Title
Essays in Resource Economics

Author
Stevens, Reid Blake

Publication Date
2015

Peer reviewed|Thesis/dissertation

Essays in Resource Economics

By

Reid Blake Stevens

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Agricultural and Resource Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in Charge:

Professor Gordon C. Rausser
Professor Brian Wright
Professor Terrence J. Hendershott

Spring 2015

Essays in Resource Economics

By

Reid Blake Stevens

Abstract

Essays in Resource Economics

by

Reid Blake Stevens

Doctor of Philosophy in Agricultural and Resource Economics

University of California, Berkeley

Professor Gordon C. Rausser, Chair

This dissertation consists of three chapters. The first chapter investigates whether U.S. oil market policy succeeded in lowering the price of crude oil. The second chapter examines the effect of private sponsorship of university research on the allocation of expenditures between public good research and commercial applications. The third chapter estimates the effects of congestion pricing in the Bay Area on public transportation ridership and air pollution levels.

The first chapter, "The Strategic Petroleum Reserve and Crude Oil Prices", uses a structural vector autoregression (VAR) model of the U.S. oil market to estimate the effect of the Strategic Petroleum Reserve (SPR) on crude oil prices. In contrast to existing literature, I show that unanticipated oil releases from the SPR have no measurable effect on oil prices. However, unanticipated oil purchases for the SPR raise oil prices 1.5 percent. I confirm these estimates with two alternate VAR identification strategies that use external information, rather than the traditional recursiveness assumption. First, the SPR purchase schedule is used as an instrument in the VAR model to identify the effect of SPR purchase shocks on oil prices. I verify the statistically significant purchase effect and show that standard identifying assumptions have a downward bias. Second, I identify SPR release shocks with high-frequency futures market data. Using several futures contracts to estimate unanticipated SPR policy shocks, I find that SPR releases do not have an effect on oil prices. To explain this puzzle, I estimate an interacted VAR with time-varying policy effects. This model shows that SPR purchases only increase oil prices when uncertainty is high, but SPR releases do not lower oil prices at any level of uncertainty.

The second chapter, "Public vs. Private Good Research at Land Grant Universities", examines the effect of private sponsorship of university research on the allocation of expenditures between public good research and commercial applications. Throughout the land-grant university system, there is much concern that as a result of reduced government funding, fundamental research will be neglected at the expense of research that is geared toward commercial applications. This paper attempts to shed some light on the relationship between

research priorities and the availability of public funding for university research. In particular, we use both a static and a dynamic model to investigate the conditions under which university/private research partnerships can crowd-in or crowd-out basic science research as public funding becomes scarcer.

In the third chapter, "The Effects of Congestion Pricing on Public Transportation Demand and Air Pollution", I examine the effects of congestion pricing in the Bay Area. On July 1, 2010, congestion pricing was implemented on the San Francisco-Oakland Bay Bridge. Using panel data on hourly commuter rail (BART) and commuter bus (AC-Transit) ridership from June 1, 2009 through July 31, 2011, we examine the effect of congestion pricing on public transportation usage and air pollution. Using a difference-in-difference model to control for unobserved variables, we estimate that BART ridership rose 4-8 percent during peak hours, while AC-Transit ridership rose 11-15 percent. Our estimated cross-price elasticities for public transportation, 0.2 for commuter rail and 0.15 for commuter bus, are slightly below estimates from the transportation economics literature. We also use panel data on air pollution levels in the Bay Area during the same time period to examine the effect of congestion pricing on air pollution. Using a similar difference in difference model, we find air pollution levels are flat during peak hours, but rise during off peak hours. As a robustness check, we use regression discontinuity models, which suggest similar, but slightly smaller effects of congestion pricing on air pollution and public transportation.

For Charlotte.

1 The Strategic Petroleum Reserve and Crude Oil Prices

1.1 Introduction

The Strategic Petroleum Reserve (SPR) was created in response to the 1975 Arab Oil Embargo, and was intended to, "Store petroleum to reduce the adverse economic impact of a major petroleum supply disruption to the United States" (U.S. Congress, 1975). Since the SPR was constructed in 1977, over 150 million barrels of crude oil have been released from the reserve. During these releases, crude oil from the SPR made up 5 to 12 percent of weekly domestic oil supply. The government has purchased over 850 million barrels of crude oil for the reserve, which accounted for between 2 to 4 percent of U.S. crude oil consumption during purchase weeks.

Despite this long history of purchases and releases, little is known about the effect of the SPR on crude oil prices. Much of the existing literature on the SPR focuses on optimal purchase and release strategies, and does not directly estimate the effect of the SPR on crude prices (see Tolley and Wilman, 1977; Nichols and Zeckhauser, 1977; Teisberg, 1981; Chao and Manne, 1983; Zhang et al., 2009). These papers generally assume a SPR crude oil release during a supply disruption will lower oil prices, but SPR purchases when the oil market is calm will not affect prices. Unsurprisingly, these models show that under optimal management, the benefits of the SPR far outweigh the costs of construction, oil acquisition, and storage. The estimated benefits of the SPR are even higher when the negative macroeconomic effects of oil price spikes are taken into account (Weimer, 1982; Bohi and Toman, 1996; Leiby and Bowman, 2000). When Congress recently approved plans to expand the reserve from its current level of 727 million barrels to 1 billion barrels, SPR proponents praised its ability to lower oil prices and stabilize the economy during oil supply disruptions (Government Accountability Office, 2006; Congressional Research Service, 2006).

The few empirical estimates of the SPR price effect vary widely: SPR releases are estimated to lower the price of crude oil by 3 to 32 percent, while SPR purchases are estimated to increase the price of crude oil by 0.4 to 32 percent (U.S. Senate, 2003; Verleger, 2003; Considine, 2006). The lack of consistency among these estimates is evidence that identifying the effect of SPR policy on oil prices is difficult. Some SPR crude oil releases occurred in response to severe oil market supply disruptions (e.g., 1991 Desert Storm sale) while others occurred because oil prices were relatively low and stable (e.g., 1996 sales to reduce the deficit). Similarly, the largest SPR purchases were a response to high oil prices caused by the Arab Oil Embargo. Purchases for the SPR were also made in the late 1980s, while there was excess capacity in oil markets and prices were relatively low.

Isolating the effect of SPR policy is challenging because the policy depends, in part, on the state of the oil market. Changes in the price of oil following an SPR policy action reflects both the policy and market conditions to which the policy is responding. The endogeniety of crude oil supply, crude oil demand, and SPR policy has likely led to the widely varying estimates of the SPR's effect on the price of oil.

Given these difficulties, I use several structural vector autoregression (VAR) models of the U.S. oil market to estimate the effect of SPR policy. Since Kilian (2009), structural VAR models have been used to model both the global and U.S. oil markets (Apergis and Miller, 2009; Lippi and Nobili, 2012; Baumeister and Peersman, 2013b; Jo, 2014). This paper contributes to this literature by adding oil market policy variables to the VAR model, following the voluminous monetary policy literature (Bernanke and Blinder, 1992; Sims, 1992; Bernanke and Mihov, 1998; Christiano et al., 1999). As a starting point, I identify SPR policy effects by restricting the structural shocks through the recursiveness assumption. For example, the exclusion restrictions in this model imply that oil supply and oil demand respond only with a lag to SPR policy shocks. To justify this assumption, I estimate the model with weekly data, rather than the monthly or quarterly data typically used in the oil market literature. In this benchmark model, an unanticipated SPR purchase raises the price of the price of oil 1.5 percent over 20 weeks following purchase, but an unanticipated SPR release does not have a statistically significant effect on the price of oil.

Though VAR models in both the oil market and monetary policy literature are identified by assuming recursiveness (see Kilian, 2009; Christiano et al., 1996), this identifying assumption is controversial (Rudebusch, 1998). Papers on oil markets have increasingly relied on sign restrictions to identify structural shocks (Baumeister and Peersman, 2013b; Kilian and Murphy, 2012). I take a different approach and identify policy shocks using information that is external to the VAR model, but correlated with unanticipated policy actions (Stock and Watson, 2008; Olea et al., 2012). Following work by Romer and Romer (1989), who use the minutes from Federal Reserve meetings to identify monetary policy actions that were exogenous with respect to market conditions, many papers have constructed external instruments to identify structural shocks (Stock and Watson, 2012). I construct two separate instruments to partially identify the effects of SPR purchases and releases in a VAR framework.

First, I use the SPR purchase schedule as an instrument to identify unanticipated SPR purchases. From 1999 to 2010, the Department of Energy (DOE) hired private contractors to purchase oil for the SPR. These contractors were required to deliver a specified amount of oil over a fixed time period (e.g., a contractor would agree deliver 9.4 million barrels of oil from April to July 2005). The purchase schedules were not immediately announced and exemptions from the schedule were rarely granted, which makes the purchase schedule correlated with unanticipated SPR purchases. The schedule was also set well ahead of the delivery window, so the purchase schedule is uncorrelated with other structural shocks at the time of purchase. Using the purchase schedule as an instrument to partially identify the VAR model, I estimate that SPR purchase shocks cause crude oil prices to rise 6 percent over 15 weeks following purchase.

Next, I use crude oil futures data to identify the effect of SPR releases on crude oil prices. This method follows work by Bagliano and Favero (1999) and Cochrane and Piazzesi (2002), who estimate monetary policy shocks with Federal Funds futures data. In these models, the change in futures prices immediately following a policy announcement is used to estimate the unanticipated component of the policy. For example, an unanticipated SPR release announcement would be immediately followed by a drop in the crude oil futures price, whereas a wholly anticipated SPR release announcement would elicit no response from the crude oil futures market. Using the benchmark U.S. crude oil futures contract, West Texas Intermediate (WTI), to estimate policy shocks confirms that unanticipated crude oil releases from the SPR do not lower the price of crude oil.

To ensure the change in oil futures prices following an SPR release reflects only the policy effect, and not other shocks to the oil market, I construct a novel financial instrument: the spread between the U.S. (WTI) and Canadian (WCS) crude oil futures. Given the similarities and proximity of the U.S. and Canadian crude oil markets, a drop in the U.S. crude futures price relative to Canadian crude futures immediately following a SPR release announcement should capture only the effect of SPR policy, and not other shocks to the oil market. This instrument also shows that unanticipated SPR releases do not lower the price of crude oil.

Under a variety of empirically plausible identifying assumptions, SPR purchase shocks raise oil prices, but SPR release shocks have no measurable impact on prices. The asymmetric effect of SPR policy is puzzling. Following the growing uncertainty literature (Aastveit et al., 2013; Bekaert et al., 2013), I augment the VAR model with an uncertainty variable and policy-uncertainty interaction terms, which allows SPR policy to have time-varying price effects. I find that an unanticipated SPR purchase has no measurable effect on prices when oil market uncertainty is low, while an unanticipated purchase increases the price of oil 3 percent when oil market uncertainty is high. The price of crude oil is unaffected by SPR release shocks at all levels of oil market uncertainty. These results are robust to using alternate measures of uncertainty based on stock market volatility and policy uncertainty. The asymmetric effect of SPR purchases under uncertainty is consistent with the literature on the value of real options under uncertainty (Dixit and Pindyk, 1994; Bloom, 2009).

The central result of my paper–SPR release shocks do not lower oil prices but SPR purchase shocks raise oil prices–is at odds with the academic literature and has a number of implications for SPR management. First, SPR purchases should be avoided during periods of high oil market uncertainty. Though policymakers tend to be interested in filling the reserve when oil markets are volatile (Government Accountability Office, 2006), I find those purchases have a large, immediate cost and have no measurable effect on prices when the oil is eventually released. Second, crude oil releases from the SPR should not be relied on to lower oil prices. A policymaker facing a spike in oil prices should explore other policy options to lower the price of crude oil.

The remainder of the paper is organized as follows. Section 1.2 provides an overview of the related literature. Section 1.3.1 presents a benchmark VAR model of the U.S. oil market with SPR policy variables. Sections 1.3.2 and 1.3.3 show that the results of the benchmark model hold under alternative identifying assumptions that incorporate external information. Section 1.4 presents an uncertainty model with time-varying policy effects. Concluding comments are contained in Section 3.4.

1.2 Literature Review

The increasing frequency of oil supply disruptions from the 1950s to the 1970s led economists to propose a government controlled crude oil reserve to protect the economy from insecure oil suppliers (Cabinet Task Force on Oil Import Control, 1970; National Petroleum Council, 1973; Nordhaus, 1974). These calls for the U.S. government to create a publicly managed crude oil reserve intensified during the Arab Oil Embargo (Tolley and Wilman, 1977). The theory behind these policy proposals is simple: the government could offset supply disruptions by increasing oil supply with the SPR rather than by restricting oil demand with unpopular quotas, tariffs, or taxes. Since oil price spikes had consistently preceded economic downturns in the postwar economy (Hamilton, 1983), a policy mechanism that controlled oil prices had broad political appeal (Weimer, 1982).

Much of the early work on the SPR focused on optimal management strategy (Teisberg, 1981; Balas, 1981; Chao and Manne, 1983; Samouilidis and Berahas, 1982). Most of these papers assumed that crude oil released from the SPR during a supply disruption would lower prices, but crude oil purchases during normal market conditions would not significantly raise prices. These assumptions could not be tested with data available at the time. Few additional papers were written on the topic until China began building its strategic oil reserve in 2004, and economists estimated optimal management strategies for China's strategic reserve using similar, untested assumptions (see Wu et al., 2008; Zhang et al., 2009; Bai et al., 2012).

Only a few papers have directly estimated the effect of strategic reserves on crude oil prices. Considine (2006) used a simultaneous equation model of the SPR to estimate its effect on oil prices between 1992 and 2005. Consistent with the assumptions in the literature, Considine found that filling the reserve increased prices imperceptibly (0.4 percent), while releasing oil during a supply shock can lower prices by 3.5 percent. Other estimates, which use models of the forward price curve, find that SPR purchases and releases have a symmetric effect on prices and can change the price of oil by as much as 32 percent (Verleger, 2003).

The key difficulty in the empirical literature is to isolate the effect of SPR policy from the state of the oil market to which the policy responds. I contribute to the literature by using a structural VAR model of the U.S. oil market to disentangle these effects. Beginning with Killian's (2009) structural VAR model of the global oil markets, VAR models have been used to study country-level oil markets (Peersman and Van Robays, 2009, 2012; Baumeister and Peersman, 2013a; Kilian and Murphy, 2014), volatility in oil markets (Van Robays, 2012; Baumeister and Peersman, 2013b; Jo, 2014), and storage in oil markets (Kilian and Murphy, 2014). I extend these models by explicitly including oil market policy variables in a VAR model of the U.S. oil market.

To estimate the effect of SPR policy in a VAR framework, I rely on the extensive work on VAR models of monetary policy (see Christiano et al., 1999). In this literature, the effects of monetary policy have been estimated using many different identification strategies and estimation methods (Sims, 1986; Bernanke and Blinder, 1992; Leeper et al., 1996; Bernanke and Mihov, 1998). Applying the time series models developed in this field to oil markets allows me to answer the question, "How do oil prices respond to unanticipated SPR purchases and releases?" with several sets of plausible identifying assumptions.

This paper also applies identification methods recently developed to incorporate external information in VAR models (Stock and Watson, 2008; Olea et al., 2012; Stock and Watson, 2012). Beginning with Romer and Romer (1989), who used minutes from Federal Reserve meetings to identify unanticipated monetary policy actions, many papers have constructed exogenous instruments to identify structural shocks (Ramey and Shapiro, 1999; Barro and Redlick, 2011; Stock and Watson, 2012; Mertens and Ravn, 2013). However, oil market VAR models are typically identified using recursive timing or sign restrictions. These models rely on restrictions on the coefficients or residuals in the model–what Stock and Watson (2008) call "internal restrictions." I construct two separate external instruments to partially identify the effects of SPR purchases and releases in a VAR framework, without internal restrictions.

A final contribution of this paper is an analysis of SPR policymaking under uncertainty. The uncertainty literature has shown the effects of economic policy vary with uncertainty (Bloom, 2009; Bekaert et al., 2013; Baker et al., 2013; Aastveit et al., 2013). Specifically, the stimulative effect of monetary policy can be dampened by high economic uncertainty when producers face non-convex adjustment costs. I extend VAR models of oil market uncertainty

(Jo, 2014; Baumeister and Peersman, 2013a) by including SPR policy variables and policy-uncertainty interaction terms. This model allows me to show that unlike other economic policies, the effect of SPR purchases is amplified by uncertainty.

1.3 Vector Autoregression Models

1.3.1 Benchmark Model

In this section, I use a system of equations to model the U.S. oil market that include SPR policy variables. Endogeneity between SPR policy and other oil market variables complicates estimation. The exogenous components of SPR policy, SPR policy shocks, are identified using restrictions on the contemporaneous effects of oil market variables. These exogenous policy shocks are then used to estimate the effect of SPR policy on crude oil prices.

The SPR policy function is modeled as:

$$\mathrm{SPR}_t = P(\Omega_t) + \varepsilon_t^{\mathrm{SPR}} \tag{1}$$

where SPR_t represents the government's oil market policy instruments: crude oil purchases and crude oil releases. Following Christiano et al. (1999), SPR policy at time t is related to the policymaker's time t information set, Ω_t, by some linear function P. The portion of SPR policy that reflects the policymaker's systematic reaction to changes in the oil market is given by $P(\Omega_t)$. The residual component of SPR policy that is not accounted for by changes in the oil market is the SPR policy shock, $\varepsilon_t^{\mathrm{SPR}}$. This unanticipated component of policy is assumed to be exogenous to the information set, Ω_t.

In the remainder of the paper, I rely on this distinction between SPR policy *shocks* and SPR policy *actions* (Bagliano and Favero, 1999). A SPR policy action is determined by the state of the oil market through the policy function. A SPR policy shock is the component of observed SPR policy that is not predicted by the policy function. Since SPR policy shocks are deviations from the policy rule that are exogenous to the state of the oil market, they can be used in causal analysis. Observed SPR policy is the sum of a SPR policy action and a SPR policy shock.

There are several economic interpretations of these SPR policy shocks (Christiano et al., 1999). The shocks could reflect an exogenous change in policymaker preferences. For example, a new presidential administration could change the weight given to oil supply shocks relative to oil demand shocks in their policy function. An exogenous policy shock could also reflect the policymaker's response to market expectations of SPR policy. If the President announced an SPR release due to private sector expectations, rather than oil market conditions, that would reflect an exogenous SPR shock in this model. Finally, policymakers may rely on preliminary oil market data that is measured with error. The exogenous variation in the policymaker's reaction function induced by measurement error would be an exogenous policy shock. Regardless of interpretation, these shocks are the exogenous component of SPR policy that will be used to identify the effects on the oil market.

To estimate SPR policy shocks, the vector of oil market variables, Y_t, is partitioned as follows,

$$Y_t = \begin{bmatrix} Y_t^{\text{slow}} \\ \text{SPR}_t \\ Y_t^{\text{fast}} \end{bmatrix} \tag{2}$$

These vectors are contained in the policymaker's information set, Ω_t. I assume the variables in the Y_t^{slow} vector do not respond contemporaneously to SPR policy shocks, and SPR policy does not respond immediately to the variables in the Y_t^{fast} vector. But, the response of the variables in the Y_t^{fast} vector to policy shocks is unconstrained. These recursive timing restrictions are the key assumptions to identify policy shocks in this model, but will be relaxed in Sections 1.3.2 and 1.3.3.

The U.S. oil market is modeled with variables that capture on crude oil supply, demand, and prices (Kilian, 2009). The Y_t^{slow} vector contains the oil supply and demand variables for the U.S. market. Oil supply is the sum of domestic oil production and imports, expressed as a log difference. This measure of oil supply captures domestic and foreign supply disruptions. Oil demand is measured by the Aruba, Diebold, and Scotti (ADS) business conditions index published by the Philadelphia Federal Reserve (Aruoba et al., 2009). This index combines high- and low-frequency economic activity measures into a daily estimate, and is regularly used to estimate economic activity at high-frequencies (Berge and Jordà, 2011; Andreou et al., 2013; Monteforte and Moretti, 2013).

The SPR_t vector, plotted in Figure 1, contains separate policy variables for SPR purchases and releases. Since the SPR opened in 1977, more than 850 million barrels have been purchased over 769 weeks, which covers 46 percent of the sample period. There have been 19 separate SPR releases that have taken place over 243 weeks, about 15 percent of the sample period. The variables in the SPR_t vector are the log differences of the SPR inventory during purchases and sales, respectively. Finally, the Y_t^{fast} vector contains the oil price variable. The price of crude oil is measured by the log spot price of West Texas Intermediate (WTI) crude, deflated by the Consumer Price Index.

The vector of oil market variables is modeled as a finite order structural VAR:

$$A_0 Y_t = \alpha + \sum_{i=1}^{40} A_i Y_{t-i} + \varepsilon_t \tag{3}$$

where ε_t are mean zero and serially uncorrelated structural shocks,

$$E(\varepsilon_t | Y_{t-1}, Y_{t-2}, \dots) = 0 \tag{4}$$

$$E(\varepsilon_t \varepsilon_t') = \Sigma_\varepsilon = \begin{bmatrix} \sigma_1^2 & & 0 \\ & \ddots & \\ 0 & & \sigma_5^2 \end{bmatrix} \tag{5}$$

This benchmark model is estimated with 40 lags using weekly data from August 1, 1983 through October 5, 2014.[1] The 9 month lag length is consistent with Hamilton and Herrera (2004), who show the importance of long lag length in models of oil markets. Other VAR models of oil markets use lag lengths ranging from 4 months (Baumeister and Peersman, 2013b; Jo, 2014) to 24 months (Kilian, 2009; Kilian and Murphy, 2014).

