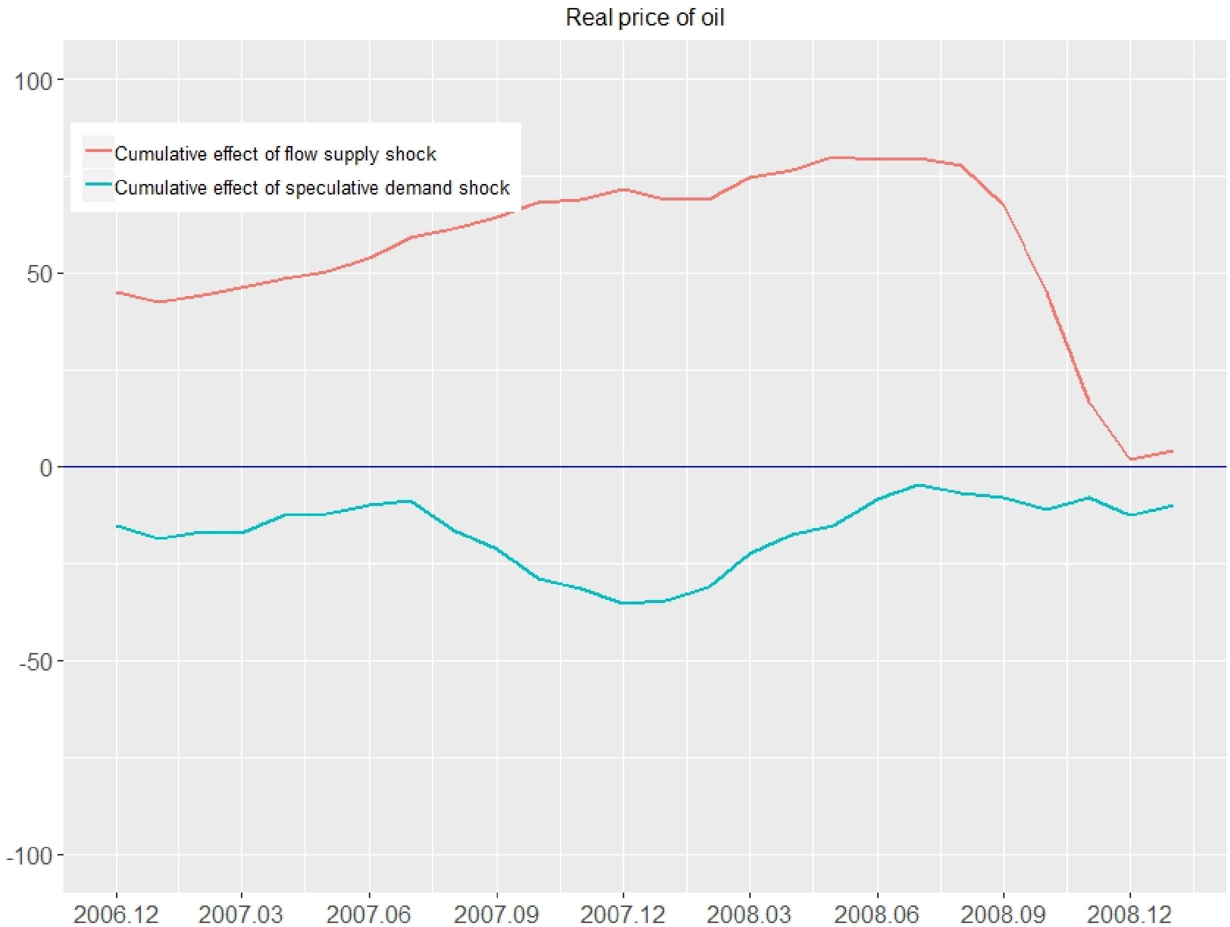

Figure A.9: Kilian and Murphy Replication-2008 Financial Crisis

Historical decomposition of the real spot price of brent oil from Jan. 2007 to Dec. 2008 showing the cumulative percentage change in spot price due to flow demand shocks, and inventory demand shocks respectively. Estimated using the original Kilian and Murphy (2014) model. Estimated by replicating Kilian and Murphy (2014) using their original variables, sample period, and restrictions.

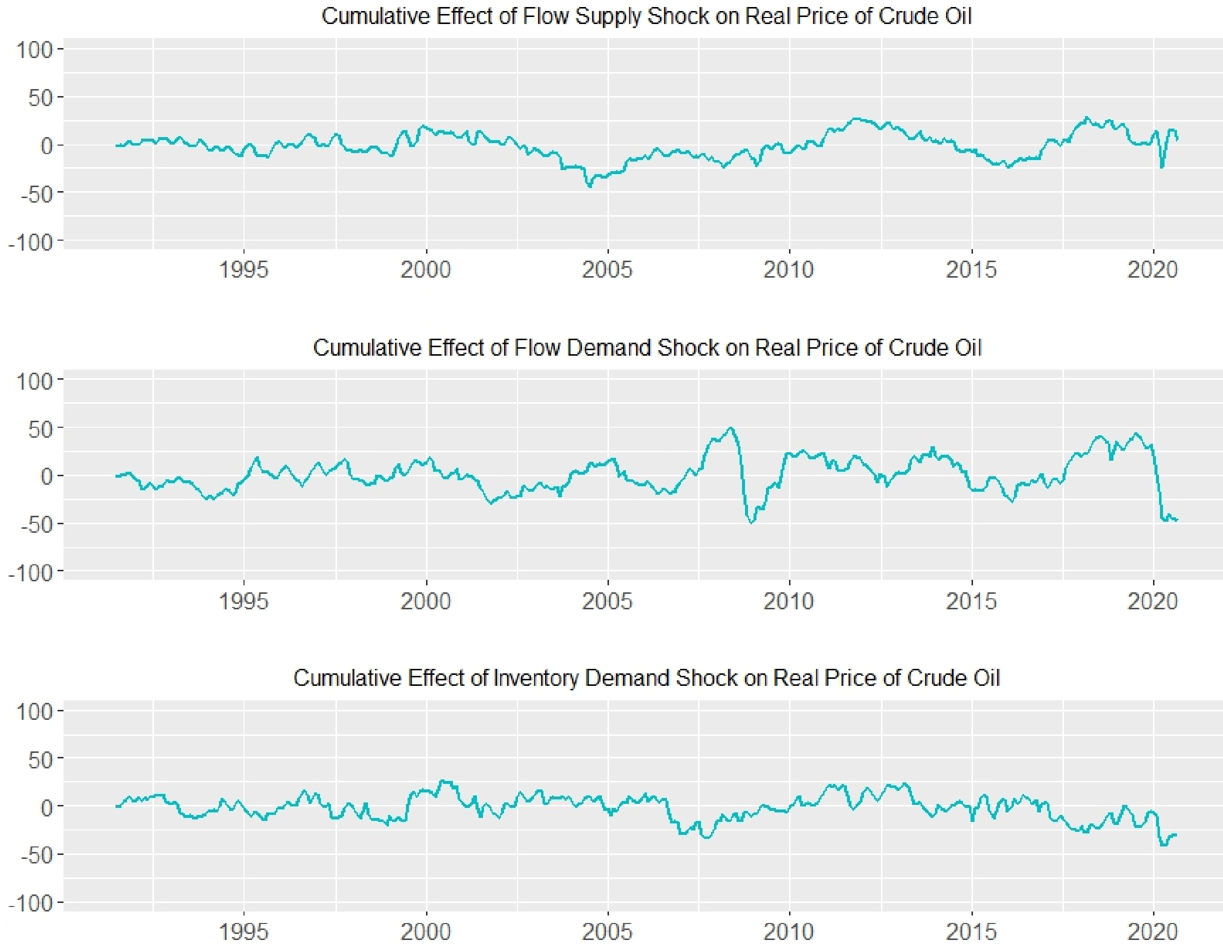

Figure A.10: Kilian and Murphy Method

Historical decomposition of the real spot price of brent oil from July 1991 to September 2020 showing the cumulative percentage change in spot price due to flow supply, flow demand, inventory demand, and futures demand shocks, respectively. Estimated using the original Kilian and Murphy (2014) model with my data and sample period.

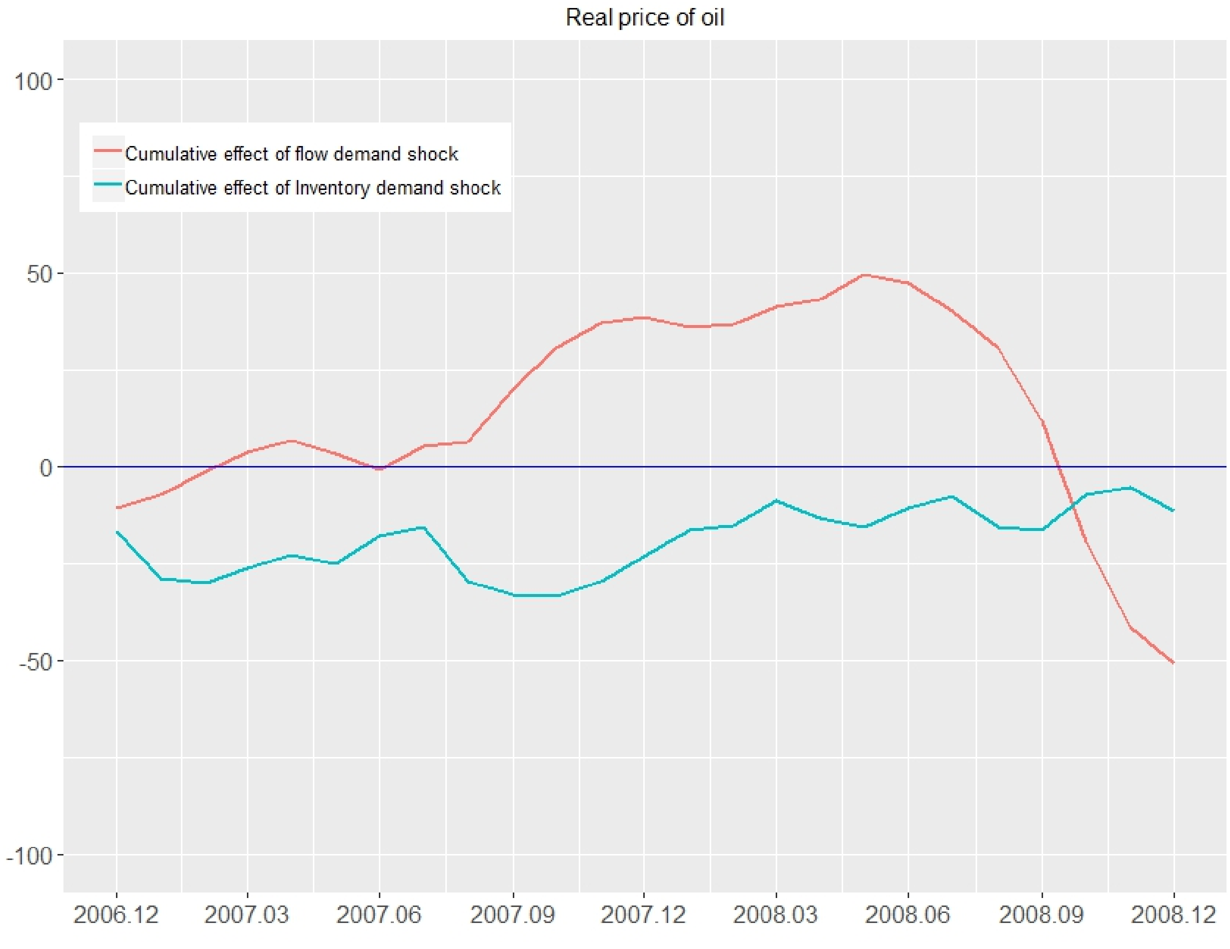

Figure A.11: Kilian and Murphy Method- 2008 Financial Crisis

Historical decomposition of the real spot price of brent oil from Jan. 2007 to Dec. 2008 showing the cumulative percentage change in spot price due to flow demand shocks, and inventory demand shocks respectively. Estimated using the original Kilian and Murphy (2014) model with my data and sample period.