The reduced form of the model,

$$Y_t = \underbrace{A_0^{-1}\alpha}_{\beta} + \sum_{i=1}^{40} \underbrace{A_0^{-1}A_i}_{B_i} Y_{t-i} + \underbrace{A_0^{-1}\varepsilon_t}_{e_t} \tag{6}$$

can be estimated with standard least-squares methods. However, the reduced form coefficients (β's) and residuals (e_t) cannot be used to estimate the causal effect on SPR policy on oil prices because each of the reduced form residuals is a weighted average of all structural shocks. Parameter values from the structural model are required to estimate causal effects.

The coefficients of the structural model (A_i's) cannot be recovered, but the structural shocks (ε_t) can be estimated using the reduced form residuals. I assume A_0, and by extension A_0^{-1}, is lower triangular, so the reduced form errors can be decomposed using $e_t = A_0^{-1}\varepsilon_t$,

$$e_t \equiv \begin{bmatrix} e^{\text{oil supply}} \\ e^{\text{econ activity}} \\ e^{\text{SPR purchase}} \\ e^{\text{SPR release}} \\ e^{\text{oil price}} \end{bmatrix} = \begin{bmatrix} a_{11} & 0 & 0 & 0 & 0 \\ a_{21} & a_{22} & 0 & 0 & 0 \\ a_{31} & a_{32} & a_{33} & 0 & 0 \\ a_{41} & a_{42} & a_{34} & a_{44} & 0 \\ a_{51} & a_{52} & a_{35} & a_{45} & a_{55} \end{bmatrix} \begin{bmatrix} \varepsilon^{\text{supply shock}} \\ \varepsilon^{\text{demand shock}} \\ \varepsilon^{\text{purchase shock}} \\ \varepsilon^{\text{release shock}} \\ \varepsilon^{\text{price shock}} \end{bmatrix} \tag{7}$$

[1]Though the SPR purchases began in 1977, weekly data on the SPR are only available beginning in 1983. These results are robust to including weekly SPR data interpolated from the monthly 1977–1982 data series.

The restrictions on A_0^{-1} are called the *recursiveness* assumption (Christiano et al., 1999). This assumption identifies the model by limiting the simultaneous responses among variables. For example, oil supply and demand are assumed to take at least one period to respond to SPR policy shocks.

The recursiveness assumption used to identify the elements of A_0^{-1} cannot be tested directly in this VAR model; rather, its validity comes from economic theory. As shown in equation 2, these exclusion restrictions divide the oil market variables into three categories: slow-moving variables, policy variables, and fast-moving variables. As in Killian (2009), oil supply and oil demand are slow-moving with respect to other variables in the model. U.S. oil supply (the sum of domestic production and imports) is fixed at the weekly level with respect to aggregate demand, SPR policy, and oil price shocks. This assumption is consistent with the high costs of short-term changes to oil production and import schedules. Similarly, U.S. oil demand (domestic economic activity) can respond immediately to oil supply shocks, but does not respond within a week to SPR policy or oil price shocks. These assumptions are reflected in the zeroes in the first two rows of A_0^{-1}.

SPR policymakers can respond immediately to unanticipated shocks to oil supply and economic activity, but are assumed to respond to oil price innovations with a one week lag. Since SPR policy decisions are the products of meeting-filled political processes, this assumption is reasonable. Purchases are usually scheduled well ahead of actual delivery and releases typically take more than one week to approve. The fastest SPR response to a specific oil supply disruption occurred on September 2, 2005 in response to Hurricane Katrina, seven days after the Hurricane entered the Gulf of Mexico and oil prices spiked. The zeroes in the rows three and four of A_0^{-1} reflect these assumptions. Finally, the fast-moving variable of the model, the price of oil, can immediately respond to all shocks, hence there are no zeros in the last row of A_0^{-1}.

Table 1 summarizes the SPR policy shocks over time. During the initial SPR fill (1983–1989), the average purchase shock was somewhat negative (-0.02) and the average release shock was negligible. During this period, crude oil deliveries to the SPR were slowed because of construction delays, which made the market overestimate the amount of oil to be purchased. And the only crude oil release during this period was a small test sale in 1985. Negative purchase shocks continued into the 1990s, though the SPR releases during the first Gulf War (1990–1993) were somewhat larger than markets anticipated. In the 2000s, purchase shocks were nonnegative. During that period, SPR purchases were consistently larger than market expectations and SPR release shocks were mixed.

The impulse response functions in Figures 2 and 3 display the causal effect of an unanticipated SPR policy shock over time, holding all other structural shocks constant at zero.

The 95 percent confidence interval is included in the figures, and is constructed with the semiparametric fixed-design wild bootstrap, accounting for conditional heteroskedasticity of an unknown form (Goncalves and Killian 2004). The paths plotted in these figures are robust to reordering the variables within Y_t^{slow} and SPR_t vectors.

Figure 2 shows the impulse response function for a one-standard deviation unanticipated SPR purchase. SPR purchase shocks have a sustained, statistically significant effect on the real price of oil which peaks at 1.5 percent 20 weeks after purchase before declining. However, Figure 3 shows that an unanticipated SPR release has a small (less than 0.25 percent) negative effect on the real price of oil that is not statistically significant. These estimates can be scaled to give the oil price response in terms of dollars per barrel released or purchased for the SPR. The peak effect of a one-million barrel unanticipated SPR purchase raises the price of a $100 barrel of oil by $1.76, while an unanticipated SPR release of the same magnitude lowers the price of a $100 barrel of oil by $0.28.

1.3.1.1 The SPR and Private Crude Oil Stocks Early work on the SPR pointed out that increases in public oil storage could influence behavior of private oil storers (Wright and Williams, 1982). If private crude oil stocks declined in response to SPR acquisitions, or rose in response to SPR releases, the effect of SPR policy could be dampened. The gradual decline in private petroleum storage since the SPR opened in 1977 suggests that the SPR growth has been offset by private storage declines. But, there has been no empirical work on the relationship between the SPR and private storage. I use the policy shocks from the VAR model to estimate the causal effect of SPR purchases and releases on private stocks.

Firms purchasing oil from, or for, the SPR pay the transportation cost to, or from, the delivery sites in the Gulf Coast. The SPR is directly connected to the oil pipeline network that runs throughout the Midwest (PADD 2) and the Gulf Coast (PADD 3).[2] Transporting oil to, or from, the SPR through these pipelines is fast and relatively inexpensive. However, other PADDs can only move oil to, or from, the SPR by ship, which is more expensive and takes 8 to 24 days. In the most recent SPR releases, over 90 percent of SPR sales were delivered to oil companies in the Gulf Coast or Midwest. Likewise, oil purchased for the SPR tends to come from oil companies in PADD 2 and PADD 3. For this reason, I examine the effects of SPR purchases and releases on PADD 2 and PADD 3 stocks.

PADD 2 and PADD 3 stock data are only available at the monthly frequency over the entire 1983 to 2014 time period. Following Kilian (2009), I convert the weekly policy shock

[2]Data published by the Energy Information Administration (EIA) divides the U.S. into five Petroleum Administration for Defense Districts (PADD): East Coast (PADD 1), Midwest (PADD 2), Gulf Coast (PADD 3), Rocky Mountain (PADD 4), and West Coast (PADD 5).

series estimated in the VAR model into a monthly series, $\hat{\eta}_m$, by averaging the weekly shocks by month,

$$\hat{\eta}_m^{\text{Purchase}} = \frac{1}{4}\sum_{j=1}^{4} \hat{\varepsilon}_{j,m}^{\text{Purchase}} \tag{8}$$

$$\hat{\eta}_m^{\text{Release}} = \frac{1}{4}\sum_{j=1}^{4} \hat{\varepsilon}_{j,m}^{\text{Release}} \tag{9}$$

where the week j structural shock in month m is given by $\hat{\varepsilon}_{j,m}$. These monthly policy shocks can be used to estimate the causal effect of SPR policy on private stocks if they are exogenous with respect to private stocks. The process used to construct the weekly shocks in Section 1.3.1 guarantees endogeniety with respect to all variables in Y_t. Since private stocks are not in the information set, additional assumptions are required to use the monthly shocks to estimate the causal effect on stocks.

Identifying the causal effect of SPR policy on private stocks requires the assumption that the monthly policy shocks are exogenous with respect to private stocks. In other words, there can be no feedback between policy shocks and private stocks within a month. Though this cannot be directly tested, evidence can be provided by comparing unanticipated policy changes to innovations in private stocks. Again, following Kilian (2009), I model private stocks as an autoregressive process, and find the residuals are uncorrelated with SPR policy shocks. Since this correlation is low, I assume SPR policy shocks are fixed at the monthly level with respect to private stocks.

The response of private stocks, s_t, to policy shocks, $\hat{\eta}_m$, is modeled as,

$$s_t = \gamma + \sum_{m=1}^{9} \phi_m \hat{\eta}_m + \theta_m \tag{10}$$

where ϕ_m is the month m stock response and θ_m is the residual. The coefficients in these regressions give the response of private crude oil stocks in PADD 2 and PADD 3 to SPR purchase and release shocks. The impulse response functions based on these coefficients are plotted in Figures 4 and 5, along with the 95 percent confidence interval.

Unanticipated SPR purchases have no measurable effect on private stocks in PADD 2 and PADD 3. Given the transportation costs discussed above, unanticipated SPR purchases are likely filled by domestically produced or imported oil and not private oil stocks in PADD 2 and PADD 3. However, private storers respond to unanticipated SPR releases by raising stocks nearly 2 percent, and this effect is statistically significant at the 90 percent level. This model shows that private storers dampen the immediate effect of SPR releases by increasing their reserves.

1.3.2 Identification Using External Instruments

Though the estimated SPR policy effects are robust to alternate specifications that include private stocks and reordering variables within Y_t^{slow} and SPR_t, this approach is still subject to Rudebusch's (1998) critique of VAR model of policy shocks. Namely, the policy shocks estimated from linear, time-invariant functions with a small information set may not be related to actual unanticipated policy changes. Alternate identifying assumptions used to model the oil market, like sign restrictions, are subject to the same critique. In Sections 1.3.2 and 1.3.3, I depart from the oil market literature by using external information, which is related to unanticipated policy changes, to identify SPR policy shocks (Stock and Watson, 2008). These instruments confirm the key insights from this benchmark model: unanticipated SPR releases do not have a measurable effect on oil prices but SPR purchase shocks have a statistically significant, positive effect on oil prices.

The recursiveness assumption identifies unanticipated policy changes by restricting the transmission of structural shocks in the oil market. If, for example, oil production responds immediately to policy shocks, structural shocks could not be identified. Even if the recursiveness assumption holds, the SPR policy function in equation 1 could be misspecified by omitting variables. If policymakers used unmodeled variables in their "true" policy function, the policy shocks estimated in Section 1.3.1 would not be exogenous. In either case, the policy shocks estimated in Section 1.3.1 would confound the effects of SPR policy changes and the oil market conditions that cause the policy changes, and could not be used to estimate the causal effect of policy shocks (Rudebusch, 1998).

To address these issues, I partially identify the VAR presented in equation 6 with an external instrument for SPR purchases. This identification method follows the narrative approach developed by Romer and Romer (1989, 2004, 2010) to identify monetary policy shocks, which has been used regularly in subsequent research (see Ramey and Shapiro, 1999; Barro and Redlick, 2011; Olea et al., 2012). The omitted variables and simultaneity critiques do not apply to the policy shocks identified in this model if the instrument is correlated with unanticipated purchases and uncorrelated with other shocks. This approach only identifies the effects of SPR purchase shocks; in the next section I use a separate instrument to identify SPR release shocks.

Identifying the causal effect of SPR purchases on oil prices requires an instrument with the following properties: the instrument is not included in Y_t, the instrument is correlated with unanticipated purchases, and the instrument is uncorrelated with contemporaneous oil market conditions (excluding policy effects). The SPR purchase instrument exploits the procedures used by the Department of Energy (DOE) to acquire crude oil for the SPR. From 1999 to 2010, DOE did not directly purchase crude oil from the SPR, rather they negotiated contracts with private companies to purchase and deliver crude oil to the SPR (U.S. Senate, 2003; Congressional Research Service, 2006).

The DOE's crude oil acquisition procedure gives the purchase schedule two useful properties: the schedule is correlated with unanticipated SPR purchases but uncorrelated with other structural shocks. The DOE deliberately kept oil market participants in the dark about the timing and amount of SPR purchases. If information about the purchase schedule were available to market participants, speculators could front-run the contractors and drive up prices. Since it was in the interest of the DOE and its contractors to purchase crude cheaply, it is reasonable to assume that the purchase schedule is correlated with unanticipated purchases.

The purchase schedule is assumed to be exogenous with respect to oil market conditions at the time of purchase because the schedule was set well in advance of actual purchases. The DOE rarely granted contractors exemptions from the schedule when the oil market tightened, so contractors were forced to provide oil for the SPR despite their requests to postpone purchases (Government Accountability Office, 2006).

These data are collected and coded from several official sources. To start, I obtained data on DOE purchase contracts from 1999 to 2010 through a Freedom of Information Act request. I confirm these data using the documents submitted pursuant to the 2003 Senate hearing on the SPR (U.S. Senate, 2003) and information on the DOE website. In these documents, a purchase schedule specifies the total amount of crude oil to be purchased and the time period of delivery. For example, in early 2005, DOE contracted for 9.4 million barrels of oil to be delivered from April to July 2005. These schedules are converted into a weekly series by assuming purchases are divided evenly over the delivery window.[3]

Between 1999 and 2010, the SPR grew by 154 million barrels and I have contracts for 126 million of those barrels.[4] Figure 6 plots SPR purchases along with the constructed purchase schedule, which shows that the purchase schedule has reasonable predictive power for SPR purchases. A regression of SPR purchases on the purchase schedule gives an R-squared of

[3]The delivery window gives contractors some flexibility in purchasing crude oil. However, the physical infrastructure of the SPR limits the maximum amount of oil that can be delivered in a single day. Typical crude oil deliveries take place over several months, even when filling the SPR at the maximum daily rate.

[4]The DOE also acquires crude oil for the SPR via short-term oil loans that are repaid with interest. The repayment of loans are not subject to the same procedures as SPR purchases and are not included in this instrument.

0.35.

To identify the effect of SPR purchases on oil prices using this instrument, I estimate the model in equation 6 with standard least-squares methods. Then, following the approach outlined by Stock and Watson (2008), I partition the data into the into SPR purchases (S_t) and the other variables (X_t),

$$Y_t = \begin{bmatrix} X_t \\ S_t \end{bmatrix} \tag{11}$$

where X_t is a (4×1) vector and S_t is a scaler. The order of variables within X_t is not relevant because the recursive timing assumptions (used in Section 1.3.1) are not used in this model. The structural errors and reduced form residuals are similarly partitioned,

$$e_t = \begin{bmatrix} e_t^X \\ e_t^S \end{bmatrix}, \varepsilon_t = \begin{bmatrix} \varepsilon_t^X \\ \varepsilon_t^S \end{bmatrix} \tag{12}$$

Analogous to the A_0 matrix in the benchmark model, I assume there exists some matrix Q, such that $Qe_t = \varepsilon_t$, where Q can be partitioned as

$$Q = \begin{bmatrix} Q_{XX} & Q_{XS} \\ Q_{SX} & Q_{SS} \end{bmatrix} \tag{13}$$

Solving $Qe_t = \varepsilon_t$ for the reduced form errors yields,

$$e_t^X = -Q_{XX}^{-1} Q_{XS} e_t^S + Q_{XX}^{-1} \varepsilon_t^X \tag{14}$$

$$e_t^S = -Q_{SS}^{-1} Q_{SX} e_t^X + Q_{SS}^{-1} \varepsilon_t^S \tag{15}$$

There are four steps to recover the structural purchase shock, ε_t^S in equation 15, from this system of equations. First, the purchase schedule (Z_t) is used as an instrument for e_S^t in equation 14. This gives an unbiased estimate of $-Q_{XX}^{-1} Q_{XS}$ because the instrument is, by assumption, uncorrelated with all other structural shocks ($E(Z_t \varepsilon_t^X) \neq 0$) and correlated with unanticipated purchases ($E(Z_t e_t^S) \neq 0$). Next, ε_t^x is estimated up to scale using, $\widehat{\varepsilon_t^X} = e_t^X + \widehat{Q_{XX}^{-1} Q_{XS}} e_t^S$. This estimate of ε_t^X is then used as an instrument for e_t^X in equation 15, which gives an estimate of $-\widehat{Q_{SS}^{-1} Q_{SX}}$. Finally, exogenous purchase shocks are estimated, up to scale, using $\widehat{\varepsilon_t^S} = e_t^S + \widehat{Q_{SS}^{-1} Q_{SX}} e_t^X$.

With these estimated structural shocks, the dynamic causal effect of unanticipated SPR purchases on the price of oil can be estimated. The impulse response function plotted in

Figure 7 shows the positive, statistically significant effect of SPR purchases on oil prices identified using the purchase instrument is quite similar to the effect identified with the recursiveness assumption. Both effects are estimated to peak around 20 weeks, though the effect estimated with the instrument is somewhat larger than the effect estimated in the benchmark model–6 percent as opposed to 1.5 percent. In this partially identified model, a one-million barrel unanticipated SPR purchase raises the price of a $100 barrel of oil by $7.12.

Though the SPR purchase effect identified with the instrument is larger than the purchase effect identified in Section 1.3.1, it is important to keep in mind that this instrument is only available between 1999-2010. During that period, oil prices jumped sharply, then collapsed during the financial crisis and recession. Excluding the commodity price boom and financial crisis from the sample, by dropping data from 2006 to 2010, lowers the peak effect on oil prices to about 3.5 percent.

It is also important to note that this method assumes that there is no heterogeneity among SPR purchase shocks. That is, an unanticipated SPR purchase shock during the 2003 Iraq War has the same effect as an unanticipated purchase while oil markets were calm in the late 1990s. If purchase shocks are heterogenous, then this model estimates the *local* average treatment effect of SPR purchase shocks, not the average treatment effect. I address the issue of heterogenous policy shocks in Section 1.4, where the effect of SPR policy is modeled as a function of oil market uncertainty.

1.3.3 Identification Using Daily Futures Market Data

As a robustness check of the SPR release effect estimated in Section 1.3.1, I use crude oil futures market data to identify SPR release shocks. This section offers a different approach to the simultaneity and omitted variables problems discussed in Section 1.3.2. Recall that the identification restrictions used in the benchmark model do not allow SPR policy to respond within a week to oil price shocks. If policymakers respond immediately to oil price changes, the release effect estimated in Section 1.3.1 would be biased. The benchmark VAR model also assumes a simple SPR policy function that only takes into account oil supply, demand, and price. All other factors that determine SPR policy (e.g., political opinion polling or gasoline prices) are omitted. Though monetary policy functions are similarly modeled (Christiano et al., 1999), there are concerns that these policy functions omit key variables used by policymakers (Rudebusch, 1998).

Following Bagliano and Favero (1999) and Cochrane and Piazzesi (2002), I use daily oil futures data to overcome these difficulties. Crude oil futures market expectations incorporate far more information about oil markets than could be modeled in a VAR. This identification strategy uses changes in crude oil futures prices immediately after SPR release announcements to estimate whether the policy announcement was unanticipated. This structural shock estimate is then used to construct an impulse response function.

This model for estimating unanticipated SPR releases has two key assumptions: intraday changes in the crude oil futures price measure SPR policy shocks and the SPR release shock is the only oil market shock at the time of announcement. The first assumption is based on the idea that futures prices aggregate information relevant to oil prices and futures prices only respond to unanticipated information. If the market anticipated some portion of the SPR release prior to announcement, I assume that would be included in the futures price. The second assumption is based on the idea that at a small enough time interval, the SPR release announcement is the only shock to the oil market. This assumption may not hold if the SPR release announcement reveals information about the government's view of the oil market. This method places no other restrictions on the determinants of oil price expectations or oil market dynamics.

Figure 8 illustrates this method with two examples from the WTI futures market around two SPR release announcements: the 1996 deficit reduction sales and the 2011 Libya sale. On April 25, 1996, one day prior to the deficit reduction sale announcement, the one-month WTI futures contract closed at $14.20 (Figure 8b). The futures price did not respond to the April 26 release announcement, closing the day at $14.28, which suggests the futures market anticipated this release. Since this SPR sale was ordered by Congress to reduce the deficit and the release amount and approximate timing had been discussed publicly for months, it is not surprising that the release was already priced into futures contracts.

On June 22, the day before the 2011 SPR release announcement, the WTI futures market closed at $42.27. Following the 9:30 am release announcement, the futures price dropped sharply, and the market closed at $40.32 on June 23 (Figure 8a). This release came at the end of the Arab Spring after Libyan oil production dropped sharply. Though there was speculation about the release, the size and timing were not clear prior to the announcement. The nearly $2 drop in the crude oil futures price suggests this policy announcement was unanticipated by the market.

The validity of the assumptions relies on the choice of futures contract. Two financial instruments are used to estimate the structural SPR release shocks. The first is the one-month WTI futures price. An unanticipated SPR release occurs when the WTI futures price falls following the announcement. Since the movements in the WTI futures price could reflect any oil market supply or demand shock, I use another financial instrument to isolate the SPR effect: the one-month US-Canadian crude futures (WTI-WCS) spread. The U.S.

and Canadian oil markets are similar and subject to similar shocks. However, only the U.S. has a strategic oil reserve, and Canadian oil companies cannot purchase oil from the SPR. Differencing U.S. and Canadian oil should remove the effects of non-SPR supply and demand shocks to the North American oil market, which allows me to directly estimate the magnitude of the SPR release shock.

The US-Canadian crude futures spread identifies unanticipated releases that were not followed by a decline in the WTI futures price. For example, the WTI futures price edged up $0.40 (1.6 percent) following the Hurricane Ivan release (Figure 9). This SPR announcement would not be considered an SPR release shock using the WTI futures price. However, the WCS futures price rose $2.05 (5.3 percent) following the release announcement. Since the futures trends pre- and post-announcement are nearly identical, the change in the futures spread indicates the WTI futures price rose less than expected. Since WTI became less expensive, relative to WCS, following an SPR release announcement, the release was unanticipated according to the WTI-WCS spread.

There are two steps to estimate the causal effect of the SPR releases in this model. First, SPR release shocks are estimated for the entire sample of 19 SPR releases and the shock variable is set to zero on weeks with no announcement.[5] I construct the structural shock variable by taking the difference between the futures price at market close the day prior the announcement and the futures price at market close the day of the announcement.[6] Using the WTI futures contract, seven SPR releases were unanticipated, while the WTI-WCS futures spread estimates that eight of the releases were unanticipated. Second, I regress the change in the spot price of crude oil, $P_{t+j} - P_t$, on the series of estimated structural shocks, $\varepsilon_t^{release}$. The coefficients of this regression give consistent estimates of the causal effects of unanticipated SPR releases on the price of oil, which I use in the impulse response functions.

Figures 10 and 11 plot the impulse response functions using the SPR release shocks estimated using WTI futures and the WTI-WCS futures spread, respectively, along with the 95 percent confidence interval. Consistent with the results from the benchmark model in Section 1.3.1, unanticipated SPR releases do not have a statistically significant, negative effect on the price of oil using either futures contract. In fact, the effect of unanticipated SPR releases increases the price of oil nearly 2 percent using shocks identified with the WTI-WCS futures spread. This is similar to the results from Cochrane and Piazzesi (2002), who show that unanticipated monetary policy shocks, as measured by futures prices, that are intended

[5] To ensure that the release announcement is the only effect on crude futures, I drop any SPR release announcements that occur during OPEC meetings or on days when the Energy Information Administration releases oil market data.

[6] To allow for information leaks preceding the announcement, the two-day change was also used. The results are qualitatively identical.

to drive down interest rates can have the opposite effect.

1.4 Uncertainty Model

In this section I explore the mechanism driving the SPR purchase and release effects. I begin by testing whether the effect of SPR policy varies with uncertainty. The growing empirical literature on uncertainty has shown that uncertainty shocks dampen the effects of monetary and fiscal policy (Bloom, 2009; Bekaert et al., 2013; Baker et al., 2013; Aastveit et al., 2013). These results are built on real options theory which has shown that uncertainty lowers the value of options when investments are not reversible (Dixit and Pindyk, 1994). This section presents a framework for SPR policy effects to vary with oil market uncertainty by augmenting the model VAR in Section 1.3.1 with an oil market uncertainty variable and policy-uncertainty interaction terms.

I cannot directly measure oil market uncertainty, so I use several different variables to capture uncertainty regarding oil prices, economic policy, and economic activity. The 90-day average volatility of the WTI spot price is my primary measure of oil market uncertainty. This common index of oil market uncertainty is the 90-day moving average of the standard deviation of intraday spot oil price changes. The volatility of oil prices is an imperfect measure of oil market uncertainty, so I also use other indices of policy and overall economic uncertainty that have been proposed in the literature. Baker et al. (2013) construct an index of policy uncertainty using a variety of news-based uncertainty measures. This variable is designed to capture uncertainty regarding economic policy. I also use an index of stock market uncertainty, VXO, which estimates market volatility based on S&P 500 options prices, to capture general economic uncertainty. These uncertainty indices are plotted in Figure 12.

It is important to note that oil market uncertainty is not always associated with high prices. For example, oil price volatility spiked following the 2007 financial crisis, but oil prices were at their lowest level in five years. SPR releases have also not been restricted to times of high-uncertainty. Crude oil has been released through several different programs, including emergency drawdowns, infrastructure test sales, sales to reduce the deficit, and short-term leases. Some sales have occurred during severe global oil market supply disruptions with high uncertainty (e.g., 1991 Desert Storm sale), while others have occurred during localized supply disruptions with low uncertainty (e.g., 2004 Hurricane Ivan release). SPR releases have also been conducted while oil prices were relatively low and stable in 1985 to test the SPR infrastructure and in 1996 to reduce the Federal deficit. Likewise, crude oil purchases have occurred during periods of high and low oil market uncertainty. This variation in SPR policy and oil market volatility is used to estimate whether the effect of SPR policy on oil price fluctuates with uncertainty.

I formalize this approach by including oil market uncertainty and SPR policy-uncertainty interaction terms in the benchmark VAR model from Section 1.3.1. These uncertainty terms add flexibility to the model allowing the effect of oil market interventions to vary over time with uncertainty, which can shed light on puzzling asymmetric response of oil prices to SPR purchases and releases. Unlike other studies of oil markets with time varying coefficients (Baumeister and Peersman, 2013b), the time varying effects of SPR policy shocks are driven by a specific exogenous variable, uncertainty.

Following Aastveit et al.'s (2013) work on monetary policy under uncertainty, I include an exogenous uncertainty variable, V_t, and uncertainty interaction terms for SPR purchases and releases in the oil price equation:

$$A_0 Y_t = \alpha + \sum_{i=1}^{40} A_i Y_{t-i} + C_i V_t \, Y_{t-i} + D V_t + \varepsilon_t \tag{16}$$

where V_t is the uncertainty measure and $V_t \, Y_t$ is the policy-uncertainty interaction term. To study the effect of uncertainty on the SPR policy effect on prices, the coefficients on both the volatility and volatility interaction variables are constrained as follows,

$$C = \begin{bmatrix} 0_{4\times 2} & 0_{4\times 1} & 0_{4\times 1} & 0_{4\times 1} \\ 0_{1\times 2} & c_1 & c_2 & 0 \end{bmatrix}, D = \begin{bmatrix} 0_{4\times 1} \\ c \end{bmatrix} \tag{17}$$

The C matrix limits the uncertainty interaction terms to the SPR policy variables in the oil price equation. The D matrix limits the uncertainty terms to the oil price equation. Given the sample size, these constraints are necessary to restrict the number of variables in the model.

The reduced form of this model,

$$Y_t = \underbrace{A_0^{-1}\alpha}_{\beta} + \sum_{i=1}^{40} \underbrace{A_0^{-1}A_i}_{\beta_i} Y_{t-i} + \underbrace{A_0^{-1}C_i}_{G_i} V_t \, Y_{t-i} + \underbrace{A_0^{-1}D}_{H} V_t + \underbrace{A_0^{-1}\varepsilon_t}_{e_t} \tag{18}$$

is estimated with standard least-squares methods. As with the VAR model in Section 1.3.1, the coefficients in the structural model (A_i, C_i, and D) cannot be recovered, but the reduced form residuals (e_t) can be converted into structural shocks (ε_t) with recursive timing restrictions. These estimated shocks are used to construct impulse response functions.

The SPR policy effects shown in the impulse response functions in Figures 13 and 14 use oil market volatility to capture uncertainty and are similar to those estimated in the benchmark model without uncertainty.[7] In both models, SPR purchases increase the price of oil 1 to 2 percent and SPR releases have a small negative, but not statistically significant, effect on oil prices.

The reduced form coefficients can be used to estimate the effect of SPR purchases and releases during periods of high and low uncertainty. Having estimated equation 18, I use the coefficients for V^{high} and V^{low} to simplify the model as follows,

$$Y_t^{high} = \underbrace{\widehat{D_0}^{high}}_{\widehat{B}+\widehat{H}V^{high}} + \sum_{i=1}^{40} \underbrace{\widehat{D_i}^{high}}_{\beta_i+G_i} Y_{t-i} + \widehat{e}_t \tag{19}$$

$$Y_t^{low} = \underbrace{\widehat{D_0}^{low}}_{\widehat{B}+\widehat{H}V^{low}} + \sum_{i=1}^{40} \underbrace{\widehat{D_i}^{low}}_{\beta_i+G_i} Y_{t-i} + \widehat{e}_t \tag{20}$$

These equations allow the impulse response functions to be calculated with fixed uncertainty levels. Specifically, V^{high} is defined as the 90$^{\text{th}}$ percentile volatility and V^{low} is defined as the 10$^{\text{th}}$ percentile volatility.

SPR releases have no measurable effect on oil prices at any either high or low uncertainty, regardless of the uncertainty measure. However, Figures 15 through 20 show the effect of SPR purchase shocks on oil prices increases with uncertainty, across all uncertainty measures. When volatility is low, unanticipated SPR purchases do not have a statistically significant effect on prices (Figures 15, 16, and 17). This confirms claims by the Department of Energy that crude oil purchases during periods of low oil market volatility do not affect oil prices (U.S. Senate, 2003). However, when uncertainty is high, SPR purchase shocks increase the price of oil from 2 percent, using political uncertainty, to 4 percent, using oil market or economic uncertainty (Figures 18, 19, and 20). Unfortunately, SPR purchases tend to occur during periods of high oil market volatility.

1.5 Conclusion

This paper analyzes SPR policy in a time series model of the U.S. oil market. I begin with the observation that much of the SPR literature has focused on optimal SPR management and does not directly estimate the effect of SPR policy on oil prices. Rather, these papers assume that SPR crude oil releases lower the price of oil during supply disruptions while purchases only raise the price slightly, if at all. Contrary to these assumptions, I show in a structural VAR model with recursive timing restrictions that an unanticipated SPR release

[7] The impulse response functions estimated using the stock market volatility and political uncertainty indices are nearly identical and omitted.

has no measurable effect on oil prices, while an unanticipated SPR purchase increases the price of oil 1.5 percent over 20 weeks following purchase.

I explore alternate identification methods for SPR policy using information that is not contained in the VAR model. Oil market VAR models are typically identified by restricting the timing or signs of structural shocks (Kilian, 2009; Baumeister and Peersman, 2013a,b; Jo, 2014; Kilian and Murphy, 2014). I show that a VAR model of the oil market can be partially identified under different, less restrictive assumptions. I begin by constructing an instrument for SPR purchases from the purchase schedules set by the Department of Energy. These schedules are set months ahead of actual purchases, making them exogenous with respect to other oil market shocks at the time of purchase. The purchase schedules were not immediately announced publicly and exemptions were rarely granted, so the schedules are correlated with unanticipated SPR purchases. This instrument identifies SPR purchase shocks in the VAR framework without other restrictions on the structural shocks. An unanticipated SPR crude oil purchase, identified with the purchase schedule, increases the price of crude 6 percent over 15 weeks following purchase.

I use a different model to identify the effect of SPR releases with high-frequency crude oil futures data. The change in crude oil futures prices immediately following release announcements is used to measure whether SPR releases are unanticipated by the market. Estimating policy shocks with this method avoids problems with endogeneity and omitted variables that have been raised concerning policy shocks estimated in traditional VAR models (Rudebusch, 1998). Several different crude oil futures contracts are used to isolate the unanticipated policy shocks. This model confirms that unanticipated SPR releases do not have a statistically significant effect on oil prices.

Finally, I explore a mechanism driving the SPR purchase and release effects by testing whether the effect of SPR policy varies with uncertainty. I augment the structural VAR with an oil market uncertainty measure and policy-uncertainty interaction terms. This allows the effects of SPR policy to vary over time with uncertainly. I find that SPR policy shocks only impact oil prices when purchases are made during periods of high uncertainty. SPR purchase shocks increase oil prices 2 to 4 percent when uncertainty is high, depending on the uncertainty measure. Unanticipated purchases when uncertainty is low, and releases at any time, do not have an effect on prices.

The main contribution of this paper is to develop a framework for estimating the effect of U.S. oil market interventions on oil prices. I show that under a variety of identifying assumptions, SPR releases have not lowered the price of oil, while SPR purchases have increased its price. These results have strong implications beyond the U.S. and global oil markets. Over half of global crude oil stocks are held in government-controlled reserves in more than 40 countries. Many of these countries are following the U.S. by expanding their government-controlled crude oil reserves (International Energy Agency, 2014). In light of these results, strategic reserve acquisitions should be limited to periods of low oil market uncertainty to minimize the price impact, and unanticipated releases from these reserves should not be expected to lower oil prices.

2 Public vs. Private Good Research at Land Grant Universities[8]

2.1 Introduction

The land-grant university system, which is the centerpiece of the agricultural science establishment in the United States, has been one of the most successful innovations in the history of education (Kerr, 1987; Rasmussen, 1989). The future viability of the system is, however, in jeopardy. Bruce Gardner has shown that while the U.S. agricultural sector was once fairly uniform, composed largely of family farms, and the benefits of new technologies were widely distributed, its benefits have become increasingly more concentrated while its costs have remained widely dispersed (Gardner, 2002). Since the Smith-Lever legislation augmenting the land-grant university system in 1914, the farm share of the population has declined from 33 percent to 2 percent and the farm size distribution has become highly skewed.[9] Correspondingly, the agricultural science establishment has received a declining share of the public research budget, roughly in line with the farm sector's declining share of overall economic activity: agriculture received almost 40 percent of the federal R&D budget in 1940 but by 2007 this share had declined to 1.4 percent.

Since the rate of return on agricultural $R\&D$ has been estimated at well over 40 percent (Alston et al., 2000), it is not surprising that private agricultural research has more than offset declining government funds (Huffman and Evenson, 1993). Many land-grant administrators, looking to supplement their shrinking public budgets, have found eager partners in the private sector to participate in research collaborations and partnerships. These administrators have become increasingly successful forming industry research partnerships; in 2006 land-grant universities accounted for more than half of the top 20 schools in attracting industry $R\&D$ sponsorship.

This corporate sponsorship has caused controversy in some states where deans and directors have tended to "direct" faculty to perform short-term profit-enhancing service (Beattie, 1991). Moreover, choice of research projects is often controlled by internal competitive grants, and/or the administration of extension has been separated from the administration of teaching and research, presumably to concentrate more directly on serving special interests. In other states, a rift has developed between extension and teaching/research faculty. As a consequence, extension activities have been slow to evolve beyond the farm sector in

[8]Joint research with Gordon Rausser and Leo Simon

[9]The largest 7.1 percent of farms produce 75.1 percent of output value and the smallest 78.7 percent produces only 6.8 percent of output value (Department of Agriculture, 2002).

the same way that research activities have. In both cases, the fundamental premise of the land-grant system is violated as research conducted for the private sector has tended to "crowd-out" public good research (Greenberg, 2007).

University administrators increasingly encourage researchers to replace declining formula funding by research grants, many of which are motivated by private interests either directly through private grants or indirectly through public research grants generated by private lobbying and pork-barrel politics. As a consequence, researchers' marginal research time and the generation of ideas are increasingly focused more on specific private interests. The ultimate concern is that private interests will leverage their funds to crowd-out public good research at land-grant universities.

The motivations to acquire private funding have been amplified by changes in intellectual property (IP) ownership and any royalty distribution within the university Washburn (2005). Following the passage of the Bayh-Dole Act, IP rights were assigned to the universities and research scientists have typically received some fraction of the royalty stream associated with the commercialization of their research. This allocation of royalty revenues provides university scientists with incentives to pursue lines of research that are likely to lead to commercially profitable discoveries. Moreover, since private industry funds are directed to commercially appropriable research, university researchers can increase the likelihood that their research will be licensed by pursuing private sponsorship. Again, this research funded by the private sector crowds-out public-good research that is not generated elsewhere. Under-emphasized research products include fundamental scientific knowledge and, in the field of social science, research on new institutions and policies, analysis of labor displacement effects of new technologies, and safety and environmental research on new biotechnologies and chemicals.

As land-grant universities work more collaboratively with private interests, questions are raised about the need for continued public funding and, more fundamentally, public justification for the land-grant research system. Some argue that privatization of university activities offers new potential for encouraging socially relevant research and facilitating transfer of technology. However, without proper policies and incentives, universities can become pawns of powerful private interests, and the unique and separable contribution that universities can make to the public good may be lost (Just and Rausser, 1993). Public-private partnerships cannot be allowed to leverage university resources and divert research from public-good outputs not produced elsewhere.

The capture of land-grant universities can be expected to lead to increased public criticism and possibly more dramatic reductions in funding. However, capture by the private sector is not the inevitable outcome of public-private research collaboration. In fact, it is conceivable that with proper policies and incentives, universities can use public-private partnerships to leverage industry resources to crowd-in public good research. Without such partnerships, there is little prospect that the private sector will replace the public-good research that would otherwise take place at land-grant universities. Sufficient incentives simply do not exist for the private sector to effectively replace fundamental investments in public research. Even though private investment in agricultural research is substantial, it quite obviously is aimed toward applied commercial research with transparent potential for profitability: chemicals, hybrid or genetically engineered seeds.

In this paper, we concentrate on the bargaining that must necessarily take place in establishing public-private joint ventures and the incentives among the two agents: research administrators and private industry representatives. The focus is on the potential for crowding-in or crowding-out of public-good research. In addition to private-sector incentives, a governance structure is specified for university research administrators. Crowding-in and/or crowding-out is shown to depend critically on the structure of the bargaining problem between these two parties.

Two alternative frameworks will be used to model university research: static and dynamic with feedback effects. The single-shot static research framework assumes research funding only affects the sponsored research, that is, there are no knowledge spill-overs. This 'worst-case scenario' in which there is no feedback between public good and private good research is overly simplistic. An alternative framework is then employed that admits the nonlinearities and feedback between public good research and applied research that characterizes the research process at land-grant universities. By accounting for knowledge spill-overs this alternative framework blurs the distinction between between public good research and commercial research and, as a result, also blurs the boundary between public land-grant universities and the private sector research.

2.2 Model I: Single-Period With No Feedback Stock Accumulation

At the center of our model is a public research institution. To fix ideas we'll label this research institution "the university." However, the model could apply with few, if any, modifications to a wide variety of public research institutions at any level of governmental organization.

The university produces two kinds of research: *theorems* and *mousetraps*. *Theorems*, denoted by k, are basic research that generates public goods. As such, they can not be appropriated for direct commercial benefit. *Mousetraps*, denoted by m, are technologies and products resulting from applied research which are private goods that can be fully appropriated for direct commercial benefit.

In this model the technology available to the university to produce new theorems and mousetraps is simply a function of expenditure, denoted by e. These production technologies have the associated cost function $C(m, k) = m^{\beta_m} + k^{\beta_k}$, where $\beta_m, \beta_k > 1$. This cost function gives rise to a production possibilities frontier for the university which describes the set of feasible combinations of theorem and mousetrap research for a given level of expenditure, $\bar{e} \geq m^{\beta_m} + k^{\beta_k}$. This production possibilities frontier is shown in figure 21.

The university's research allocation decisions are made by an administrator who is given a *performance function* by the university's regents. This function is used to evaluate the quality of the administrator's decisions. Initially we presume a simple linear performance function, viz.,

$$P(m, k) = \mu m + (1 - \mu)k \qquad \mu \in [0, 1] \tag{21}$$

This performance function gives the administrator μ *performance units* for each mousetrap she produces and $(1 - \mu)$ *performance units* for each theorem she produces. This function establishes a relative price between mousetraps and theorems equal to the ratio $p = \mu/(1 - \mu)$. For a fixed level of expenditure, the administrator will select the research mix along her production possibilities frontier defined by the usual condition that the rate of product transformation between mousetraps and theorems just equal their relative price, p. We consider later the effect of generalizing the performance measure to allow for imperfect substitution between mousetraps and theorems.

To produce any research the university must obtain funding. The university has both public (government) and private commercial company funding sources.[10] Government funding, g, is presumed to be costlessly and exogenously obtained. Funding from the commercial company, denoted by f, is given in exchange for property rights over the produced mousetraps and is obtained through a bargaining process.[11] The outcome of this bargaining process is an (m, k) pair which is produced using funding from the commercial company and government, $C(m, k) \leq f + g$. If we normalize the price of mousetraps to unity,[12] the commercial

[10] We broaden our analysis to consider investment dynamics in the second model and consider a third funding source, viz., the royalties associated with any university mousetraps that have been commercialized.

[11] In reality the university would not fully cede its property rights to the mousetraps, but would instead negotiate a royalties agreement. This complication is addressed in the second model.

[12] We assume the commercial company can costlessly transform the university's mousetrap research into a marketable product. For simplicity, we assume every unit of the mousetrap research produced by the university can produce a single unit of marketable product for the commercial company. Relaxing this assumption by allowing for a costly, nonlinear transformation of mousetrap research does not significantly

company's profit function is given by

$$\pi(m,k) = m - f \tag{22}$$

In the current static model, any investment in theorems is worthless to the commercial company; they only capture value from the mousetraps. However, depending on the bargaining outcome, they may also have to produce funding support for some theorems.