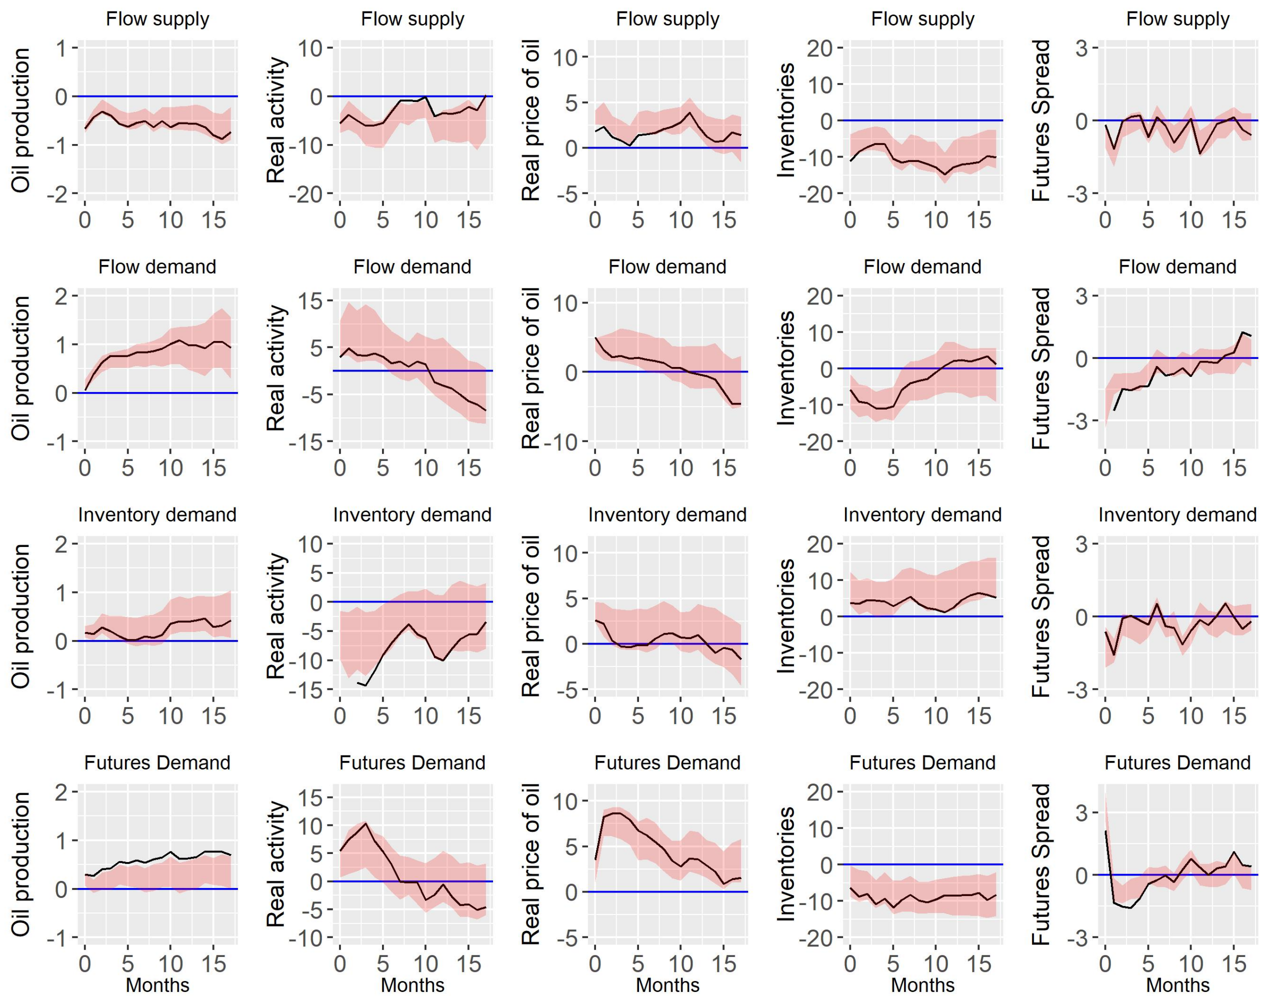

Figure A.12: Structural IRFs: 6mo Brent Futures

Structural Impulse Response Functions showing the response of each variable to a one standard deviation innovation to each structural shock. Responses are the cumulative % change for production, real activity, and the spot price, and cumulative level change for inventories. The Spread response is the difference in the futures and spot responses. The red band illustrates the 68% error band from the posterior distribution of the IRF's.Obtained as described in section 3, and in Appendix B. The 3mo Brent Oil Futures contract is replaced with the 6mo contract.

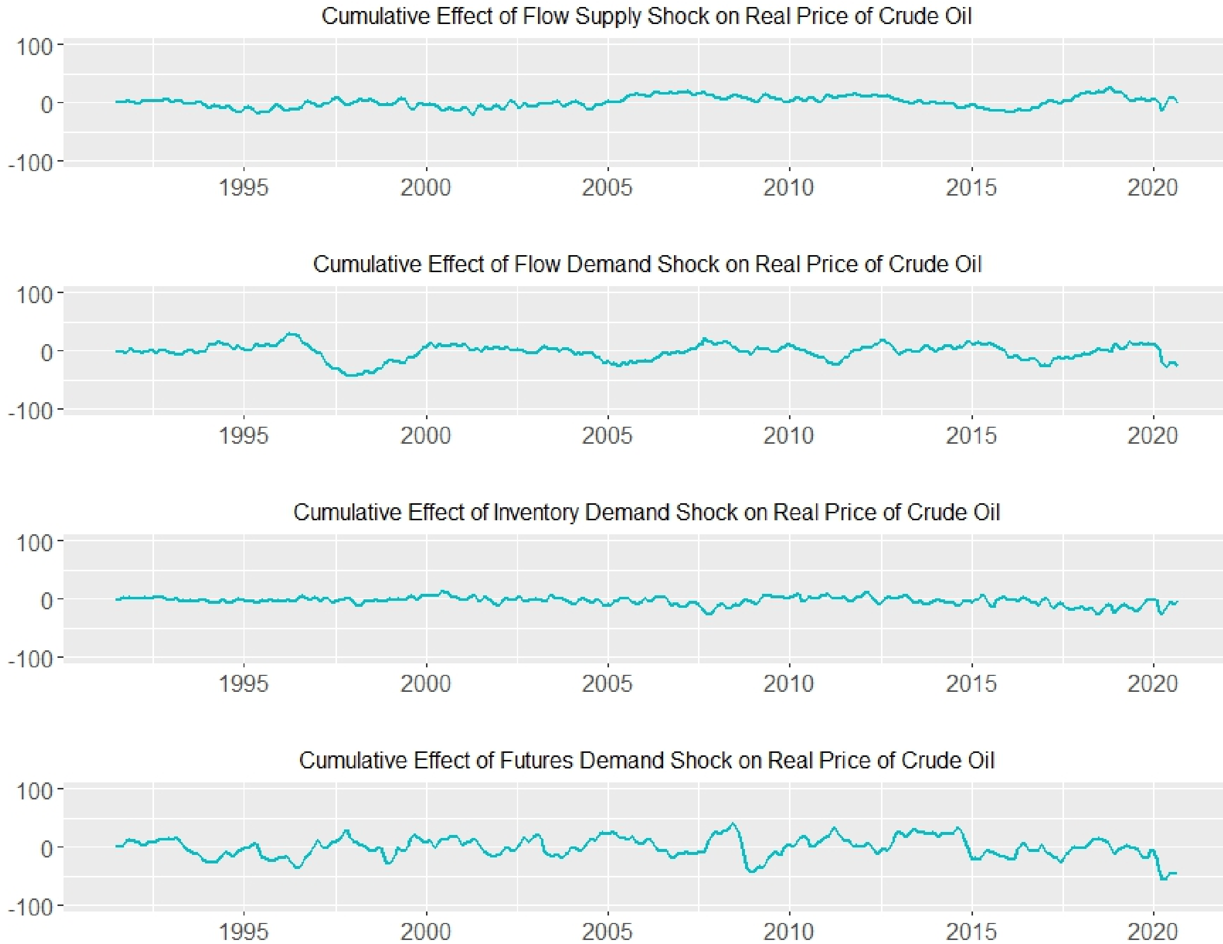

Figure A.13: Historical Decomposition: 6mo Brent Futures

Historical decomposition of the real spot price of brent oil from July 1991 to September 2020 showing the cumulative percentage change in spot price due to flow supply, flow demand, inventory demand, and futures demand shocks, respectively. The 3mo Brent Oil Futures contract is replaced with the 6mo contract.