We use the Nash cooperative bargaining solution concept to solve this problem.[13] Undeniably, the Nash solution is a good place to start due to its simplicity. In particular, we will be able to neatly define a precise measure of what we mean by "crowding-out" or "crowding-in" and provide simple geometric interpretations.

Recall that the Nash solution is computed as the point in the bargaining set, $\mathcal{B}$, that maximizes the product of the players' utility gains from cooperating:

$$(\hat{m}, \hat{k}) = \arg\max\Big\{\big[\pi(m,k) - \pi_d\big]\big[P(m,k) - P_d\big] : \forall (m,k) \in \mathcal{B}\Big\} \tag{23}$$

π_d and P_d represent the disagreement outcomes (or threat points) for the commercial company and university, respectively. In the event of disagreement, the commercial company gets nothing—its threat point π_d is zero. In the event of disagreement, the university still has a funding level of g from governmental grants and can produce any allocation in the feasible set labeled g in Fig. 2. The iso-performance line associated with this level of governmental funding is labeled P_d. The administrator would allocate this funding in accordance with the price, p, induced by its performance function. Let the chosen allocation be denoted by (m_g, k_g). The threat point for the university is $P_d = \mu m_g + (1-\mu)k_g$.

For the problem in 23 to be well-defined, we must specify the bargaining set $\mathcal{B}$. $\mathcal{B}$ is the set of efficient points that lie between the ideal research allocations for the university and company. This set is constructed as follows. First consider commercial company's iso-profit curves labeled π in 22. The profit associated with these iso-profit lines is decreasing along the vertical axis. The commercial company's ideal point is on the iso-profit curve with the highest profit, which is the allocation (m_c^*, k_c^*), where $k_c^* = 0$ and the level of mousetraps is defined by the standard marginal condition $\frac{df}{dm} = 1$. The university has no global ideal point, as its preferences are monotonic, but does have a local ideal point for every iso-profit line. This local ideal point is determined by the tangency between the administrator's iso-

alter our conclusions.

[13]There are good reasons to doubt the validity of this solution concept for modeling collective decision making processes in general. We won't discuss these here, but instead refer the reader to Rausser, Simon and van't Veld (1995).

performance line and the company's iso-profit line. For the commercial company to fund university research, the administrator must offer an allocation that gives the company at least its zero-profit disagreement outcome, π_d. Such allocations are found in figure 22 along the iso-profit line labeled π_d that begins at the origin. The university's local ideal point along this iso-profit curve is the allocation (m_u^*, k_u^*) that lies at the tangency with its iso-performance line. The bargaining set $\mathcal{B}$ is the set of all such efficient points between this local ideal point for the university, (m_u^*, k_u^*), and the company's global ideal point, $(m_c^*, 0)$. This set of points is also the *core* of the bargaining problem, in which the solution must lie.

It is relatively straightforward to map this bargaining set from $m - k$ space to utility space to produce (a slight variant of) the well-known Nash bargaining picture, as depicted in figure 23.[14]

2.2.1 Measuring the Degree of Crowding-Out

As noted in the introduction, a basic concern of this paper is the allocation of research expenditures between public good research (theorems) and commercial applications (mousetraps). There is concern that as a result of reduced government funding, public good research will be neglected and research will be expanded that is more short-term in orientation, and more geared toward commercial applications. This paper attempts to shed some light on the relationship between research priorities and the availability of public funding for university research. In particular, we use the model to investigate the claim that privately funded research can crowd-out or, in the alternative, crowd-in fundamental research, and that this can occur even when public funding becomes scarcer.

Two notions of crowding-out are easily captured, one a *research ratio* measure, the other a *negotiating leverage* measure. The research ratio measure focuses on the ratio of expenditure on mousetrap research to expenditure on theorem research. This ratio will be referred to as the *m-k ratio.*

Definition: *A fall in the level of government funding exacerbates the research ratio measure of crowding-out if it results in a increase in the m–k ratio. That is, research ratio crowding-out exists if* $\frac{d(m/k)}{dg} < 0$. *Research ratio crowding-in exists if* $\frac{d(m/k)}{dg} > 0$*, i.e. a decrease in the level of government funding decreases the* $m - k$ *ratio.*

This is a very natural notion that seems, at first blush, to capture the essence of the crowding-out metaphor. The first result from our model is that it exhibits research ratio crowding-out.

[14]To maintain concavity of the Nash product requires some mild restrictions on the cost functions which essentially guarantee that the contract curve does not rise out of the commercial company's ideal point too quickly.

Proposition 1: *In the context of the specified model, a decline in the level of government funding will increase the ratio of mousetraps to theorems produced by the university.*

The intuition for this result is quite straightforward. Under extremely weak restrictions, publicly funded research in our model will generate a lower mousetrap to theorem ratio than privately funded research. The $m - k$ ratio compares *total* expenditure on mousetrap research to *total* expenditure on theorem research, aggregating across privately-funded and publicly-funded activities. As public research funds decline, so does the contribution of its relatively theorem-rich research mix to the aggregate mix, and so the $m - k$ ratio declines.

The above result is hardly surprising. Avoiding this research ratio form of crowding-out seems to be too much to expect. While administrators might wish that private agencies would share their enthusiasm for fundamental research, they can hardly expect this result. More realistically, administrators should expect that some dilution of the purity of their research activities is a necessary price that they must pay if they are to augment their public funds with private ones. For this reason, the research ratio measure of crowding-out result is of little practical use as it tells us nothing about how policy makers could restructure incentives to ameliorate the deleterious effects of crowding-out.

For the remainder of this section we focus on a more subtle notion of crowding-out, which is intimately connected to the view that the research fund raising process should be viewed as a bargaining problem. The fundamental issue is: how will the decline in government funding affect the relative bargaining strength of the university administrator as she enters into negotiations with private funding sources for research contracts. A very natural concern is that as government funding declines, the administrator's need for funds will become more desperate, and hence more willing to compromise in the bargaining process in order to secure funding. In this context, compromise will take the form of skewing research contracts in favor of mousetraps rather than theorems.

To operationalize this idea, we introduce the following negotiation leverage measure of crowding-out. The very best that the administrator can hope for in a negotiation with a private funding source is to drive the other bargainer to her reservation utility, i.e. to extract *all* of the surplus from the bargaining relationship so that the other party is just indifferent as to whether or not an agreement is reached. We have already defined this as the administrator's local ideal point, (m^*_u, k^*_u). In general, of course, such a good bargain will never be struck, and the commercial company will secure some of the surplus. A natural way to measure the university's negotiation leverage, then, is to consider the gap between the realized bargaining outcome and the best possible outcome, as a fraction of the total potential surplus that is available to the administrator from the bargaining relationship Rausser and Zusman (1992).

Definition: *Formally, this measure of the university's leverage is:*

$$\zeta = \left[\frac{P^* - \hat{P}}{P^* - P_d}\right] \tag{24}$$

where, as illustrated in Fig. 3, P^* is the highest utility that the administrator can obtain from the bargaining relationship, $\hat{P}$ is the utility she obtains from the actual solution to the bargaining and P_d is the administrator's default utility.

This measure of leverage ranges between zero and one. When $\zeta = 0$, all the negotiation leverage resides with the administrator, while if $\zeta = 1$, the administrator is completely at the whim of the commercial company.

Given this definition of leverage, a negotiation leverage measure of crowding-out may be defined that relates directly to the incentive structure of the underlying bargaining problem that can generate the crowding-out phenomenon:

Definition: *The degree of crowding-out is measured by the change in the administrator's negotiation leverage as government funding falls. If* $\frac{d\zeta}{dg} = 0$, *we say the bargaining problem is* neutral*. The bargaining problem exhibits* negotiation leverage crowding-out *if* $\frac{d\zeta}{dg} < 0$. *Conversely, the bargaining problem exhibits* negotiation leverage crowding-in *if* $\frac{d\zeta}{dg} > 0$.

2.3 Neutrality

The key issue to be addressed is: can a decline in the level of government research funding increase the university administrator's negotiation leverage? We begin by constructing a "neutrality" result. That is, we identify conditions under which the extent of negotiation leverage crowding-out is independent of the level of government funding. The theorem below is intended as a benchmark rather than as a positive result. Because the three assumptions upon which the result depends are all extremely restrictive, we can conclude that in reality, the extent of crowding-out is indeed quite sensitive to the level of government funding.

Proposition 2: *The following three separability conditions are necessary and sufficient for the bargaining problem to exhibit neutrality, i.e. for the degree of negotiation leverage crowding-out to be independent of the level of government funding.*

1. Benefit Separability*: The commercial company only has property rights over the mousetraps that result from research that it funds. The university retains property rights over the mousetraps it produces with government funds.*

2. Cost Separability*: The cost* functions *for privately and publicly funded research are independent of each other.*

3. Performance Separability*: Theorems and mousetraps are separable in the administrator's performance function.*

Under condition (1), the commercial company derives utility from any mousetraps resulting from the research program that it funds, but derives none from mousetraps resulting from publicly funded research. More generally, this condition reflects the idea that association with the university provides the commercial company with no externality whatsoever. In reality, of course, this is not the case. Private funding sources generally gain a great deal from these associations, over and above the utility that the research outputs generate. These benefits take many forms, ranging from prestige and the benefits of public visibility through access to research ideas and potential employees at all levels.

An implication of condition (2) is that the marginal cost of producing an additional mousetrap or theorem depends *only* on the number of mousetraps and theorems already being produced *within the given research program* as opposed to depending on the outputs of other research projects. In particular, privately-funded and publicly-funded research programs cannot compete for the same scarce resources. Once again, this is an extremely restrictive assumption. In fact, it is typically the case that some fraction of the private research funds are earmarked for operating expenses rather than for infrastructure, so that at the very least, the university's opportunity cost of publicly funded research *does* increase with the level of privately funded research.

Condition (3) will be satisfied if and only if theorems and mousetraps are perfect substitutes in the administrator's performance function. To the extent that the performance function is intended as a realistic proxy for "social welfare," this condition is just as restrictive as any other assumption about perfect substitutability.

The proof of this neutrality result is illustrated by figure 24. The figure depicts the two bargaining problems confronting the administrator and the commercial company, with different levels of public funding, as exact translates of each other. The neutrality result then follows immediately from the fact that the Nash Bargaining Solution is translation invariant.[15]

The key intuition for the proof lies in the demonstration that the problems are indeed translates of each other. To see this, note that under assumption (1), the commercial company obtains no benefit whatsoever from the outputs of publicly-funded research. Nor do its costs of doing business depend in any way on the level of public research. Hence, from the commercial company's point of view, all aspects of the bargaining problem are completely unchanged as the level of public funding decreases from g to g'. Similarly, the administrator

[15]A monotonic decrease in government funding simply shifts the bargaining problem and the properties of the Nash Bargaining Solution do not change.

benefits from publicly-funded research by exactly the same amount, irrespective of whether or not it reaches an agreement with the commercial company, and irrespective of the nature of this agreement. In particular, note that by condition (3), the manner in which the administrator allocates public funds between theorem and mousetrap research is completely independent of whether or not, and how, it negotiates with the commercial company.

2.4 Nonneutrality

In this section, we relax each of the assumptions of Proposition 2 in turn, and consider the effect of reduced government funding on the degree of crowding-out. We begin by omitting condition (1). For concreteness, imagine that the university administrator is authorized to offer to the commercial company property rights over all government-funded mousetraps as a "side-payment" that may induce more private participation. Specifically, consider a bargaining contract which assigns the commercial company property rights to *every* mousetrap produced, provided that *some* agreement is reached. If no agreement is reached, then the commercial company gets nothing.

As government funding declines, the size of the "pot" of bonus mousetraps available for side-payments also declines. What is the implication of this for crowding-out?

Proposition 3: Nonseparable Benefits *If conditions (2) and (3) hold but not (1), then a decrease in the level of government funded research may either increase or decrease the degree of crowding-out.*

This indeterminacy is at first sight unexpected. Intuitively, it would seem obvious that the larger is the pot of bonus mousetraps, the more the commercial company has to lose in the event that it fails to reach an agreement. It would seem that this device for enticing the commercial company into a bargaining relationship should strengthen the administrator's hand in the bargaining process. After all, the more she brings to the table, the more leverage she has in bargaining with the commercial company. Hence, we would expect that a decrease in the level of government funding would *decrease* the degree of crowding-out. Clearly, this is not the case, however.

The intuition for Proposition 3 is provided by figure 25 and figure 26 for a nonlinear university performance function. To highlight the problem, we consider two very extreme cases in which the university administrator's iso-performance lines are convex. Note however that the bargaining problems illustrated in these figures do not correspond to the problem at hand. In particular, their Pareto loci are not smooth. On the other hand, it will be apparent to the reader that we could construct smoothed versions with the same properties that would be consistent with our model. In figure 25, the original bargaining frontier is the line *abe*. After government funding has been reduced, the new frontier is *cde*. The explanation for this difference is that before government funding is cut, the university can offer a side-payment to the commercial company that the company values at $\pi^* - \pi^{*\prime}$ but the university values at zero. The commercial company receives this side payment in full, provided *some* agreement is reached, but receives nothing in the event of disagreement. Note that the bargaining frontier is initially downward-sloping but then vertical.[16] After the cut in government funding, government-funding research dries up entirely, and the administrator cannot offer the commercial company a side-payment. Consequently the entire bargaining frontier shifts down by precisely the amount $\pi^* - \pi^{*\prime}$. Note that because of the extreme specification of the frontier, the the maximal utility level that the administrator can obtain from the bargaining relationship, P^*, is unaffected by the decline in funding. On the other hand, the *realized* outcome shifts back down along the ray through the default outcome, so that $\hat{P}$ exceeds $\hat{P}'$. Clearly, this example exhibits crowding-out.

Now consider figure 26. The structure is exactly parallel except that the bargaining frontier starts out horizontal and becomes downward sloping. In this case, however, P^* exceeds $P^{*\prime}$ while $\hat{P}$ is unaffected by the decline in funding. This example exhibits crowding-in.

Since the "intuitive" result proved to be false, the reader will not be surprised to learn that the effects of relaxing the other two conditions are also indeterminate. First, consider what happens when costs are nonseparable. For concreteness, assume that research costs for mousetraps (similarly theorems) strictly increase with the *aggregate* number of mousetraps (theorems) produced. Clearly, in this case an issue of cost-sharing arises: should private or public funding have access to the flatter part of the cost curve? To simplify the computations, we shall assume that public research has first access to the technology. That is, we will assume that public research costs are independent of private research activity, but the reverse relationship does not hold. In extensions to this work, we will consider the effect of including cost-sharing as an additional dimension along which bargaining can occur.

Proposition 4: Nonseparable Costs *If conditions (1) and (2) hold but not (3), then a decrease in the level of government funded research may either increase or decrease the degree of crowding-out.*

Finally, consider the effect of introducing curvature into the administrator's performance measure. Specifically, assume that the administrator's performance measure is Cobb Douglass in *aggregate* mousetraps and *aggregate* theorems. In this case, intuition strongly suggests a determinate result. We have already noted in Proposition 1 that as the level of public

[16]While this extreme specification is useful for heuristic purposes, the argument clearly holds for more general specifications.

funding declines, the *aggregate* theorem to mousetrap ratio will decline also. Once we move to a Cobb Douglas performance measure, a decline in public funding will have the effect of increasing the administrator's "hunger" for theorems. Being more "needy" in this regard, we would expect that the administrator's bargaining position would be weakened by the decline, and that crowding-out would increase. Once again, however, this intuition proves to be spurious. Specifically,

Proposition 5: Nonseparable Performance *If conditions (1) and (3) hold but not (2), then a decrease in the level of government funded research may either increase or decrease the degree of crowding-out.*

2.5 Model II: Multi-Period With Feedback Effects

The success of the land-grant university system is attributed to two-way interaction whereby research scientists are informed of field problems and relevant solutions based on scientific developments are communicated to the field (Bonnen, 1986). Phenomenal productivity has been due to dispersing and commercializing these technologies widely in the agricultural sector of the economy (Ruttan, 1982; Evanson et al., 1979). Moreover, the system has supplied products with public good characteristics such as new crop varieties, improved breeding stock, and improved management practices. These products have been easily reproducible and thus have not lent themselves to private market development and appropriation. The following model focuses on the joint dependencies between basic and applied research.

In our dynamic model, the production functions of mousetraps and theorems are dependent on the stocks of both basic and applied research. Since money spent on research in the current period increases the stock of theorems available for future research, this funding should be considered investment. The creation of theorems will also increase the stock of knowledge available to the creators of mousetraps; thus the stocks are public goods. However, there is also potential feedback from the accumulated mousetrap stocks to new discoveries in basic science.

Each period the university administrator chooses m_t and k_t and output is determined by:

$$M_t = \alpha M_{t-1} + m_t + h(K_{t-1}) \tag{25}$$

$$K_t = \gamma K_{t-1} + k_t + l(M_{t-1}) \tag{26}$$

Now the number of mousetraps (theorems) produced in period t is a function of the stock of mousetraps (theorems) available to researchers, M_{t-1}, mousetrap (theorem) research funding, m_t, and feedback from the stock of theorems (mousetraps), $h(K_{t-1})$. The feedback effects have the following properties for all $M, K > 0$:

$$h(K), l(M) > 0 \tag{27}$$

$$h', l' > 0 \tag{28}$$

$$h'', l'' < 0 \tag{29}$$

Conditions (27), (28), and (29) capture the strictly positive, convex nature of feedback effects in university research.

Each period the administrator's performance is evaluated by a performance function, now stated in terms of the stocks of both applied (M_t) and basic (K_t) knowledge, i.e.

$$P_t(M_t, K_t) = \mu M_t + (1-\mu)K_t \qquad \mu \in [0,1] \tag{30}$$

Thus over T periods the administrator's objective is to maximize

$$P = \sum \beta^t P_t + \beta^T \psi(M_T, K_T) \tag{31}$$

where β is the administrator's discount factor and $\psi(M_T, K_T)$ is the value to the administrator of research conducted during the final year. This objective function establishes the relative price between mousetraps (applied research) and theorems (basic research) in every period is $p = \mu/(1-\mu)$.

As in the single-period model, the university must obtain funding to sponsor research, but now the university has three potential funding sources: the government, a commercial company, and royalties from the university's mousetraps that have been commercialized by the company. As before, government funding, g_t, is presumed to be costlessly and exogenously obtained. Private funding, denoted by f_t, is obtained through a bargaining process that also determines royalties, $\mathcal{R}(M_t)$. The outcome of this bargaining process is a pair of vectors ($m = m_1, \ldots, m_T; k = k_1, \ldots, k_T$) and a royalty rate, $\mathcal{R}(M_t)$: the company agrees to provide a level of funding sufficient to produce a portion of this allocation, f_t, and to pay the university some fraction of the revenue from the commercialization of mousetrap research from the previous period as a royalty. We assume the administrator will only have $\theta \in [0,1]$ of the royalties paid by the company at their disposal, while the remaining $(1-\theta)$ is distributed to other groups in the university and has no impact on research. Note, the price of mousetraps is again normalized to unity.

The production of theorems and mousetraps in period t again has the following cost function: $C_t(m_t, k_t) = m_t^{\beta_m} + k_t^{\beta_k}$, where $\beta_k, \beta_m > 1$. This cost function gives rise to a production possibilities frontier for the university which describes the set of feasible combinations of theorem and mousetrap research for a given level of expenditure in period t,

$$C(m_t, k_t) \leq g_t + f_t + \mathcal{R}(M_{t-1}) \tag{32}$$

where $\mathcal{R}(M_{t-1})$ are the royalties paid by the company to the university in period t for the commercialization of mousetrap research in period *t-1*.

If a commercial company reaches a funding agreement with the university, then each period that company sells mousetraps licensed from the university, funds university research, and pays royalties based on the commercialization of mousetrap research in the previous period. The company's profit in period t is therefore

$$\pi_t(m_t, k_t) = M_t - f_t - \mathcal{R}(M_{t-1}) \tag{33}$$

where the private funding level in period t, f_t, is equal to the university's cost of producing the company's share of the total research output. The company's total profit is the sum of the profits from each period and the scrap value of an agreement, $W(M_T)$, discounted by the company's weighted average cost of capital, r:

$$\Pi = \sum_{t=0}^{T-1} (\frac{1}{1+r})^t \pi(m_t, k_t) + (\frac{1}{1+r})^T W(M_T) \tag{34}$$

Again, we use the Nash cooperative bargaining solution concept to solve this problem. Using the multi-period extension of the Nash bargaining framework from the first model, the solution is the pair of vectors $\hat{m} = \hat{m_1}, \ldots, \hat{m_T}$ and $\hat{k} = \hat{k_1}, \ldots, \hat{k_T}$ in the bargaining set, $\mathcal{B}$, that maximize the product of the players' utility gains from cooperating:

$$\begin{aligned}(\hat{m}, \hat{k}) = \text{arg max}[\Pi(m_1, \ldots, m_T, k_1, \ldots, k_T) - \Pi_d \\ P(m_1, \ldots, m_T, k_1, \ldots, k_T) - P_d \\ : \ \forall \ (m_1, \ldots, m_T, k_1, \ldots, k_T) \in \mathcal{B}]\end{aligned} \tag{35}$$

Again, Π_d and P_d represent the disagreement outcomes for the commercial company and university, respectively. The company earns no revenue if it does not reach an agreement with the university, since no mousetraps are sold, and has no costs, since no research is funded, so its threat point, Π_d, is zero. The university still has a funding level of g from the government in the event of disagreement, which it will allocate in accordance with the price, $p = \mu/(1-\mu)$, induced by its performance function. Let the allocation chosen by the administrator for a funding level of g be denoted by $(m_{g_1}, \ldots, m_{g_T}; k_{g_1}, \ldots, k_{g_T})$. In the event no agreement is reached, the university administrator's performance function will be $P_d = \sum_{t=0}^{T-1} \beta^t[\mu M_{g_t} + (1-\mu)K_{g_t}] + \beta^T \psi(M_{g_T}, K_{g_T})$.

A new bargaining set, $\mathcal{B}$, must be defined for this multi-period game for (35) to be well-defined. The company's ideal allocation is now the set of points $({m_{c_1}}^*, \ldots, {m_{c_T}}^*, {k_{c_1}}^*, \ldots, {k_{c_T}}^*)$ which are determined by the first-order conditions given in the appendix. Again, the university administrator has no global ideal point and must only offer the company an allocation that gives the company at least its disagreement outcome, if there is to be an agreement. The set of efficient points between the company's global ideal allocation and disagreement outcome form the bargaining set $\mathcal{B}$, in which the solution must lie.

2.6 Measuring Crowding-Out With Feedback

We will use the two notions of crowding-out from section 2.1 to determine the conditions under a research partnership with the private sector can improve a university's public good research. Recall, the first notion, the research ratio measure of crowding-out, focuses on the ratio of applied research to basic research. We extend the single-period definition to this multi-period model by comparing the ratio of *total* expenditure on mousetrap research, $(m_1 + \ldots + m_T)$, to *total* expenditure on theorem research, $(k_1 + \ldots + k_T)$.

Definition: *A fall in the level of government funding exacerbates the research ratio measure of crowding-out if it results in an increase in the M–K ratio. That is, crowding-out exists if* $\frac{d((m_1+\ldots+m_T)/(k_1+\ldots+k_T))}{dg} < 0$.

Though the first-order conditions, see appendix, indicate the profit-maximizing company will want to invest in some basic research because of the feedback loops, this model still exhibits research ratio crowding-out.

Proposition 6: *For the specified model, there exists $\delta, r > 0$, such that for all $\beta > 1 - \delta$, a decline in the level of government funding will increase the ratio of total expenditure on mousetraps to theorems produced by the university.*

The university administrator's linear performance function guarantees that for any $\mu \epsilon [0, 1]$, the administrator values the production of theorems more than the commercial company, which leads the administrator to always prefer more basic research than the company (see appendix). Once again, public funding leads to a relatively theorem-rich research mix and, as private funding increases, relatively fewer theorems are produced. Even with large feedback effects the private sector will not value basic research as much as the administrator.

As before, the negotiation leverage measure of crowding-out, in which private fund raising is viewed as a bargaining problem, focuses on change in the university administrator's leverage in her bargaining negotiations with the private sector as government funding falls. Note that we extend the negotiation leverage measure of crowding-out to this multi-period problem by simply using the weighted sum of the administrator's utility in the local ideal, actual, and default bargaining outcomes:

$$P^* = \sum_{t=0}^{T-1} \beta^t(\mu M^*_{u_t} + (1-\mu)K^*_{u_t}) + \beta^T \psi(M^*_{u_T}, K^*_{u_T}) \tag{36}$$

$$\hat{P} = \sum_{t=0}^{T-1} \beta^t(\mu \hat{M}_{u_t} + (1-\mu)\hat{K}_{u_t}) + \beta^T \psi(\hat{M}_{u_T}, \hat{K}_{u_T}) \tag{37}$$

$$P_d = \sum_{t=0}^{T-1} \beta^t(\mu M_{g_t} + (1-\mu)K_{g_t}) + \beta^T \psi(M_{g_T}, K_{g_T}) \tag{38}$$

The university's leverage, 24, is again measured as the distance between the realized bargaining outcome and the local ideal outcome, as a fraction of the administrator's total potential surplus from the bargaining outcome.

Again, *negotiation leverage crowding-out* occurs if $\frac{d\zeta}{dg} < 0$ and *negotiation leverage crowding-in* occurs if $\frac{d\zeta}{dg} > 0$. If $\frac{d\zeta}{dg} = 0$, the bargaining problem is *neutral*.

6. Neutrality With Feedback Two more neutrality conditions are needed to establish the "neutrality" result for the multi-period model with two feedback loops. These assumptions, like conditions (1), (2), and (3) from Proposition 2, are extremely restrictive, so the extent of crowding-out in this multi-period model is likely sensitive to the level of governmental funding.

Proposition 7: *When conditions (4) and (5) are added to conditions (1), (2), and (3), the separability conditions are necessary and sufficient for the multi-period bargaining problem to exhibit neutrality, i.e.* $\frac{d\zeta}{dg} = 0$.

4. Equal Weighting: *The university administrator's performance function assigns equal weight to basic and applied research. That is,* $\mu = 1/2$.

5. Equal Feedback: *The feedback functions are equal,* $h(\cdot) = l(\cdot)$.

Conditions (4) and (6) ensure the symmetry required for the neutrality result. Under condition (4) the administrator derives an equal amount of utility from the production of basic and applied research. Because the administrator's performance function is a proxy for "social welfare," constraint (4) on μ is just as restrictive as the other substitutability assumptions. Condition (5) ensures the administrator's production possibilities frontier is symmetric about the 45° line. Though necessary to guarantee neutrality of governmental funding, this assumption is highly unlikely to be satisfied in practice.

Figure 27 illustrates why conditions (4) and (5) are needed to prove Proposition 7. The symmetry imposed by these neutrality conditions ensures that the administrator's leverage does not change as governmental funding changes. With feedback effects, a decrease in funding from the government has two effects: conditions (1), (2), and (3) guarantee the end points of the production possibilities frontier shift down along their respective axis and condition (5) guarantees the frontier becomes less convex in a symmetric way so that the derivative of the frontier at the 45° line is -1. Condition (4) establishes that the allocation that lies at the intersection of the production possibilities frontier and the 45° line is optimal for the administrator. Figure 27 depicts a change in the university's $m-k$ space as funding decreases. Since the ray along which the administrator's preferred research allocations lie does not change with changes in the level of governmental funding, the bargaining problem is neutral.

2.7 Nonneutrality with Feedback

We now consider the effect of relaxing condition (d) of Proposition 7 on the degree of negotiation leverage crowding-out.

Proposition 8: *If conditions (1), (2), (3), and (5) hold but not (4), a performance function that weights applied research more heavily than basic research, $\mu < 1/2$ will increase the degree of negotiation leverage crowding-out, and a performance function that weights basic research more heavily than applied research, $\mu > 1/2$ increases the degree of negotiation leverage crowding-in*

The effect of changing the relative weights of basic and applied research in the administrator's performance function is illustrated in figure 28 and figure 29 the administrator's performance function gives more weight to theorems, $\mu > 1/2$. A decline in funding leads to crowding-in as the university administrator's optimal allocation includes relatively more basic research. Figure 29 illustrates that for an administrator with $\mu < 1/2$, a similar decline in government funding leads to crowding-out. The magnitude of crowding-in/out increases with the size of the drop in governmental funding.

Relaxing condition (5) of Proposition 7 can similarly lead to either negotiation leverage crowding-in or out.

Proposition 9: *If conditions (1), (2), (3), and (4) hold but not (5), then a decrease in the level of government funded research will increase the degree of negotiation leverage crowding-out if $h(\cdot) > l(\cdot)$, or increase the degree of negotiation leverage crowding-in if $l(\cdot) > h(\cdot)$.*

First, assume the feedback from mousetraps to theorems is larger than the feedback from theorems to mousetraps, $l(\cdot) > h(\cdot)$. Figure 30 illustrates the effect of a decrease in government funding. For a university administrator who weights basic and applied research equally ($\mu = 1/2$), it is optimal to fund relatively more basic research when government funding declines. The intuition for this crowding-in of basic research is simple; the university administrator gets relatively more bang for every buck invested in applied research and, as funding declines and the asymmetric feedback effect becomes less pronounced, relatively more basic research is funded.

Figure 31 illustrates that a decrease in governmental funding has the opposite effect when the feedback from theorems to mousetraps is larger than the feedback from mousetraps to theorems, $h(\cdot) > l(\cdot)$. Since basic research gives relatively more bang for the buck, the administrator will fund relatively more applied research as governmental funding declines and the asymmetry becomes less pronounced.

2.8 Conclusion

Concerns that commercial sponsorship of university research will crowd-out basic research conducted for the public good, "theorems", are legitimate. Both models of the university research process used in this paper show that commercial sponsorship agreements can lead to the crowding-out of public good research. The likelihood of crowding-out depends on the framework used to model university research. In the first model, there is an implicit linear evolution from public good or basic research to applied or private good research (Bush, 1945). For the case of no dynamic relationships crowding-in only occurs under a set of highly restrictive assumptions. When there exists a linear decomposition of public good research and commercial or private good research, it is difficult for university administrators to form partnerships with profit-maximizing companies do not crowd-out public good research.

In the second model we use a dynamic framework to model university research that admits nonlinearities and recognizes the chaotic nature of the research and development processes. Unsurprisingly, crowding-in becomes more likely as we allow for feedback loops from discoveries in applied science to expand the opportunity set for public good research. Astute university administrators can form partnerships that magnify these feedback loops and increase public good research while working with the private sector. University administrator's should consider evaluating commercial sponsors based on the potential for significant feedback effects. For example, administrators can increase feedback effects by forming partnerships with companies that allow university researchers to access proprietary knowledge that is otherwise unavailable to the public. Identifying these complementarities will help university administrators create partnerships with the private sector that improve both public good and commercial research (Rausser, 1999).

We have structured the models presented in this paper under the simplifying assumption that university researchers simply follow direction from the administrator and do not respond to other incentives. In practice, the potential for royalties from commercialized research entices university researchers away from conducting public good research (Greenberg, 2007; Washburn, 2005). We expect these internal incentives to lead to more researchers working on applied science with commercial benefits than is socially optimal. Given these incentives underlying research pursuits, it is all the more important for administrators to acknowledge and exploit feedback loops by forming commercial research partnerships that enhance public good research.

The importance of feedback effects in scientific research has long been acknowledged. In 1871, Louis Pasteur wrote, "There does not exist a category of science to which one can give the name applied science. There are science and the applications of science, bound together as the fruit to the tree which bears it." By establishing research partnerships with the private sector that maximize feedback from applied to basic research, administrators at land-grant universities can take advantage of this unique feature of scientific research and promote the public good while managing their funding problems.

2.9 Appendix

The administrator's choices of m_t and k_t in the second model are determined by solving the following constrained optimization problem:

$$\begin{aligned} L = \sum_{t=0}^{T-1} [&\beta^t P_t(M_t, K_t) + \beta^T \psi(M_T, K_T) + \lambda_{1t}(\frac{1}{1+r})^t(\theta\mathcal{R}(M_{t-1}) + \bar{e}_t - m_t^{\beta_m} - k_t^{\beta_k}) \\ &+ \lambda_{2t}(K_t - \gamma K_{t-1} - k_t - l(M_{t-1})) + \lambda_{3t}(M_t - \alpha M_{t-1} - m_t - h(K_{t-1}))] \end{aligned}$$

The first-order conditions for the administrator's optimization problem in the second model are

$$
\begin{aligned}
\frac{dL}{dk_t} =& [\sum_{j=0}^{(T-1)-t}(\beta^{t+j}\frac{dP_{t+j}(M_{t+j},K_{t+j})}{dK_{t+j}}\frac{dK_{t+j}}{dk_t}) \\
&+ \sum_{j=0}^{(T-2)-t}(\beta^{t+j}\frac{dP_{t+j+1}(M_{t+j+1},K_{t+j+1})}{dM_{t+j+1}}\frac{dM_{t+j+1}}{dh_{t+j+1}}\frac{dh_{t+j+1}}{dK_{t+j}}\frac{dK_{t+j}}{dk_t})] \\
&+ [\beta^T(\frac{d\psi(M_T,K_T)}{dK_T}\frac{dK_T}{dk_t} + \frac{d\psi(M_T,K_T)}{dM_T}\frac{dM_T}{dh_T}\frac{dh_T}{dK_{T-1}}\frac{dK_{T-1}}{dk_t})] \\
&+ [\sum_{j=0}^{(T-2)-t}(\lambda_{1(t+j+2)}(\frac{1}{1+r})^{t+j+2}\theta\frac{d\mathcal{R}_{t+j+2}}{dM_{t+j+1}}\frac{dM_{t+j+1}}{dh_{t+j+1}}\frac{dh_{t+j+1}}{dK_{t+j}}\frac{dK_{t+j}}{dk_t}) - \lambda_{1t}(\frac{1}{1+r})^t\beta_k k_t^{\beta_k-1}] \\
&+ [[\sum_{j=0}^{T-1-t}(\lambda_{2(t+j)}\frac{dK_{t+j}}{dk_t})] - [\sum_{j=0}^{(T-2)-t}\lambda_{2(t+j)}\gamma\frac{dK_{t+j}}{dk_t}] - \lambda_{2t}] \\
&- [[\sum_{j=0}^{(T-3)-t}\lambda_{2(t+j+2)}\frac{dl_{t+j+2}}{dM_{t+j+1}}\frac{dM_{t+j+1}}{dh_{t+j+1}}\frac{dh_{t+j+1}}{dK_{t+j}}\frac{dK_{t+j}}{dk_t}] \\
&+ [\sum_{j=0}^{(T-2)-t}(\lambda_{3(t+j)}\frac{dM_{t+j+1}}{dh_{t+j+1}}\frac{dh_{t+j+1}}{dK_{t+j}}\frac{dK_{t+j}}{dk_t}) \\
&- \sum_{j=0}^{(T-3)-t}(\lambda_{3(t+j)}\alpha\frac{dM_{t+j+1}}{dh_{t+j+1}}\frac{dh_{t+j+1}}{dK_{t+j}}\frac{dK_{t+j}}{dk_t}) \\
&- \sum_{j=0}^{(T-2)-t}(\lambda_{3(t+j)}\frac{dh_{t+j+1}}{dK_{t+j}}\frac{dK_{t+j}}{dk_t})]] \leq 0
\end{aligned}
$$

$$
\begin{aligned}
\frac{dL}{dm_t} =&[\sum_{j=0}^{(T-1)-t}(\beta^{t+j}\frac{dP_{t+j}(M_{t+j},K_{t+j})}{dM_{t+j}}\frac{dM_{t+j}}{dm_t}) \\
&+\sum_{j=0}^{(T-2)-t}(\beta^{t+j}\frac{dP_{t+j+1}(M_{t+j+1},K_{t+j+1})}{dK_{t+j+1}}\frac{dK_{t+j+1}}{dl_{t+j+1}}\frac{dl_{t+j+1}}{dM_{t+j}}\frac{dM_{t+j}}{dm_t})] \\
&+[\beta^T(\frac{d\psi(M_T,K_T)}{dM_T}\frac{dM_T}{dm_t}+\frac{d\psi(M_T,K_T)}{dK_T}\frac{dK_T}{dl_T}\frac{dl_T}{dM_{T-1}}\frac{dM_{T-1}}{dm_t})] \\
&+[\sum_{j=0}^{(T-2)-t}(\lambda_{1(t+j+1)}(\frac{1}{1+r})^{t+j+1}\theta\frac{d\mathcal{R}_{t+j+1}}{dM_{t+j}}\frac{dM_{t+j}}{dm_t})-\lambda_{1t}(\frac{1}{1+r})^t\beta_m m_t^{\beta_m-1}] \\
&+[\sum_{j=0}^{(T-1)-1}(\lambda_{2(t+j)}\frac{dM_{t+j}}{dm_t})+\sum_{j=0}^{(T-2)-t}(\lambda_{2(t+j)}\alpha\frac{dM_{t+j}}{dm_t})-\lambda_{2t} \\
&-\sum_{j=0}^{(T-3)-t}(\lambda_{2(t+j+2)}\frac{dh_{t+j+2}}{dK_{t+j+1}}\frac{dK_{t+j+1}}{dl_{t+j+1}}\frac{dl_{t+j+1}}{dM_{t+j}}\frac{dM_{t+j}}{dm_t})] \\
&+[\sum_{j=0}^{(T-2)-t}(\lambda_{3(t+j)}\frac{dK_{t+j+1}}{dl_{t+j+1}}\frac{dl_{t+j+1}}{dM_{t+j}}\frac{dM_{t+j}}{dm_t}) \\
&-\sum_{j=0}^{(T-3)-t}(\lambda_{3(t+j)}\gamma\frac{dK_{t+j+1}}{dl_{t+j+1}}\frac{dl_{t+j+1}}{dM_{t+j}}\frac{dM_{t+j}}{dm_t}) \\
&-\sum_{j=0}^{(T-2)-t}(\lambda_{3(t+j)}\frac{dl_{j+t+1}}{dM_{t+j}}\frac{dM_{t+j}}{dm_t})]\leq 0
\end{aligned}
$$

where

$$\frac{dK_{t+j}}{dk_t}=\alpha^j+I_{j\geq 2}\sum_{i=2}^{j}\alpha^{j-i}\frac{dl(M_{t+i-1})}{dM_{t+i-1}}\frac{dM_{t+i-1}}{dK_{t+i-2}}\frac{dK_{t+i-2}}{dk_t} \tag{39}$$

$$\frac{dM_{t+j}}{dm_t}=\gamma^j+I_{j\geq 2}\sum_{i=2}^{j}\gamma^{j-i}\frac{dh(K_{t+i-1})}{dK_{t+i-1}}\frac{dK_{t+i-1}}{dM_{t+i-2}}\frac{dM_{t+i-2}}{dm_t} \tag{40}$$

In the second model the commercial company wants to maximize the sum of the profits from each period, $\pi_t(m_t,k_t)=M_t-f_t-\mathcal{R}(M_{t-1})$, and the scrap value of an agreement, $W(M_T)$, discounted by the weighted average cost of capital, r :

$$\Pi=\sum_{t=0}^{T-1}(\frac{1}{1+r})^t\pi(m_t,k_t)+(\frac{1}{1+r})^T W(M_T) \tag{41}$$

where f_t is the funding the company provides to the university to produce the company's share, (m_{c_t}, k_{c_t}), of the university's research allocation each period, $f_t = C(m_{c_t}, k_{c_t})$.

The company's first-order conditions are

$$\frac{d\Pi}{dm_t} = \sum_{j=0}^{(T-1)-t} (\frac{1}{1+r})^{t+j} \frac{dM_{t+j}}{dm_t} - (\frac{1}{1+r})^t \frac{dC_t(m_{c_t}, k_{c_t})}{dm_t}$$
$$- \sum_{j=0}^{(T-2)-t} (\frac{1}{1+r})^{t+j+1} (\frac{d\mathcal{R}(M_{t+j+1})}{dM_{t+j+1}} \frac{dM_{t+j+1}}{dm_t}) + (\frac{1}{1+r})^T \frac{dW(T)}{dm_t} = 0$$

$$\frac{d\Pi}{dk_t} = \sum_{j=0}^{(T-2)-t} (\frac{1}{1+r})^{t+j+1} \frac{dM_{t+j+1}}{dh_{j+t+1}} \frac{dh_{j+t+1}}{dK_{t+j}} \frac{dK_{t+j}}{dk_t} - (\frac{1}{1+r})^t \frac{dC_t(m_{c_t}, k_{c_t})}{dk_t}$$
$$- \sum_{j=0}^{(T-2)-t} (\frac{1}{1+r})^{t+j+1} \frac{d\mathcal{R}_{t+j+1}}{dM_{t+j+1}} \frac{dM_{t+j+1}}{dh_{j+t+1}} \frac{dh_{j+t+1}}{dK_{t+j}} \frac{dK_{t+j}}{dk_t} + \frac{dW(T)}{dk_t} = 0$$

3 The Effects of Congestion Pricing on Public Transportation Demand and Air Pollution[17]

3.1 Introduction

How does congestion pricing effect public transportation ridership and air pollution? The answer to this question has important implications for the design of transportation policy. Proponents of congestion pricing claim that drivers will not simply substitute off-peak driving for peak driving, rather, drivers will increase public transportation usage during peak hours, which will lower levels of air pollution (Geoghegan, 1995; Daniel and Bekka, 2000). Estimating the cross price elasticity of public transportation will improve forecasts of public transportation ridership and revenue increases following the implementation of congestion pricing. A precise estimate of the effect of decreased driving on air pollution will also improve welfare calculations for congestion pricing.

Congestion pricing on the San Francisco-Oakland Bay Bridge was implemented on July 1, 2010. Though the primary motivation for the policy was to raise funds to pay for bridge repair and maintenance, it was expected to address externalities associated with private automobile use: traffic and air pollution by decreasing the number of cars travelling on the bridge and also lowering travel times. Foreman (2012) showed that congestion pricing on the Bay Bridge did decrease traffic volume and travel times during peak hours, and some of that traffic shifted to off-peak hours. This paper extends her work by examining whether commuters shifted to public transportation and whether the shift is large enough to lower air pollution levels.

Though the cross-price public transportation demand has been studied extensively, those estimates are not necessarily useful for estimating the impact of congestion pricing. Early work on the cross-price elasticity of public transportation demand used increases in gasoline prices to capture increases to the cost of private transportation. This approach will yield biased estimates if the determinants of gasoline price (economic growth) also determine public transportation demand. Bland (1984) and Wang and Skinner (1984) used gasoline price increases to estimate the cross-price commuter bus elasticity ranges from 0.08 to 0.8. Doi and Allen (1986) use a similar approach to estimate the cross-price commuter rail elasticity of 0.11 cross price.