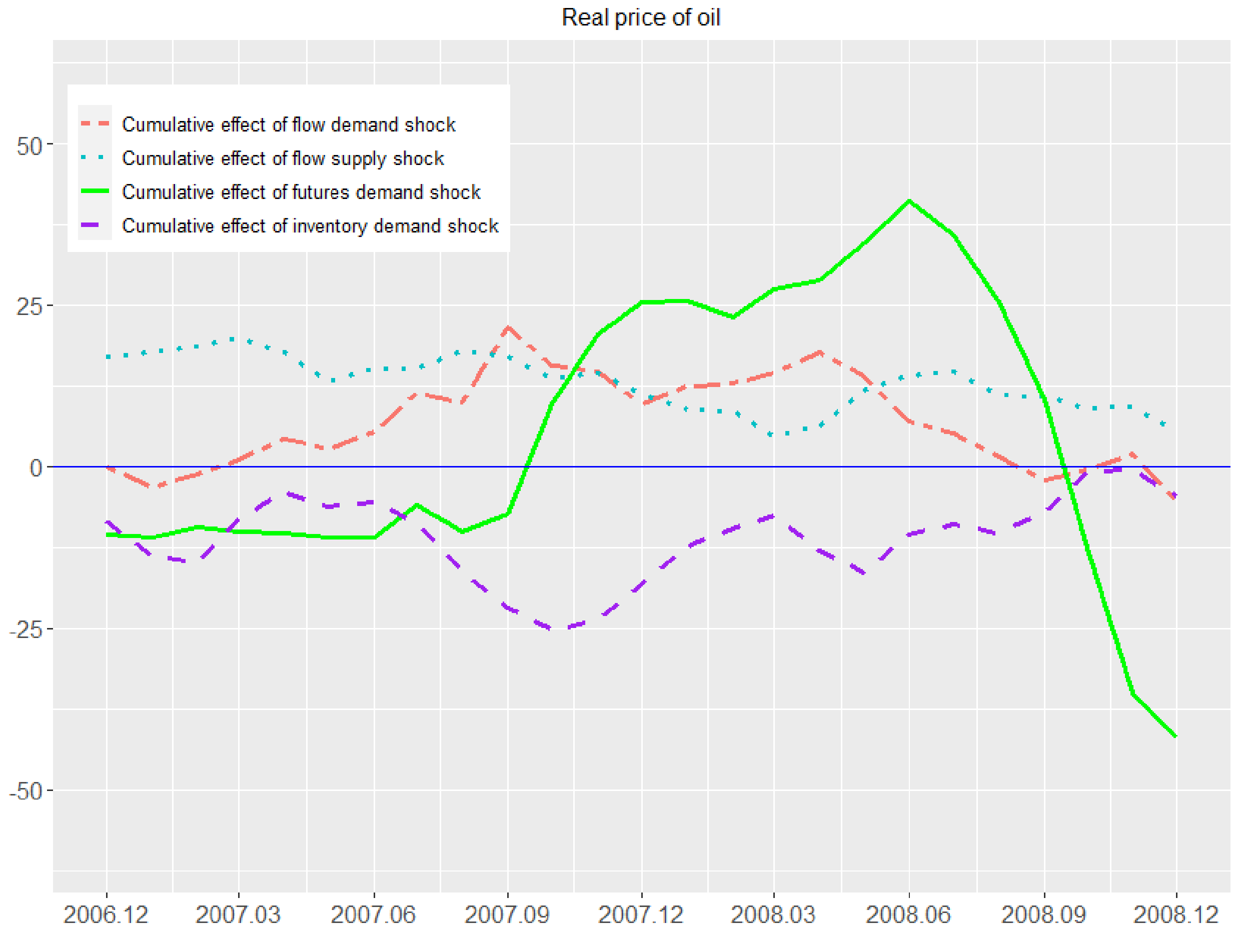

Figure A.14: Historical Decomposition-2008 Financial Crisis: 6mo Futures Contract

Historical decomposition of the real spot price of brent oil showing cumulative percentage change in spot price due to flow demand and futures demand shocks, respectively, between Jan. 2007 and Jan. 2009. The 3mo Brent Oil Futures contract is replaced with the 6mo contract.