Similarly, estimating the effect of private transportation on air pollution levels is complicated by endogeneity. The estimates from early work on the topic vary widely, likely due to

[17] Joint research with Kate Foreman

endogeneity. But, recently, a number of papers have used quasi-natural experiments to estimate the effect of driving on pollution levels (Auffhammer and Kellogg, 2011; Davis, 2008). In this paper, we use the implementation of congestion pricing on the San Francisco Bay Bridge to estimate the effect of congestion pricing on public transportation and air pollution levels. The congestion pricing policy provides a natural experiment to estimate these effects. Ours is the first study to exploit the use of congestion pricing in the San Francisco Bay Area to estimate the cross price elasticity.

The central results of this paper are that the cross price elasticity of metro rail transportation (BART) is 0.2, and for commuter bus (AC-Transit) the elasticity is 0.15. These estimates are slightly below the estimates for the transportation economics literature previous literature. We also show that following the implementation of congestion pricing, local air pollution levels were flat during peak pricing hours, but rose somewhat during off-peak hours. This suggests some drivers shifter their commutes to off-peak hours to avoid paying the increased toll during peak hours.

The rest of the paper is organized as follows. In section 2 we review the empirical strategy and data. In section 3 we present the results from the difference-in-difference and regression discontinuity models. Concluding remarks are presented in section 4.

3.2 Empirical Strategy and Data

This paper uses two empirical strategies to identify the effects of congestion pricing public transportation and air pollution: difference-in-difference and regression discontinuity. These approaches exploit the quasi-natural experiment provided by congestion pricing to estimate the effect on public transportation and air pollution.

3.2.1 Empirical Strategy

We estimate the cross-price urban public transportation demand elasticity using a linear difference in difference model, where ridership depends on a treatment dummy, time and date fixed effects. Our baseline specification is:

$$riders_t = \beta_0 + \beta_1 after_t + \beta_j after_t * h + \rho_h + \phi_d + \omega_w + \varepsilon_t \tag{42}$$

where t indexes time, $riders_t$ is the hourly ridership level for a station pair, $after_t$ is a binary treatment indicator (1 if after toll change), $after * h$ is the treatment-hour interaction, ρ_h is the hour fixed effect (excluded category is 12:00 midnight), ϕ_d is the day fixed effect (excluded category is Monday), ω_w is the week fixed effect (excluded category is first week of year), and ε_t is an error term. This model is used for both public transportation (bus and commuter rail) and air pollution.

The parameter of interest, β_j, measures the effect of treatment by hour, and is our measure of hourly price elasticity. This parameter gives the effect on point-to-point commuter bus and commuter rail ridership by hour.

The key identifying assumption of our difference-in-difference estimators is that the trend in midnight public transportation ridership is parallel to the trend in peak and off-peak ridership. That is, absent the congestion pricing policy, the trend in midnight ridership (the control time period) has a similar trajectory to what we would have seen in the peak and off-peak ridership (treatment). The difference-in-difference estimator also assumes stable unit treatment value. We assume that peak and off-peak transportation usage would follow similar trends (relative to midnight) in the absence of treatment. Given these two assumptions, the coefficients in our baseline model can be estimated using least squares regression.

We use a similar model to estimate the effect of congestion pricing on air pollution. In this model, hourly pollution levels are a function of

$$pollution_t = \gamma_0 + \gamma_1 after_t + \gamma_j after_t * h + \gamma_k X_t + \rho_h + \phi_d + \omega_w + \varepsilon_t \tag{43}$$

In addition to the hour, day, and week fixed effect, we include weather control variables (X_t) that include temperature, precipitation, wind speed, and wind direction dummies. The variable of interest, γ_j, tells us the effect, by hour, of congestion pricing on pollution.

A potential problem with the differences-in-differences methodology is that an unobserved change in the control group (midnight) trend would bias results. We address this problem with a regression discontinuity framework. This allows us to estimate the effects of congestion pricing without relying on parallel trends in the treatment and control groups. The regression discontinuity estimator used for both public transportation and air pollution is

$$y_t = \alpha + \tau after_t + \beta X_t + \varepsilon_t \tag{44}$$

where y_t is the daily average of hourly ridership volume or air pollution level, $after_t$ is the binary treatment indicator (1 if after toll change), and X_t are the control variables. This approach uses the abrupt implementation of congestion pricing on July 1, 2010 to estimate short-run effect on public transportation air pollution effects, instead of the long run effects captured by the difference-in-difference estimator. The regression discontinuity estimator allows us to see whether the policy effect declines over time.

3.2.2 Data and Descriptive Statistics

3.2.2.1 Public Transportation Data Given the cost of private transportation only increased for traffic over the San-Francisco-Oakland Bay Bridge, the analysis in this paper is restricted to the alternate modes of public transportation over between San Francisco and Oakland. This paper uses data on public transportation from two sources: Bay Area Rapid Transit (BART) and Alameda-Contra Costa Transit District (AC-Transit). The BART dataset is a comprehensive record of all 62.97 million BART trips taken between July 2009 and June 2011. Of those trips, 88 percent (55.49 million) went through the Transbay Tube between the East Bay and San Francisco. The Transbay Tube follows roughly the same route as the Bay Bridge. This data set gives us the entire population of BART riders over the Bay Bridge. The unit of observation in this dataset is the number of riders, each hour, travelling between two stations.

Figure 32 below plots hourly BART ridership between the East Bay and San Francisco before and after the toll change.

The overall ridership levels are very similar before and after the implementation of congestion pricing; during both periods transbay ridership is highest during morning and evening commute times. The median transbay BART rider enters the system during commute hours. Figure 33 breaks down this data by plotting the hourly ridership into the San Francisco Central Business District.[18] The median commuter travels to San Francisco central business district from the East Bay in the morning and returns to the East Bay in the evening.

The AC-Transit dataset, unlike the BART dataset, provides a random sample of ridership levels by hour for commuter bus routes in the Bay Area. We use only AC-Transit buses routes that cross the Bay Bridge (the transbay lines). There are transbay AC-Transit buses routes originating in Berkeley, Castro Valley, El Cerrito, Hayward, Oakland, Richmond, and San Leandro. All of these routes terminate at the bus terminal in San Francisco. Figure 34 plots AC passengers per route.

Unlike BART, AC-Transit ridership tends to peak around mid-day. It appears that AC transit carries fewer traditional commuters. However, when we break down the ridership in and out of the San Francisco Central Business District (figure 35), there appears to be more riders leaving the city during the evening commute.

3.2.2.2 Air Pollution and Weather Data Air pollution data comes from the California Environmental Protection Agencys Toxic Air Contaminants (TAC) monitoring network.

[18] Defined as the area serviced by the Embarcadero, Montgomery Street, Powell Street, and Civic Center BART stations.

The Bay Area TAC network consists of 15 non-mobile stations that measure hourly air pollution levels. The location of these TAC monitors is shown in Figure X. The TAC monitors produce hourly measures of a number of pollutants. We use the measurements of the air pollutants most produced by vehicles: CO, NOX, and ozone.

We use readings from the four TAC monitors closest to the Bay Bridge to estimate the effect of congestion pricing on Bay Bridge pollution. One monitor is in San Francisco (Arkansas Street) and three monitors are in the East Bay in Oakland (West Oakland and International Blvd.) and Berkeley (6th St.) (see figure 36). We follow EPA data standards (Auffhammer and Kellogg, 2011) and drop all days for which there are at least 9 missing observations between 9:00 am and 9:00 pm.

We augment the air pollution data with weather data for the San Francisco Bay Area. It is essential to include wind, temperature, and precipitation data in air pollution regressions. We use weather data produced by the National Oceanic and Atmospheric Administrations (NOAA) National Weather service. We use weather data from the airport nearest the TAC monitoring station: San Francisco International Airport (SFO) for the San Francisco station and Oakland International Airport (OAK) for the East Bay monitoring stations. The weather data from these airports is extremely reliable and has few missing observations. We drop missing drop hours for the missing observations from the dataset.

We use hourly measurements of three air pollutants commonly associated with driving: CO, NOX, and ozone. Average hourly CO measurements from the San Francisco and East Bay TAC stations are displayed in figure 37. Similarly, the NOX measurements from the San Francisco and East Bay TAC stations, displayed in figure 38, also peak during morning rush hour.

Average hourly ozone measurements from the San Francisco and East Bay TAC stations are displayed in figure 39 Though ozone levels are a function of automobile traffic, they are determined in large part by temperature. Ozone levels peak when temperatures are the highest, during the midday.

3.3 Main Results

We first estimate the difference-in-difference models for public transportation ridership and air pollution levels. We then estimate the regression discontinuity models to estimate the short-run effects of congestion pricing under different assumptions.

3.3.1 Difference-in-Difference Models

In table 1, we present the results for a simple a simple regression of hourly public transportation ridership before and after congestion pricing. The first specification simply compares ridership before and after the policy change, and the second specification includes peak and off-peak dummy variables and interaction terms. The first column gives the estimate of the effect of congestion pricing on the number of BART riders per hour with no control variables. The congestion pricing coefficient is positive and statistically significant at the 1 percent level, suggesting a modest policy effect. The second column, which includes control variables for the peak and off-peak time periods, estimates a slightly larger, statistically significant, congestion pricing effect. The third column gives estimates of the effect of congestion pricing on AC Transit riders, with no control variables. The estimated effect, about 2 additional riders per hour, remains after adding control variables.

Next, we estimate the full difference-in-difference model for each hour, including day of week and week fixed effect. In each of these regressions, midnight ridership is used as the control period. Figure 41 plots the estimate treatment effect by hour for westbound and eastbound ridership for all transbay Bridge BART trains. As expected, the average treatment effect on westbound ridership peaks from 8:00 am 9:00 am. The 6.3 rider per hour estimated increase represents an 8.8 percent rise in overall ridership. The effect on eastbound ridership is peaks at a somewhat lower level, +4.1 riders per hour, but the statistically significant effect lasts until 10:00 pm, long after the end of congestion pricing.

Next, we estimate the average the full model by the entry (exit) station for all BART riders into (out of) San Franciscos Central Business District. Figure 3.1.2 plots the treatment effect by hour for each station. The largest increases were in stations along the Pittsburgh-Bay Point and Oakland lines, with many stations averaging increases from 15-30 additional riders per hour during peak congestion pricing hours. The stations with the smallest increases were along the Richmond line, where the implementation of congestion pricing did not have a statistically significant effect on the level of BART riders.

In the full difference-in-difference model, the changes in AC-Transit ridership levels were not as large as those seen in BART ridership. Figure 43 plots the estimated change in AC-Transit ridership after the implementation of congestion pricing. The effect on Westbound AC-Transit routes peaked at 7.9 riders per route during the peak evening hours, whereas the effect on eastbound traffic was not statistically significant. These results are consistent with AC-Transit ridership trends (section 2.2.1), where most riders use the system during midday.

When broken down by route, the largest increases in AC ridership were in routes originating in Hayward and Oakland, with 10-15 additional riders on some routes during peak morning hours and 7-12 addition riders during peak evening hours. The smallest effect was in the Berkeley routes.

In table 2, we present the estimates from a regression of hourly pollution levels before and after congestion pricing. For each pollutant we present two regressions: the first only includes a congestion pricing dummy and the second includes peak and off-peaky dummy and interaction terms. Though these regressions do not include important determinants of pollution levels (i.e., temperature and precipitation), these coefficients compare pollution levels before and after the policy change during peak and off-peak hours. Carbon monoxide levels fell slightly following the implementation of congestion pricing (columns 1 and 2), while ozone levels rose (columns 5 and 6). The effect on nitrous oxide level varied by specification (columns 3 and 4), with a slight increase estimated in the simple specification and a small decrease estimated when controls are added.

In the following charts, we present estimates for the full difference-in-difference pollution models. These models estimate the effect of congestion pricing on hourly pollution levels, including time (day and week fixed effects) and weather fixed effects. The weather control variables include temperature, wind speed, wind direction, and precipitation. Similar to the public transportation models, midnight pollution levels are used as the control. The overall and station level treatment effect estimates for carbon monoxide are presented in figure 46. Carbon monoxide levels rise during the morning and evening off-peak hours. This result is consistent with some drivers moving their commutes to off-peak hours. This result is driven almost entirely by pollution levels in the East Bay, while the change San Francisco pollution levels are not statistically significant.

The results for the nitrous oxide regressions, plotted in figure 47, are similar to the carbon monoxide results: large increases during off-peak hours and small, statistically insignificant decreases during peak hours.

Unlike the other air pollutants, the effect of congestion pricing on ozone levels was negative and statistically significant during morning peak and off-peak hours (figure 48). The effect on evening ozone levels was slightly positive and statistically significant during evening peak hours.

3.3.2 Regression Discontinuity Models

In this section we use a regression discontinuity approach to estimate the short-run effects of congestion pricing on Bay Area public transportation and air pollution. This allows us to relax the parallel trends and stable unit treatment value assumptions used in the difference in difference models in section 3.3.1. In the regression discontinuity graphs below, we plot estimate daily averages of the outcome variable of interest along with the fitted values from a linear regression for the month before and after the policy change. For comparison, we also include the same plot for the year after congestion pricing was implemented. BART ridership tends to fall from June to July due to the July 4th holiday. Figure 49 shows the fall in ridership following the implementation of congestion pricing in 2010 was much smaller than the drop in the level of riders the following year. One explanation for this declining effect of congestion pricing could be that the initial response of drivers to congestion pricing was larger than the long-term effect.

AC-Transit ridership also tends to fall during in early July. But, in 2010, average ridership was unchanged during the month of July, suggesting a short-run impact on AC-Transit ridership that diminished over the following year.

The regression discontinuity results show an immediate impact on July ridership, which, accounting for seasonal trends, is positive. This effect appears to diminish over time. If the effect of congestion pricing does decline over time, there are important implications for policymakers.

3.4 Conclusion

Following the implementation of congestion pricing on the San Francisco-Oakland Bay Bridge, we show that public transportation usage increased and air pollution levels fell. We use panel data on hourly BART and AC-Transit riders from 2009 through 2011, we examine the effect of congestion pricing on public transportation usage and air pollution.

Our preferred estimation method is a difference-in-difference model, which uses the midnight hour trend to control for unobserved variables. In this model, we estimate that BART ridership rose 4-8 percent during peak hours, while AC-Transit ridership rose 11-15 percent. Our estimated cross-price elasticities for public transportation, 0.2 for commuter rail and 0.15 for commuter bus, are slightly below estimates from the transportation economics literature.

Using a similar approach for Bay Area air pollution levels during the same time period, we find air pollution levels are flat during peak hours, but rise during off peak hours. This suggests some drivers are adjusting their commute times to avoid driving during peak hours, rather than using public transportation. As a robustness check, we use regression discontinuity models to estimate the short-run effects of congestion pricing. These models show similar, but slightly smaller effects of congestion pricing.

References

Aastveit, Knut Are, Gisle James Natvik, and Sergio Sola (2013), "Economic Uncertainty and the Effectiveness of Monetary Policy." Technical report, Norges Bank.

Alston, J. M., C. Chan-Kang, M. C. Marra, P. G. Pardey, and J. J. Wyatt (2000), *A Meta-analysis of Rates of Return to Agricultural Research and Development*. International Food Policy Research Institute.

Andreou, Elena, Eric Ghysels, and Andros Kourtellos (2013), "Should Macroeconomic Forecasters use Daily Financial Data and How?" *Journal of Business & Economic Statistics*, 31, 240–251.

Apergis, Nicholas and Stephen M Miller (2009), "Do Structural Oil-Market Shocks Affect Stock Prices?" *Energy Economics*, 31, 569–575.

Aruoba, S Borağan, Francis X Diebold, and Chiara Scotti (2009), "Real-time measurement of business conditions." *Journal of Business & Economic Statistics*, 27.

Auffhammer, Maximilian and Ryan Kellogg (2011), "Clearing the air? the effects of gasoline content regulation on air quality." *The American Economic Review*, 2687–2722.

Bagliano, Fabio C and Carlo A Favero (1999), "Information from Financial Markets and VAR Measures of Monetary Policy." *European Economic Review*, 43, 825–837.

Bai, Y, DQ Zhou, P Zhou, and LB Zhang (2012), "Optimal Path for China's Strategic Petroleum Reserve: A Dynamic Programming Analysis." *Energy Economics*, 34, 1058–1063.

Baker, Scott R, Nicholas Bloom, and Steven J Davis (2013), "Measuring Economic Policy Uncertainty." Unpublished Working Paper.

Balas, Egon (1981), *The Strategic Petroleum Reserve: How Large Should it be?* in Energy Policy Planning, Bayraktar et al. (Eds.), Plenum Press, New York.

Barro, Robert J. and Charles J. Redlick (2011), "Macroeconomic Effects From Government Purchases and Taxes." *The Quarterly Journal of Economics*, 126, 51–102.

Baumeister, Christiane and Gert Peersman (2013a), "The Role of Time-Varying Price Elasticities in Accounting for Volatility Changes in the Crude Oil Market." *Journal of Applied Econometrics*, 28, 1087–1109.

Baumeister, Christiane and Gert Peersman (2013b), "Time-Varying Effects of Oil Supply Shocks on the US Economy." *American Economic Journal: Macroeconomics*, 5, 1–28.

Beattie, B. R. (1991), "Some almost ideal remedies for healing land grant universities." *American Journal of Agricultural Economics*, 73, 1307–21.

Bekaert, Geert, Marie Hoerova, and Marco Lo Duca (2013), "Risk, Uncertainty and Monetary Policy." *Journal of Monetary Economics*, 60, 771–788.

Berge, Travis J and Òscar Jordà (2011), "Evaluating the Classification of Economic Activity into Recessions and Expansions." *American Economic Journal: Macroeconomics*, 3, 246–277.

Bernanke, Ben S and Alan S Blinder (1992), "The Federal Funds Rate and the Channels of Monetary Transmission." *American Economic Review*, 82, 901–21.

Bernanke, Ben S. and Ilian Mihov (1998), "Measuring Monetary Policy." *The Quarterly Journal of Economics*, 113, 869–902.

Bland, BH (1984), "Effect of fuel price on final use and travel patterns." *LR1114. Transport and Road Research Laboratory, Crowthorne.*

Bloom, Nicholas (2009), "The Impact of Uncertainty Shocks." *Econometrica*, 77, 623–685.

Bohi, Douglas and Michael Toman (1996), *The Economics of Energy Security*. Norwell, MA: Kluwer Academic Publishers.

Bonnen, J. T. (1986), "A century of science in agriculture: Lessons for science policy." *American Journal of Agricultural Economics*, 68, 1065–80.

Bush, V. (1945), *Science: The Endless Frontier*. Washington DC: Office of Scientific Research and Development.

Cabinet Task Force on Oil Import Control (1970), *The Oil Import Question*. Washington D.C.: U.S. Government Printing Office.

Chao, Hung-Po and Alan S Manne (1983), "Oil Stockpiles and Import Reductions: A Dynamic Programming Approach." *Operations Research*, 31, 632–651.

Christiano, Lawrence, Martin Eichenbaum, and Charles Evans (1999), "Monetary Policy Shocks: What Have we Learned and to What End?" *Handbook of Macroeconomics*, 1, 65–148.

Christiano, Lawrence J, Martin Eichenbaum, and Charles Evans (1996), "The Effects of Monetary Policy Shocks: Evidence from the Flow of Funds." *The Review of Economics and Statistics*, 78, 16–34.

Cochrane, Jone and Monika Piazzesi (2002), "The Fed and Interest Rates - A High-Frequency Identification." *American Economic Review*, 92, 90–95.

Congressional Research Service (2006), *Energy Policy Act of 2005: Summary and Analysis of Enacted Provisions*. RL33302, Library of Congress.

Considine, Timothy J (2006), "Is the Strategic Petroleum Reserve Our Ace in the Hole?" *The Energy Journal*, 91–112.

Daniel, Joseph I and Khalid Bekka (2000), "The environmental impact of highway congestion pricing." *Journal of Urban Economics*, 47, 180–215.

Davis, Lucas W (2008), "The effect of driving restrictions on air quality in mexico city." *Journal of Political Economy*, 116, 38–81.

Department of Agriculture (2002), "2002 census of agriculture." Technical report, AC-02-A-51 Volume 1, Part 51. U.S. Department of Agriculture.

Dixit, Avinash and Robert Pindyk (1994), *Investment under uncertainty*. Princeton University Press.

Doi, Masayuki and W Bruce Allen (1986), "A time series analysis of monthly ridership for an urban rail rapid transit line." *Transportation*, 13, 257–269.

Evanson, R. E., P. E. Waggoner, and V. W. Ruttan (1979), "Economic benefits from research: An example from agriculture." *Science*, 205, 1101–07.

Foreman, Kate (2012), "Crossing the bridge: The effects of time-varying tolls on curbing congestion." *Mimeo, UC Berkeley*.

Gardner, B. L. (2002), *U.S. Agriculture in the Twentieth Century: How it Flourished and What it Cost*. Harvard University Press.

Geoghegan, Jacqueline (1995), "The road not taken: Environmental congestion pricing on the san francisco-oakland bay bridge." *Mimeo, UC Berkeley*.

Government Accountability Office (2006), *Available Oil Can Provide Significant Benefits, but Many Factors Should Influence Future Decesions about Fill, Use, and Expansion*. Publication No. GAO-06-872.