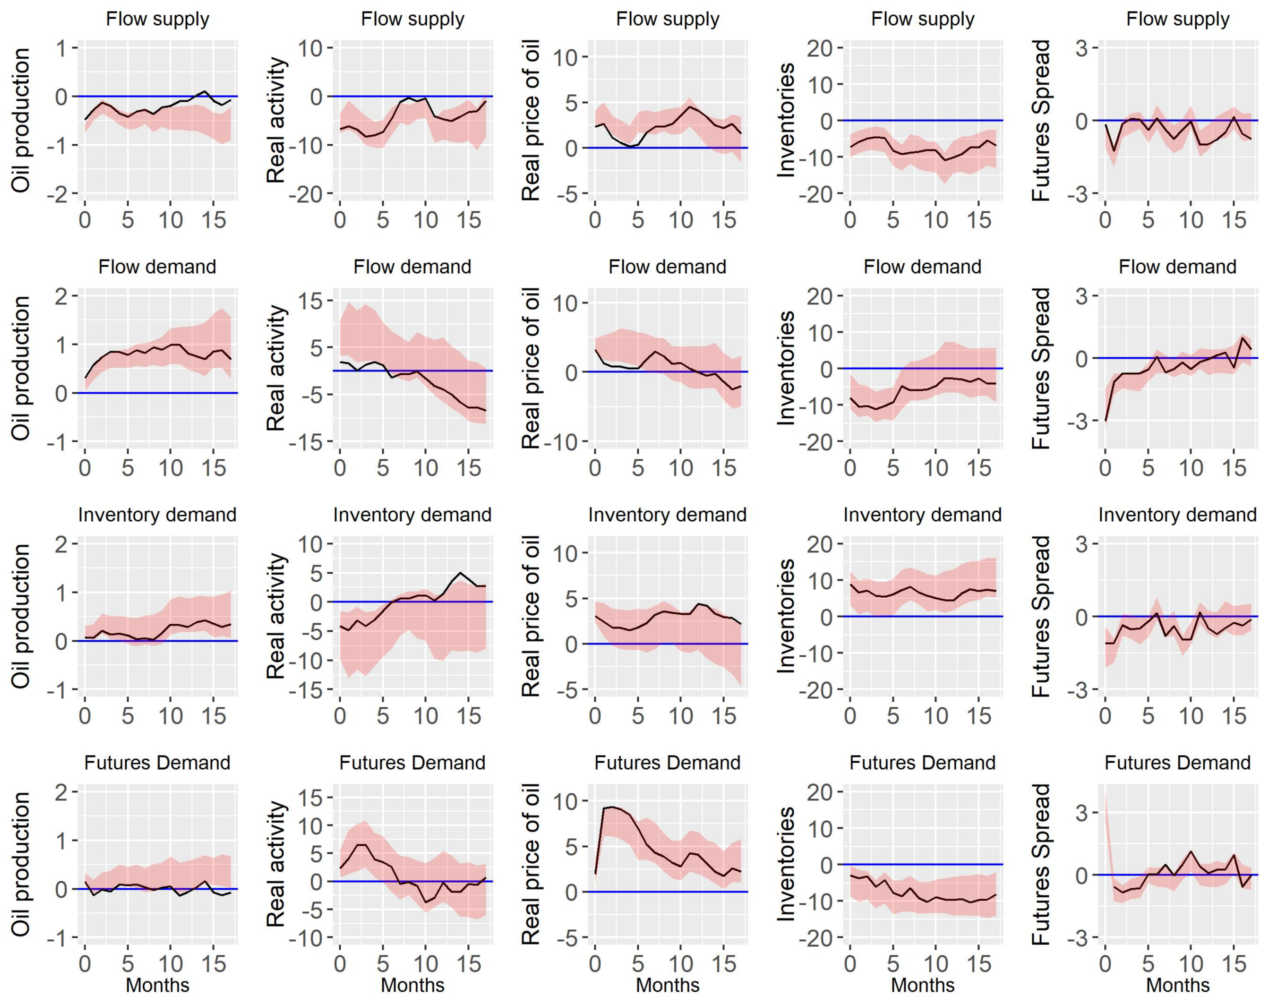

Figure A.15: Structural IRFs: Jorda

Structural Impulse Response Functions showing the response of each variable to a one standard deviation innovation to each structural shock. Responses are the cumulative % change for production, real activity, and the spot price, and cumulative level change for inventories. The Spread response is the difference in the futures and spot responses. The red band illustrates the 68% error band from the posterior distribution of the IRF's. Obtained as described in section 3, and in Appendix B, using Jorda (2005) linear projection methodology.

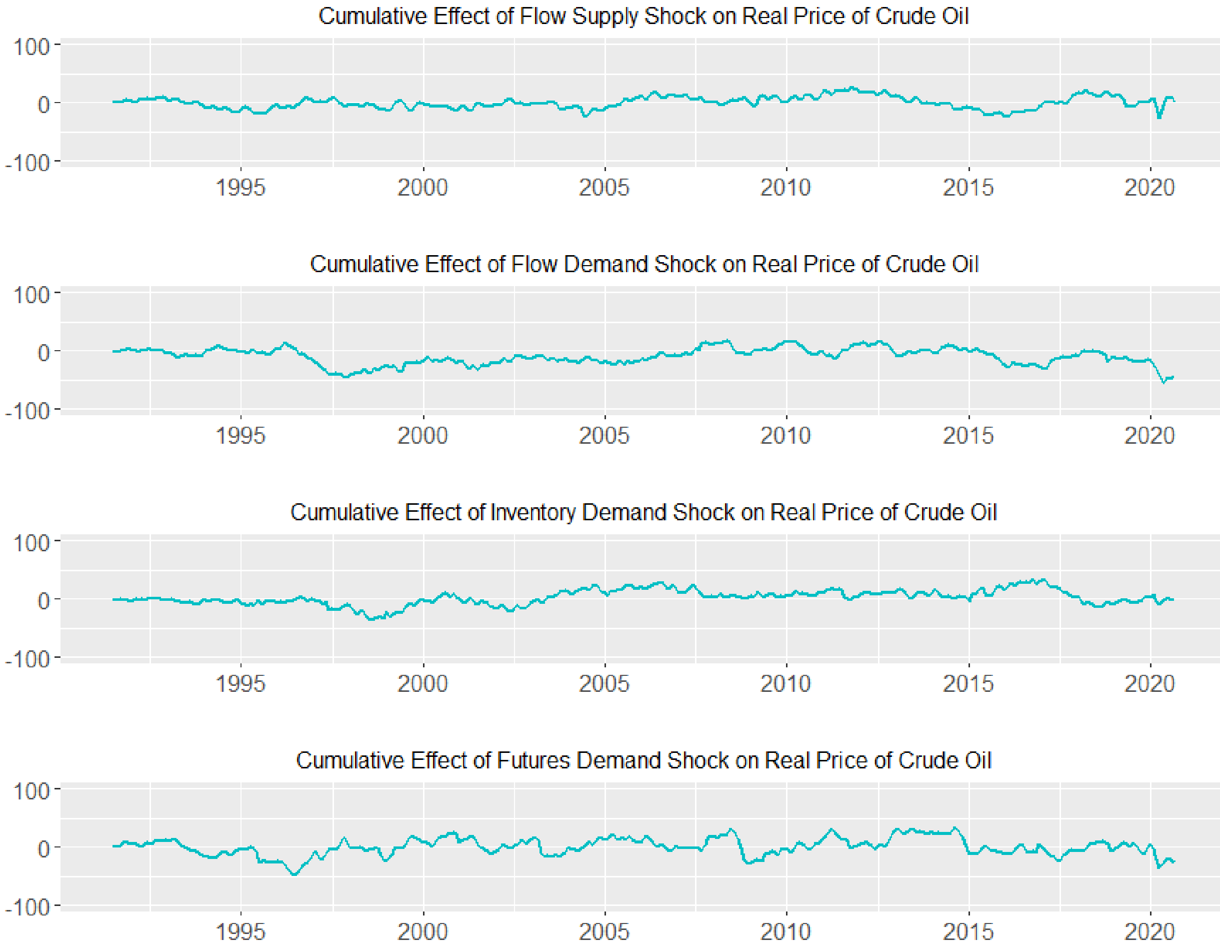

Figure A.16: Historical Decomposition: Jorda

Historical decomposition of the real spot price of brent oil from July 1991 to September 2020 showing the cumulative percentage change in spot price due to flow supply, flow demand, inventory demand, and futures demand shocks, respectively.