Greenberg, D. (2007), *Science for Sale: The Perils, Rewards, and Delusions of Campus*

Capitalism. Chicago, IL: University Of Chicago Press.

Hamilton, James D (1983), "Oil and the Macroeconomy since World War II." *The Journal of Political Economy*, 228–248.

Hamilton, James D and Ana Maria Herrera (2004), "Comment: Oil Shocks and Aggregate Macroeconomic Behavior: The Role of Monetary Policy." *Journal of Money, Credit and Banking*, 265–286.

Huffman, W. E. and R. E. Evenson (1993), *Science for Agriculture*. Ames: Iowa State University Press.

International Energy Agency (2014), *Oil Information, 2014 Edition*. Paris: International Energy Agency.

Jo, Soojin (2014), "The Effects of Oil Price Uncertainty on Global Real Economic Activity." *Journal of Money, Credit and Banking*, 46, 1113–1135.

Just, R. E. and G. C. Rausser (1993), "The governance structure of agricultural science and agricultural economics: A call to arms." *American Journal of Agricultural Economics*, 75, 69–83.

Kerr, N. S. (1987), *The legacy: a centennial history of the State Agricultural Experiment Stations, 1887–1987*. Columbia: Missouri Agricultural Experiment Station, University of Missouri.

Kilian, Lutz (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, Lutz and Daniel P. Murphy (2012), "Why Agnostic Sign Restrictions Are Not Enough: Understanding the Dynamics of Oil Market VAR Models." *Journal of the European Economic Association*, 10, 1166–1188.

Kilian, Lutz and Daniel P. Murphy (2014), "The Role of Inventories and Speculative Trading in the Global Market for Crude Oil." *Journal of Applied Econometrics*, 29, 454–478.

Leeper, Eric M, Christopher A Sims, Tao Zha, Robert E Hall, and Ben S Bernanke (1996), "What Does Monetary Policy Do?" *Brookings papers on economic activity*, 1–78.

Leiby, Paul N and David Bowman (2000), "The Value of Expanding the US Strategic Petroleum Reserve." *ORNL/TM-2000/179, Oak Ridge National Laboratory, Oak Ridge, Tennessee, November*.

Lippi, Francesco and Andrea Nobili (2012), "Oil and the Macroeconomy: A Quantitative Structural Analysis." *Journal of the European Economic Association*, 10, 1059–1083.

Mertens, Karel and Morten O Ravn (2013), "The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States." *The American Economic Review*, 103, 1212–1247.

Monteforte, Libero and Gianluca Moretti (2013), "Real-Time Forecasts of Inflation: The Role of Financial Variables." *Journal of Forecasting*, 32, 51–61.

National Petroleum Council (1973), *Emergency Preparedness for Interruption of Petroleum Imports into the United States*. Department of Interior. Washington D.C.: U.S. Government Printing Office.

Nichols, Albert L and Richard J Zeckhauser (1977), "Stockpiling Strategies and Cartel prices." *The Bell Journal of Economics*, 66–96.

Nordhaus, William D (1974), "The 1974 Report of the President's Council of Economic Advisers: Energy in the Economic Report." *The American Economic Review*, 558–565.

Olea, M, J Stock, and M Watson (2012), "Inference in Structural VARs with External

Instruments.” Technical report.

Peersman, Gert and Ine Van Robays (2009), “Oil and the Euro Area Economy.” *Economic Policy*, 24, 603–651.

Peersman, Gert and Ine Van Robays (2012), “Cross-Country Differences in the Effects of Oil Shocks.” *Energy Economics*, 34, 1532–1547.

Ramey, Valerie A. and Matthew D. Shapiro (1999), “Costly Capital Reallocation and the Effects of Government Spending.” Working Paper 6283, National Bureau of Economic Research.

Rasmussen, W. D. (1989), *Taking the university to the people: seventy-five years of Cooperative Extension*. Ames: Iowa State University Press.

Rausser, G. C. (1999), “Private/public research: Knowledge assets and future scenarios.” *American Journal of Agricultural Economics*, 81, 1011–1027.

Rausser, G. C. and P. Zusman (1992), “Public policy and consitutional prescription.” *American Journal of Agricultural Economics*, 74, 247–57.

Romer, Christina D. and David H. Romer (1989), “Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz.” Working Paper 2966, National Bureau of Economic Research.

Romer, Christina D. and David H. Romer (2004), “A New Measure of Monetary Shocks: Derivation and Implications.” *American Economic Review*, 94, 1055–1084.

Romer, Christina D. and David H. Romer (2010), “The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks.” *American Economic Review*, 100, 763–801.

Rudebusch, Glenn D. (1998), “Do Measures of Monetary Policy in a VAR Make Sense?” *International Economic Review*, 39, 907–931.

Ruttan, V. W. (1982), *Agricultural Research Policy*. Minneapolis: University of Minnesota Press.

Samouilidis, J. and S. Berahas (1982), “A Methodological Approach to Strategic Petroleum Reserves.” *Omega*, 10, 565–574.

Sims, Christopher A (1986), “Are Forecasting Models Usable for Policy Analysis?” *Federal Reserve Bank of Minneapolis Quarterly Review*, 10, 2–16.

Sims, Christopher A. (1992), “Interpreting the Macroeconomic Time Series Facts: The Effects of Monetary Policy.” Cowles Foundation Discussion Papers 1011, Cowles Foundation for Research in Economics, Yale University.

Stock, James H and Mark W Watson (2008), “NBER Summer Institute Minicourse 2008: What’s New in Econometrics - Time Series, Lecture 7: Structural VARs.”

Stock, James H and Mark W Watson (2012), “Disentangling the Channels of the 2007-2009 Recession.” Technical report, National Bureau of Economic Research.

Teisberg, Thomas J (1981), “A Dynamic Programming Model of the US Strategic Petroleum Reserve.” *The Bell Journal of Economics*, 526–546.

Tolley, George S and John D Wilman (1977), “The Foreign Dependence Question.” *The Journal of Political Economy*, 323–347.

U.S. Congress (1975), “S.622: The Energy Policy and Conservation Act of 1975.” Pub. L. No. 94-163. U.S. Government Printing Office.

U.S. Senate (2003), “U.S. Strategic Petroleum Reserve: Recent policy has increased costs to consumers but not overal U.S. energy security (108-18).” Text from: *Committee Reports.*

U.S. Government Printing Office.
Van Robays, Ine (2012), "Macroeconomic Uncertainty and the Impact of Oil Shocks." Technical report, CESifo Working Paper: Monetary Policy and International Finance.
Verleger, Phillip (2003), "Measuring the economic impact of an oil release from the Strategic Petroleum Reserve to compensate for the loss of Venezuelan oil production." Included in U.S. Strategic Petroleum Reserve (108-18), Appendix 4.
Wang, George HK and David Skinner (1984), "The impact of fare and gasoline price changes on monthly transit ridership: empirical evidence from seven us transit authorities." *Transportation Research Part B: Methodological*, 18, 29–41.
Washburn, J. (2005), *University, Inc: The Corporate Corruption of Higher Education*. Cambridge, MA: Basic Books.
Weimer, David Leo (1982), *Strategic Petroleum Reserve: planning, implementation, and analysis*. Greenwood Press, Westport, CT.
Wright, Brian D and Jeffrey C Williams (1982), "The Roles of Public and Private Storage in Managing Oil Import Disruptions." *The Bell Journal of Economics*, 341–353.
Wu, Gang, Ying Fan, Lan-Cui Liu, and Yi-Ming Wei (2008), "An Empirical Analysis of the Dynamic Programming Model of Stockpile Acquisition Strategies for China's Strategic Petroleum Reserve." *Energy Policy*, 36, 1470 – 1478.
Zhang, Xiao-Bing, Ying Fan, and Yi-Ming Wei (2009), "A Model Based on Stochastic Dynamic programming for Determining China's Optimal Strategic Petroleum Reserve Policy." *Energy Policy*, 37, 4397–4406.

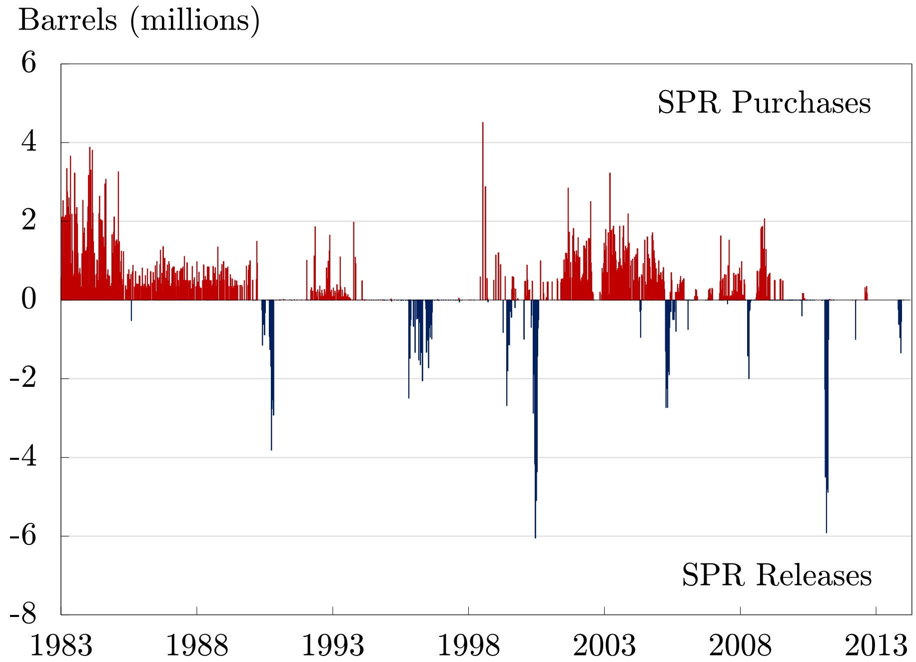

Figure 1: SPR Purchases and Releases (1983–2014)

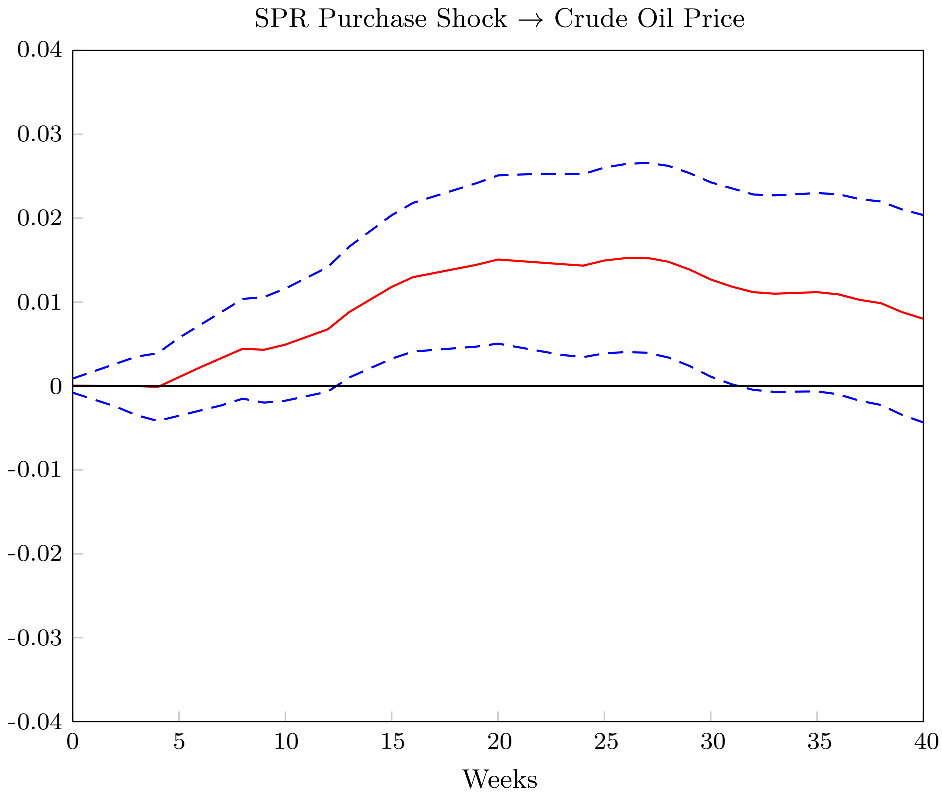

Figure 2: SPR Purchase Impulse Response Function (Benchmark Model)

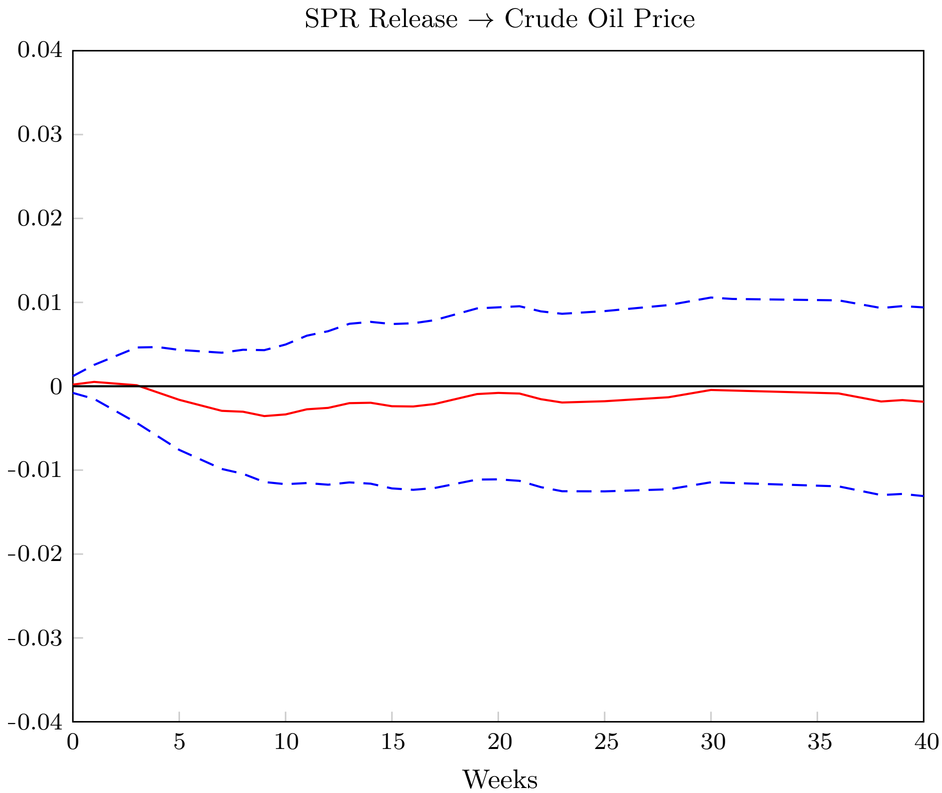

Figure 3: SPR Release Impulse Response Function (Benchmark Model)

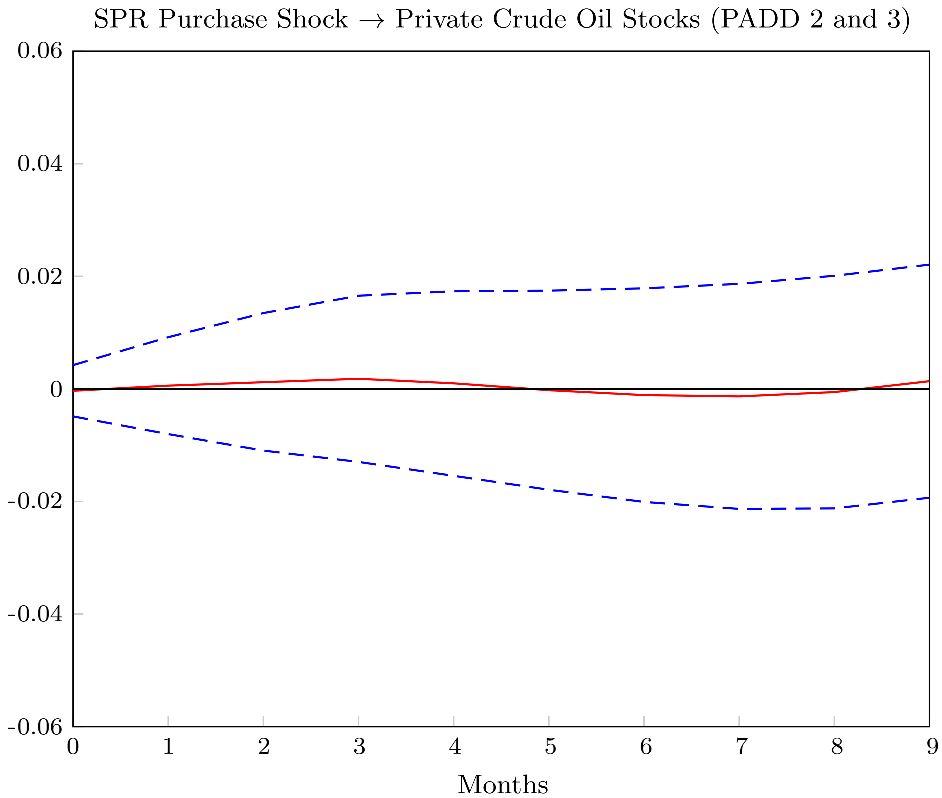

Figure 4: SPR Purchase Impulse Response Function (Private Crude Stocks)

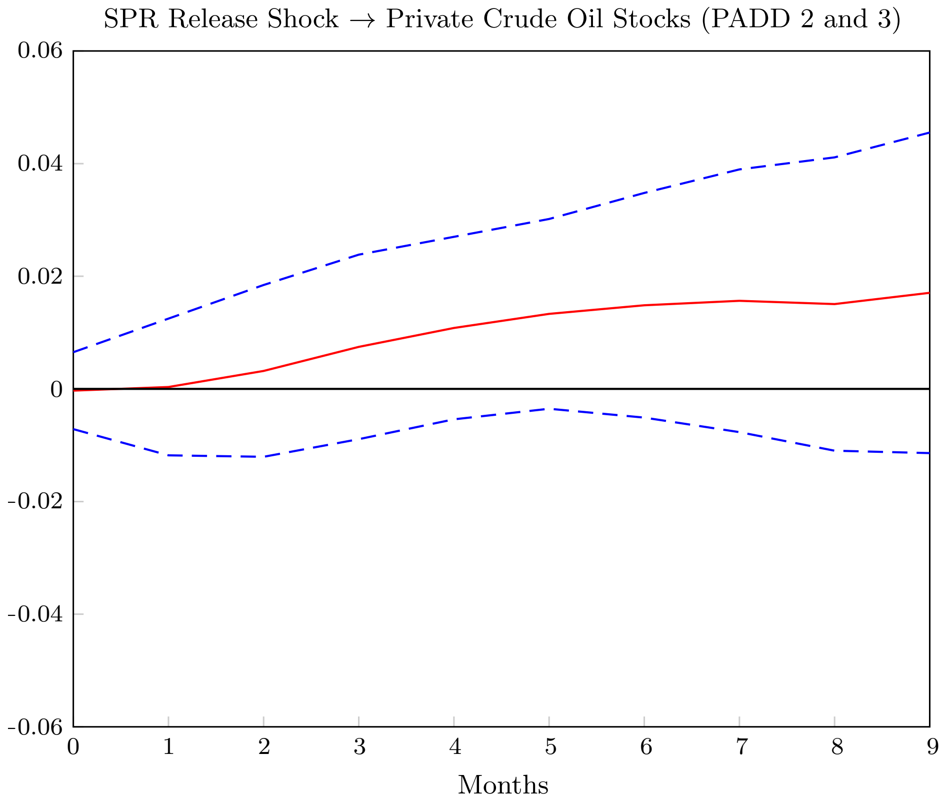

Figure 5: SPR Release Impulse Response Function (Private Crude Stocks)

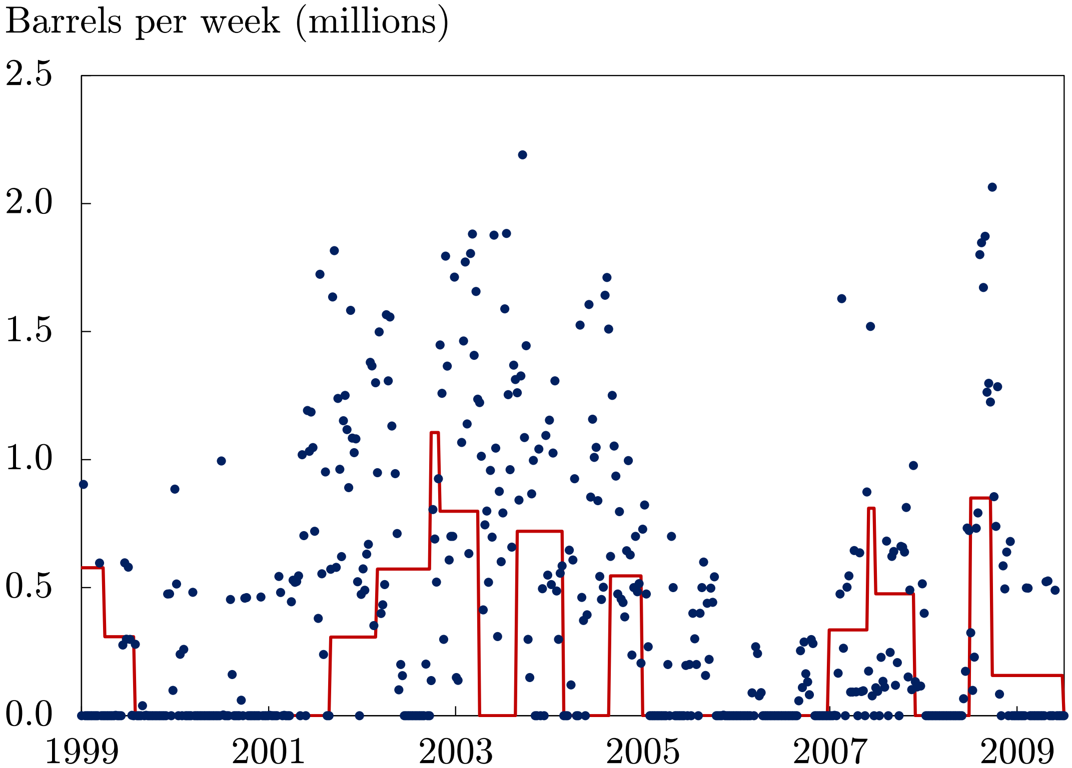

Figure 6: SPR Purchase Schedule (red line) and Actual Purchases (blue dots)