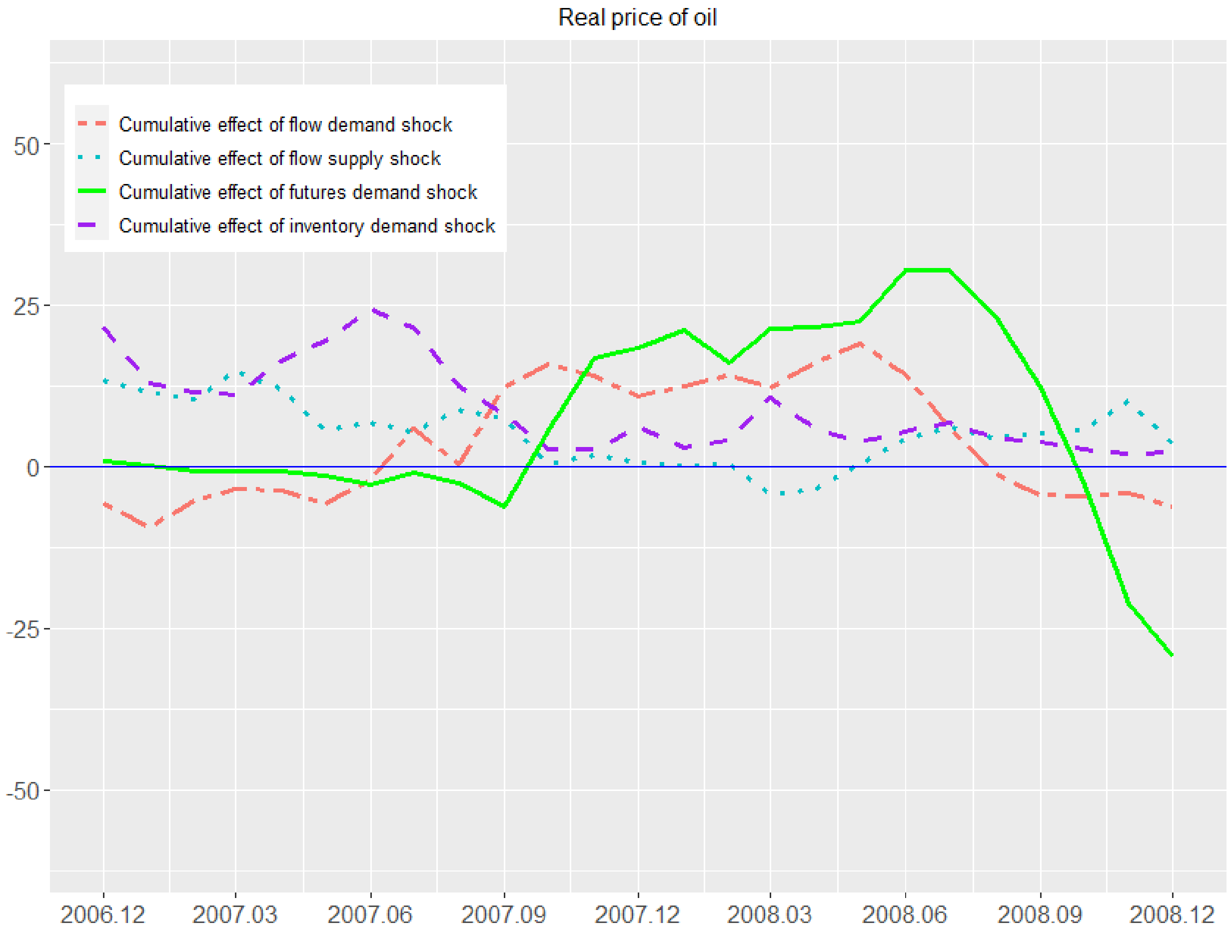

Figure A.17: Historical Decomposition-2008 Financial Crisis: Jorda

Historical decomposition of the real spot price of brent oil showing cumulative percentage change in spot price due to flow demand and futures demand shocks, respectively, between Jan. 2007 and Jan. 2009.

Table A.1: Subperiod Forecast Error Variance Decomposition of the Real Spot Price of Oil

Variance decomposition of the real spot price of oil reflecting the percentage of variation at each monthly horizon attributable to each respective shock. Estimated using main model with restrictions from Table 1 with 12 lags estimated over July 1990-Sept. 2020

			Full Sample		
Horizon	Supply	Demand	Inven.	Spec.	Resid.
1	24.78	24.11	25.20	23.72	2.19
2	20.77	22.52	20.45	33.10	3.16
3	13.69	26.25	14.42	43.45	2.20
4	11.84	27.03	11.75	47.88	1.50
5	12.17	27.28	10.50	48.88	1.17
6	11.57	28.64	9.25	49.39	1.15
7	11.31	30.29	8.43	48.88	1.08
8	11.08	30.64	8.20	49.06	1.02
9	10.69	30.69	7.86	49.81	0.96
10	10.29	31.24	7.54	50.00	0.93
11	9.90	31.83	7.23	50.13	0.92
12	9.32	33.74	7.11	48.95	0.88
13	9.98	35.93	6.88	46.42	0.78
14	11.33	37.88	6.72	43.33	0.74
15	12.68	39.42	6.63	40.56	0.72
% of Explained	12.8%	39.7%	6.7%	40.8%	
600	27.72	37.63	10.26	19.08	5.32
% of Explained	29.3%	39.7%	10.8%	20.1%	

Table A.2: Forecast Error Variance Decomposition: No narrative restrictions

Variance decomposition of the real spot price of oil reflecting the percentage of variation at each monthly horizon attributable to each respective shock. Estimated using main model with restrictions from Table 1 with 24 lags estimated over July 1991-Sept 2020.

Horizon	Supply	Demand	Inv.	Spec.	Resid.
1	24.07	36.62	10.30	28.82	0.19
2	20.31	30.51	8.88	40.14	0.16
3	15.44	21.85	7.53	53.96	1.22
4	13.24	16.66	6.96	60.18	2.97
5	12.77	13.70	7.02	62.18	4.33
6	11.14	11.82	7.27	61.69	8.07
7	9.91	10.44	7.15	61.35	11.15
8	9.38	9.81	7.27	62.25	11.29
9	9.21	9.39	7.55	62.58	11.28
10	9.41	9.17	8.05	61.68	11.69
11	9.93	8.81	8.36	60.15	12.75
12	10.34	8.70	8.43	58.15	14.38
13	10.06	9.01	9.79	55.57	15.58
14	9.46	8.93	11.46	52.75	17.40
15	8.77	8.91	12.55	49.31	20.46
% of Explained	11.03%	11.20%	15.78%	61.99%	
600	9.21	20.66	4.1	6.78	59.24
% of Explained	16.20%	11.21%	26.29%	10.09%	

Table A.3: Forecast Error Variance Decomposition: 6month Brent Oil Futures Contract

Variance decomposition of the real spot price of oil reflecting the percentage of variation at each monthly horizon attributable to each respective shock. Estimated using main model with restrictions from Table 1 with 24 lags estimated over July 1991-Sept 2020, and replacing the 3mo Brent Oil Futures contract with the 6mo contract.

Horizon	Supply	Demand	Inv.	Spec.	Resid.
1	4.49	35.58	9.46	30.70	19.76
2	3.91	31.89	8.27	38.49	17.45
3	4.63	27.21	8.51	45.71	13.95
4	5.92	22.90	9.89	48.94	12.35
5	7.80	19.88	11.02	49.68	11.61
6	7.35	17.86	12.34	48.31	14.14
7	6.78	16.81	13.59	46.95	15.86
8	6.42	16.18	13.16	47.67	16.59
9	6.17	16.15	12.68	47.72	17.29
10	6.02	16.24	12.24	46.70	18.79
11	6.07	15.93	11.77	45.19	21.04
12	6.16	16.23	11.35	43.04	23.23
13	5.78	17.19	10.68	40.52	25.83
14	5.63	17.36	10.22	38.03	28.76
15	5.51	17.52	10.25	35.17	31.54
% of Explained	8.05%	25.60%	14.97%	51.38%	
600	9.13	16.96	12.42	6.31	55.18
% of Explained	20.38%	37.84%	27.70%	14.08%	

Table A.4: Forecast Error Variance Decomposition: Jorda

Variance decomposition of the real spot price of oil reflecting the percentage of variation at each monthly horizon attributable to each respective shock. Estimated using main model with restrictions from Table 1 with 24 lags estimated over July 1991-Sept 2020, and IRFs calculated using Jorda (2005) linear projection methodology.

Horizon	Supply	Demand	Inv.	Spec.	Resid.
1	10.16	19.71	17.16	41.49	11.48
2	7.59	14.87	12.71	56.00	8.83
3	5.91	10.83	8.16	67.99	7.10
4	5.96	8.73	6.10	71.68	7.52
5	6.85	7.57	5.01	72.39	8.19
6	6.40	7.37	4.32	71.20	10.71
7	5.90	7.59	3.84	70.29	12.38
8	5.47	7.38	3.58	71.06	12.50
9	5.18	7.57	3.44	71.52	12.29
10	5.04	8.21	3.44	70.97	12.34
11	5.03	8.79	3.42	69.87	12.89
12	4.98	10.14	3.25	68.30	13.34
13	4.66	12.16	3.38	66.37	13.43
14	4.49	13.66	3.58	64.16	14.10
15	4.44	15.49	3.55	61.31	15.21
% of Explained	5.24%	18.27%	4.19%	72.31%	
600	10.15	34.15	5.93	12.48	37.29
% of Explained	16.19%	54.45%	9.46%	19.90%	

Table A.5: Granger Causality Test of Futures Returns on DCOT Swap Dealer Oil Futures Position Changes

This table presents the results of the bivariate causality test obtained by estimating the regression:

$$R_t = \alpha + \sum_{i=1}^{m} \gamma_i R_{t-i} + \sum_{j=1}^{n} \beta_j \Delta Net_{t-j} + \epsilon_t$$

and then running a partial F-test on the restriction $\sum_{j=1}^{n} \beta_j = 0$. Durbin-Watson (DW) statistic of 2 indicates no remaining serial correlation in returns in the unrestricted model. ΔNet_{t-j} represents the net weekly change in swap dealer positions in oil futures on week $(t-j)$ from the CFTC's DCOT report, while R_{t-i} is the lagged oil futures return on week $t-i$. Sample period Jan. 31, 2012 to Dec. 31, 2019.

Market	m,n	Estimate $\sum \beta_j$	partial F	p-value	DW
Brent	4,2	1.428	0.62	0.63	2.01
WTI	4,2	3.256	0.965	0.43	2.01

Table A.6: Average Coefficients for Deletions Event Study

This table reports the average coefficients and the standard error of the coefficient estimates across univariate regressions and bivariate regressions for stock i. POST is a dummy variable equal to 1 if a return is in a 12 month window after stock i is added to the S&P 500 index, and 0 if the return is in the 12 months window before stock i is added to the index. SP500 and NonSP500 are the returns on the S&P 500 and those non-SP500 firms in the SP1500 excluding the influence of stock i during the period it is within the respective index at time t. Standard errors are in parentheses.

Sample Period	Univariate	Bivariate
Constant	-0.0011	-0.0017
	(0.0002)	(0.0002)
POST	0.0018	0.0014
	(0.0004)	(0.0004)
$SP500$	1.1951	0.4359
	(0.0445)	(0.0483)
$SP500^2$	-2.8879	-7.9543
	(0.9559)	(1.7175)
$POST \times SP500$	-0.0569	-0.5094
	(0.0363)	(0.0684)
$POST \times SP500^2$	-0.7735	1.6324
	(1.7191)	(4.4823)
$NonSP500$	1.1951	0.7097
	(0.0445)	(0.0532)
$NonSP500^2$	-2.8879	4.3320
	(0.9559)	(1.5091)
$POST \times NonSP500$	-0.0569	0.4342
	(0.0363)	(0.0726)
$POST \times NonSP500^2$	-0.7735	-1.4283
	(1.7191)	(3.3210)
Obs	215	215
R-Square	0.20	0.23
	(0.0155)	(0.0109)

Table A.7: Pre vs Post S&P 500 deletion coskewness estimates

This table reports the average coskewness coefficients and across the univariate regressions of each individual stock against the returns of S&P 500 and Non-S&P 500 stocks in Panel A:

$$R_{i,t} = \alpha_i + \beta_{i1} index_{jt} + \gamma_{i1} index_{jt}^2 + \epsilon_{it}$$

and the bivariate regressions in Panel B:

$$R_{i,t} = \alpha_i + \beta_{i1} SP500_t + \beta_{i2} NonSP500_t + \gamma_{i1} SP500_t^2 + +\gamma_{i2} NonSP500_t^2 + \epsilon_{it}$$

The regressions are run seperately on the pre-period including the 12 full calendar months before the month of deletion, and the post-period including the full 12 calendar months after the month of deletion. Significance of the difference estimates at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively. Standard errors are in parenthesis.

Univariate Regression								
		S&P500			*Non-S&P500*			*Diff*
Sample Period	*Obs*	$\gamma_{1,pre}$	$\gamma_{1,post}$	$\Delta\gamma_{1,sp}$	$\gamma_{1,pre}$	$\gamma_{1,post}$	$\Delta\gamma_{1,nsp}$	$\Delta\gamma_{1,sp} - \Delta\gamma_{1,nsp}$
1995-2017	215	-2.8879	-3.5396	-0.6517	-0.8361	-1.0137	-0.1776	-0.4741
		(0.9552)	(1.4437)	(1.7354)	(0.8366)	(1.1302)	(1.4625)	(1.5927)

Bivariate Regression								
		S&P500			*Non-S&P500*			*Diff*
Sample Period	*Obs*	$\gamma_{1,pre}$	$\gamma_{1,post}$	$\Delta\gamma_1$	$\gamma_{2,pre}$	$\gamma_{2,post}$	$\Delta\gamma_2$	$\Delta\gamma_1 - \Delta\gamma_2$
1995-2017	215	-7.9543	-5.5407	2.4136	4.3320	2.6948	-1.6372	3.7856
		(1.7175)	(4.1603)	(4.5893)	(1.5091)	(2.9102)	(3.4006)	(7.7781)

Appendix B

Econometric Appendix

This appendix provides additional details on the estimation of my structural vector autoregressive model, including a derivation of the relationship between the reduced form estimates and the structural impulse response functions, the historical decomposition of oil prices into its component structural shocks, and practical details on implementation and model selection.

Computing Structural Impulse Response Functions

First, Rearranging equation 2.1 as follows:

$$y_t = B_0^{-1}\beta_0 + B_0^{-1}\sum_{i=1}^{24}\beta_i y_{t-i} + B_0^{-1}u_t \tag{B.1}$$

allows us to equate equation 2.1 and 2.2, and see that the innovations in the reduced form model ϵ_t actually represent a linear combination of all the contemporaneous structural shocks in the system, such that $\epsilon_t = B_0^{-1}u_t$.

To compute structural impulse response functions, we can then rewrite equation B.1 using lag operator notation as[1]:

$$A(L)y_t = \alpha + \epsilon_t, \quad A(L) = I_n - \sum_{i=1}^{24} A_i L^i \tag{B.2}$$

Bringing $A(L)$ to the RHS and substituting $\epsilon_t = B_0^{-1}u_t$ allows us to rewrite this as:

$$y_t = A(1)^{-1}B_o^{-1}u_t \tag{B.3}$$

[1] The following assumes that y_t is covariance stationary.

Defining $\mu = A_0 = A(1)^{-1}\alpha$ and $\psi(L) = A(L)^{-1}$ then gives us:

$$y_t = \mu + \psi(L)B_0^{-1}u_t \tag{B.4}$$

Finally, to recover the structural moving average representation we can expand out the lag operator as:

$$y_t = \mu + \sum_{k=0}^{\infty} \psi_k B_0^{-1} u_{t-k} \tag{B.5}$$

And then define $\theta_k = \psi_k B_0^{-1}$:

$$y_t = \mu + \sum_{k=0}^{\infty} \theta_k u_{t-k} \tag{B.6}$$

This form expresses the level of y_t as the cumulative evolution of all responses to past shocks. Here, the square matrices θ_k contain information on the effect of each structural shock on each reduced form variable through time. To illustrate more clearly, the matrix form for a simpler, 4 variable system, would be:

$$\begin{bmatrix} y_{Qt} \\ y_{Rt} \\ y_{Pt} \\ y_{It} \end{bmatrix} = \begin{bmatrix} \mu_Q \\ \mu_R \\ \mu_P \\ \mu_I \end{bmatrix} + \begin{bmatrix} \theta^0_{11} & \theta^0_{12} & \theta^0_{13} & \theta^0_{14} \\ \theta^0_{21} & \theta^0_{22} & \theta^0_{23} & \theta^0_{24} \\ \theta^0_{31} & \theta^0_{32} & \theta^0_{33} & \theta^0_{34} \\ \theta^0_{41} & \theta^0_{42} & \theta^0_{43} & \theta^0_{44} \end{bmatrix} \begin{bmatrix} u_{1t} \\ u_{2t} \\ u_{3t} \\ u_{4t} \end{bmatrix} + \begin{bmatrix} \theta^1_{11} & \theta^1_{12} & \theta^1_{13} & \theta^1_{14} \\ \theta^1_{21} & \theta^1_{22} & \theta^1_{23} & \theta^1_{24} \\ \theta^1_{31} & \theta^1_{32} & \theta^1_{33} & \theta^1_{34} \\ \theta^1_{41} & \theta^1_{42} & \theta^1_{43} & \theta^1_{44} \end{bmatrix} \begin{bmatrix} u_{1,t-1} \\ u_{2,t-1} \\ u_{3,t-1} \\ u_{4,t-1} \end{bmatrix} + ... \tag{B.7}$$

The elements θ^k_{sj} are scalars representing the effect of a given shock at time $t-k$ to u_j on y_s, and can be recovered from the reduced form coefficients A_i and errors ϵ_t and the identified impact response matrix B_0^{-1}. The scalar time series $(\theta^0_{sj}, \theta^1_{sj}, \ldots, \theta^h_{sj})$

defines the impulse response functions.[2]

Forecast Error Variance Decomposition

The contribution of shock j to the forecast error variance, or mean squared prediction error, $MSPE^s(k)$ at horizon k for variable y_s is decomposed as follows:

$$1 = \sum_{j=1}^{J} \frac{MSPE_j^s(k)}{MSPE^s(k)} \tag{B.8}$$

where

$$MSPE_j^s(k) = \sum_{i=0}^{k-1} (\theta_{sj}^k)^2 \tag{B.9}$$

As an example, the ratio above will provide a percentage of forecast variance of oil prices explained by supply, demand, inventory demand, futures demand shocks, and residual shocks. The total contribution of all of these shocks sum to one.

Historical Decomposition

I also decompose the cumulative contribution of a shock u_j to a variable y_s as the sum of the impacts of all past shock through time:

$$y_{st}^j = \sum_{i=1}^{t-1} \theta_{sj}^i u_{j,t-i} \tag{B.10}$$

This allows me to highlight the important drivers of any variable throughout different points in time. For example, Kilian and Murphy (2014) use this to highlight flow demand shocks as being much more important in determining oil prices than

[2]In addition to the standard SVAR methodology described above, I also confirm the robustness of all key results using IRF's obtained following the linear projection methodology of Jorda (2005). Results are in Appendix A, Tables A.15-A.17 and Figure A.4.

flow supply and inventory demand shocks throughout recent history.

Implementation

To estimate my structural model, I use the restrictions outlined in section 3 to generate a set of allowable models, or allowable iterations of B_0^{-1}, from the set of all models or B_0^{-1} that satisfy the reduced form estimates. In practice, the set of all models is very large, so I limit myself to sampling a smaller set of the models which satisfy the reduced form, and evaluate that set against the restrictions. As long as the models are sampled randomly and a sufficient number are sampled, the collection of allowable models is free from bias.

To randomly generate models which satisfy the reduced form estimates, I follow the algorithm introduced by Rubio-Ramirez, Waggoner, and Zha (2010) and subsequently used by Kilian and Murphy (2014). The algorithm works as follows:

1. Decompose the reduced form variance-covariance matrix $\Sigma_\epsilon = P\Lambda P'$ and let $B_0^{-1} = P\Lambda^{0.5}$ so that $B_0^{-1}B_0^{-1\prime} = \Sigma_\epsilon$
2. Randomly generate an NxN matrix K, where each element is generated from a standard normal distribution
3. Apply QR decomposition to matrix K, resulting in a orthogonal rotation matrix Q satisfying $QQ' = I$
4. Generate $\tilde{B}_0^{-1} = P\Lambda^{0.5}Q$.
5. The resulting $\tilde{B}_0^{-1}\tilde{B}_0^{-1\prime} = P\Lambda^{0.5}QQ'\Lambda^{0.5\prime}P' = P\Lambda P' = \Sigma_\epsilon$ thus by randomly generating Q using step 3 we can repeatedly generate new $\tilde{B}_0^{-1}$ which satisfy the initial reduced form model
6. Repeat steps 2-5 for each random draw of K

I use this algorithm with 800 million draws of the rotation matrix to obtain a large number of candidate models, and then check each model against sign, elasticity, and narrative restrictions, keeping only those models which satisfy all credibility conditions in the allowable set.

Inference

Since the model uses sign and elasticity restrictions for set identification instead of exact identification, there are multiple models which are equally valid. Furthermore, the structural parameters are recovered from the reduced form parameters, which are point estimates, and are themselves uncertain. To make the analysis more robust, I take a Bayesian approach to model selection, and create confidence bands for impulse response functions.

I use the standard approach outlined in Inoue and Kilian (2013), and randomly draw reduced form parameters A_i and Σ_ϵ from their posterior distribution assuming a Gaussian-inverse Wishart prior, and run them through the Rubio-Ramirez et al (2010) algorithm to generate impulse response functions. I utilize 100 draws from the posterior distribution, and 800,000 draws of the rotation matrix. I then follow a similar approach as Kilian and Murphy (2014), and select the model which most closely matches posterior median price elasticity of oil demand in use in response to supply shocks and posterior median price elasticity of demand in use due to futures demand shocks. Importantly, key results are highly robust to selection of other admissible models. I also use the posterior distribution of impulse responses to generate error bands for the point estimates of the impulse responses.