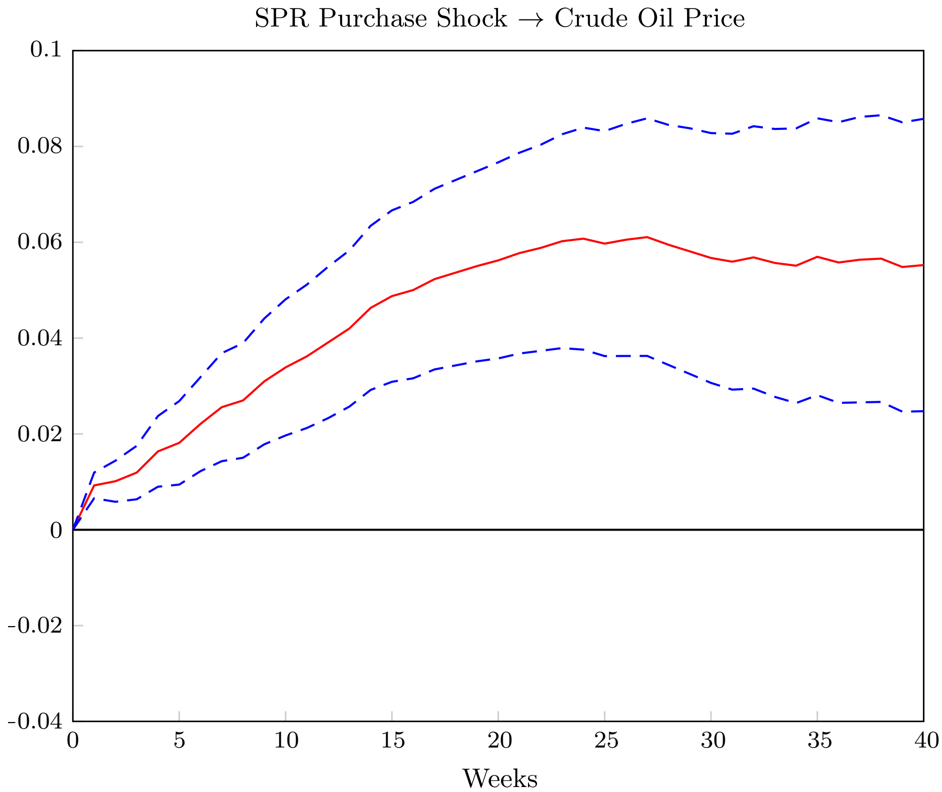

Figure 7: SPR Purchase Impulse Response Function (Identified with External Instrument)

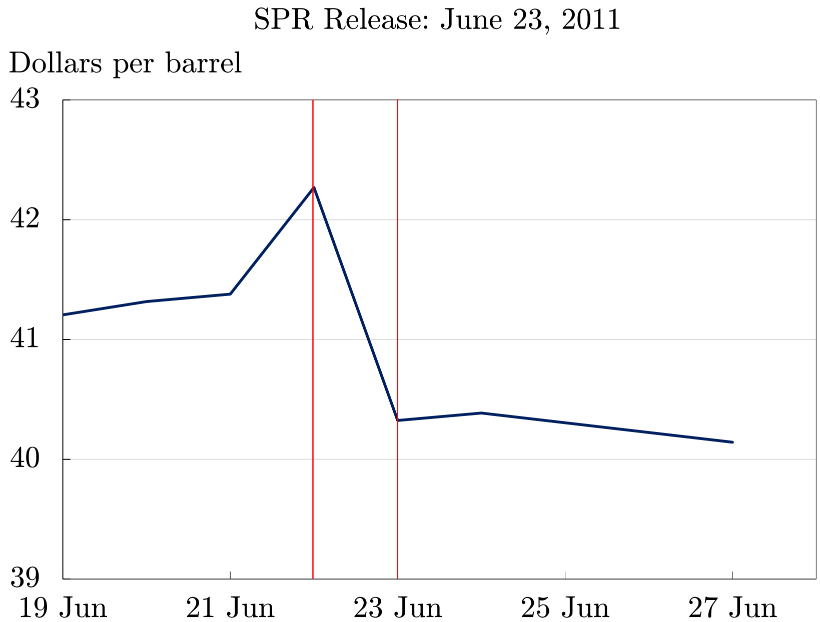

(a) Unanticipated SPR Release

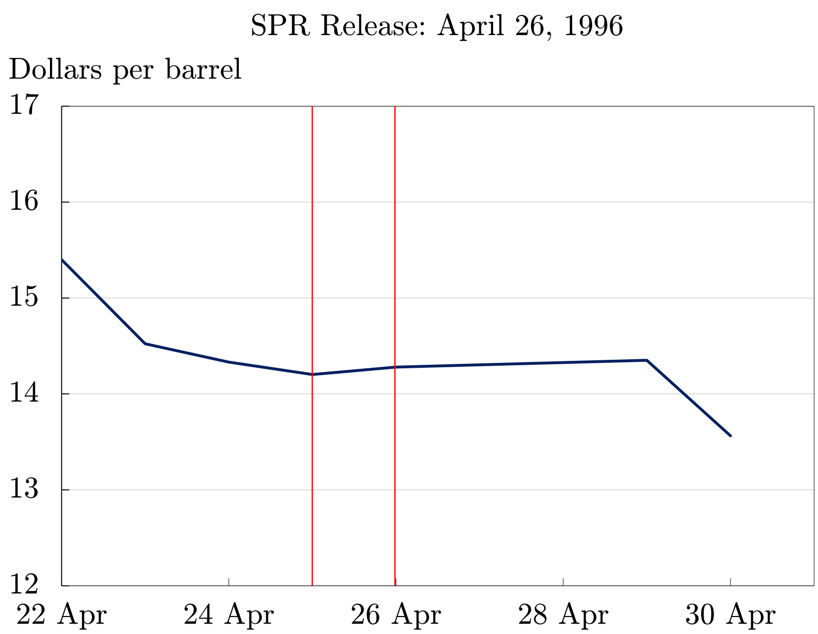

(b) Anticipated SPR Release

Figure 8: Crude Oil Futures Prices Following Release Announcements

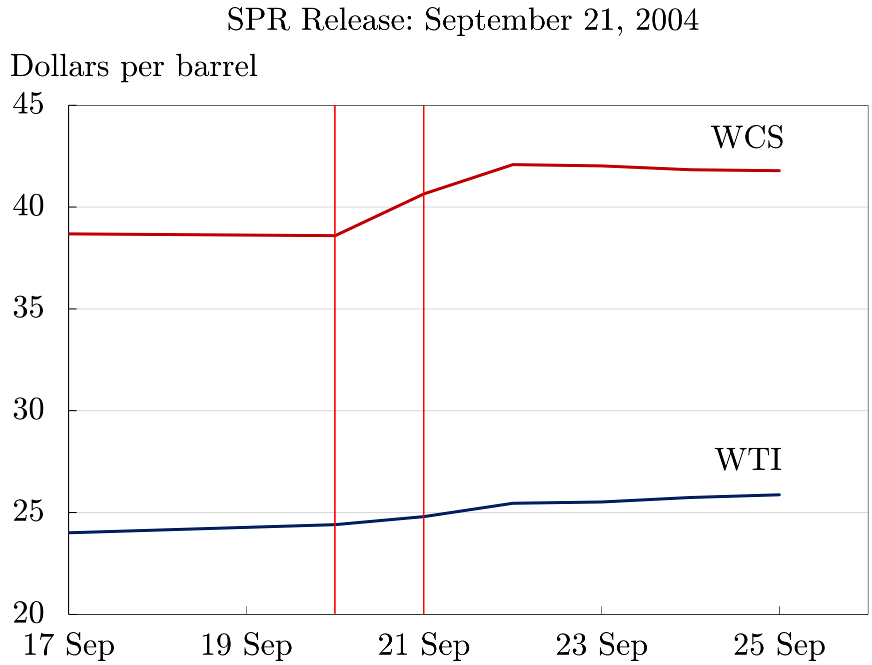

Figure 9: U.S.-Canada Crude Oil Futures Prices Following Release Announcement

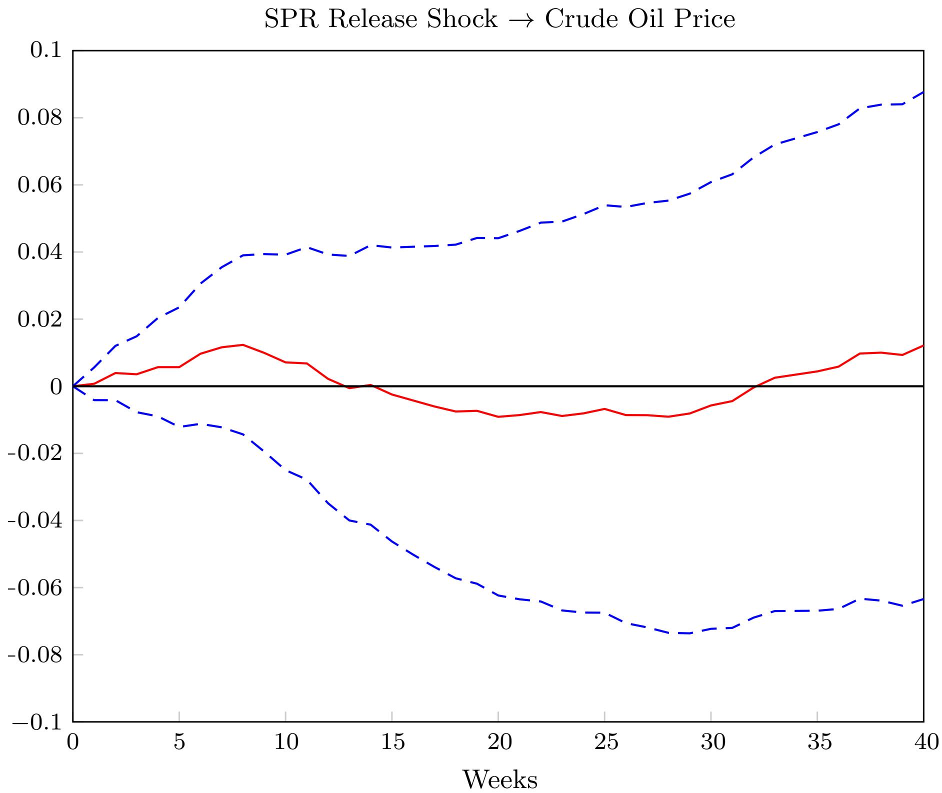

Figure 10: SPR Release Impulse Response Function (identified using WTI futures)

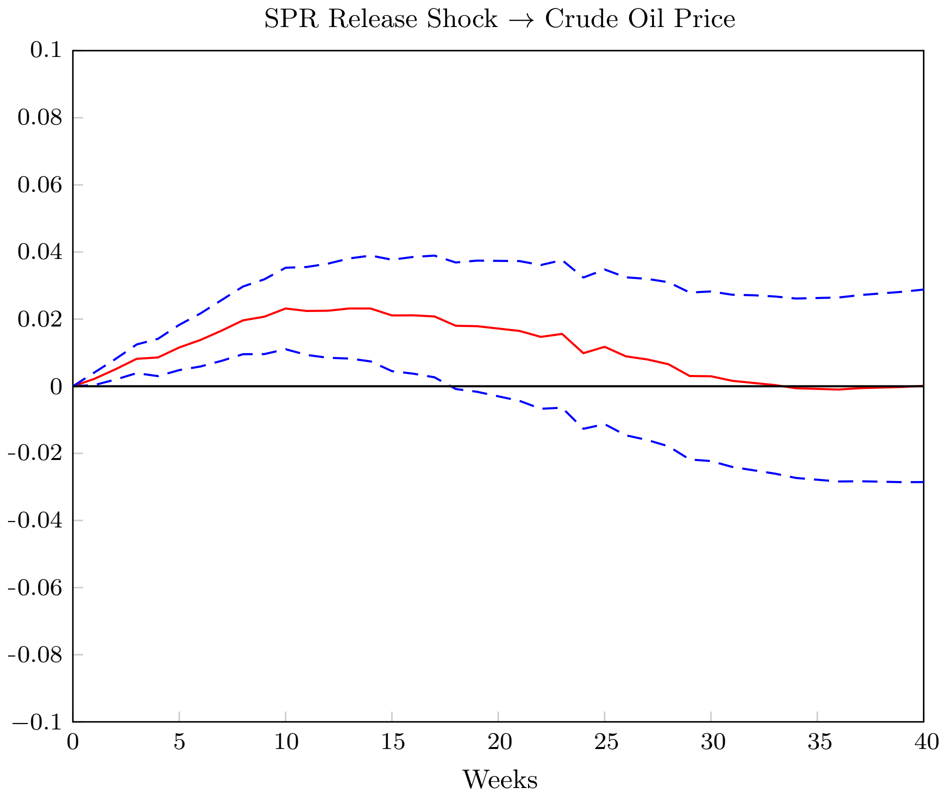

Figure 11: SPR Release Impulse Response Function (identified using the WTI-WCS futures spread)

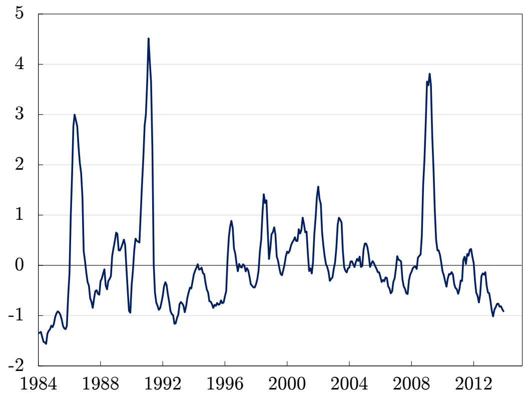

(a) Oil Market Volatility Index

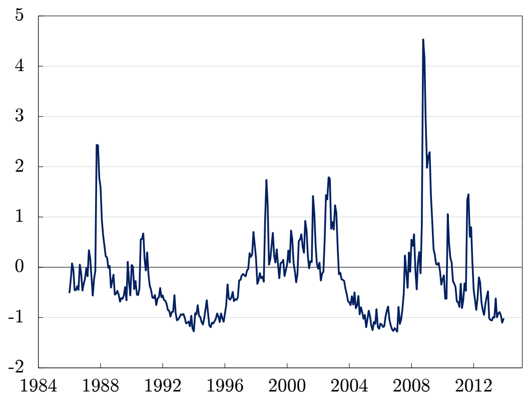

(b) Stock Market Volatility Index

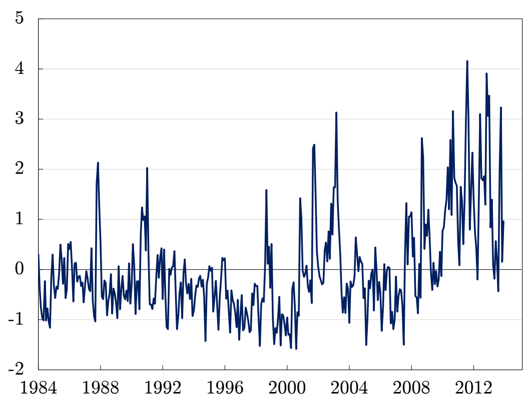

(c) Policy Uncertainty Index

Figure 12: Uncertainty Indices

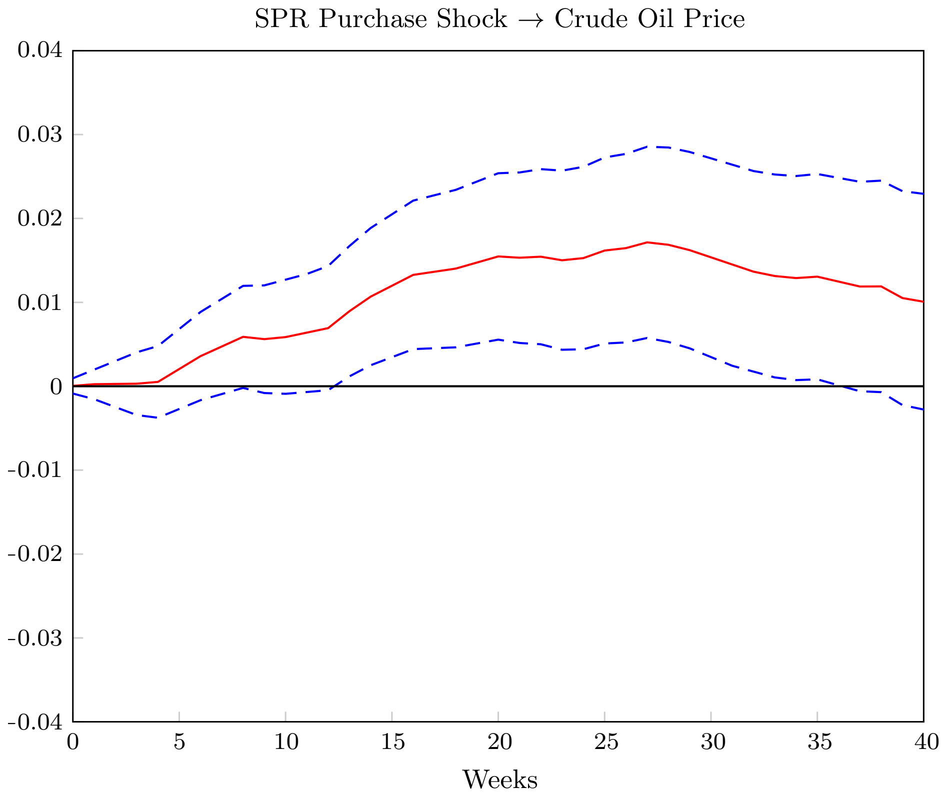

Figure 13: SPR Purchase Impulse Response Function (Mean Oil Market Volatility)

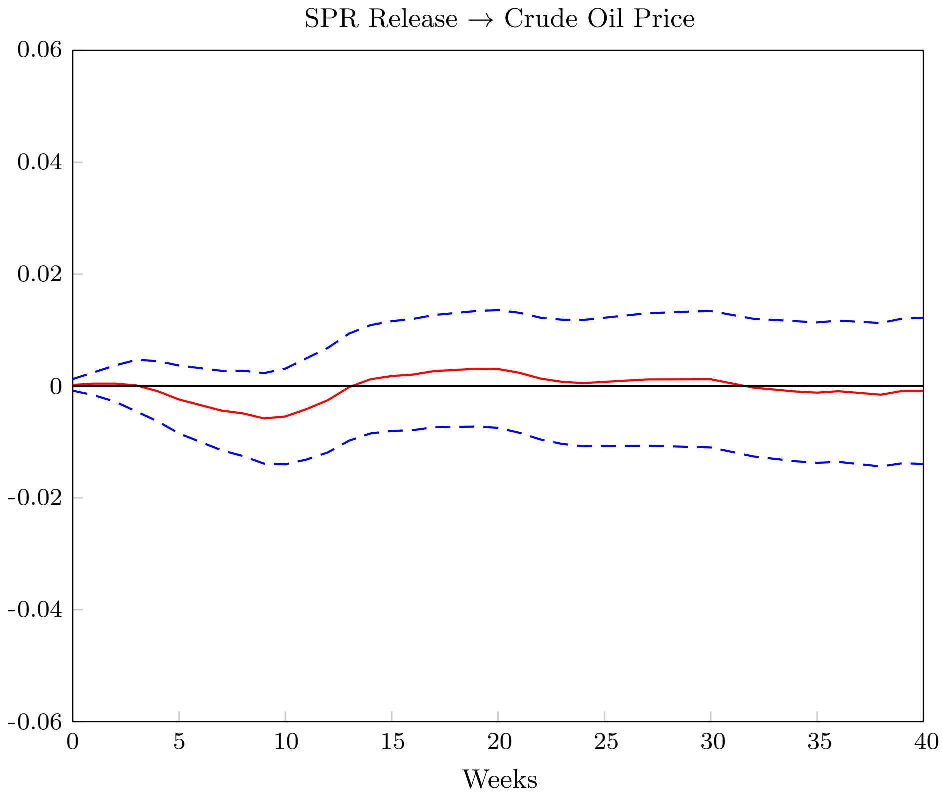

Figure 14: SPR Release Impulse Response Functions (Mean Oil Market Volatility)

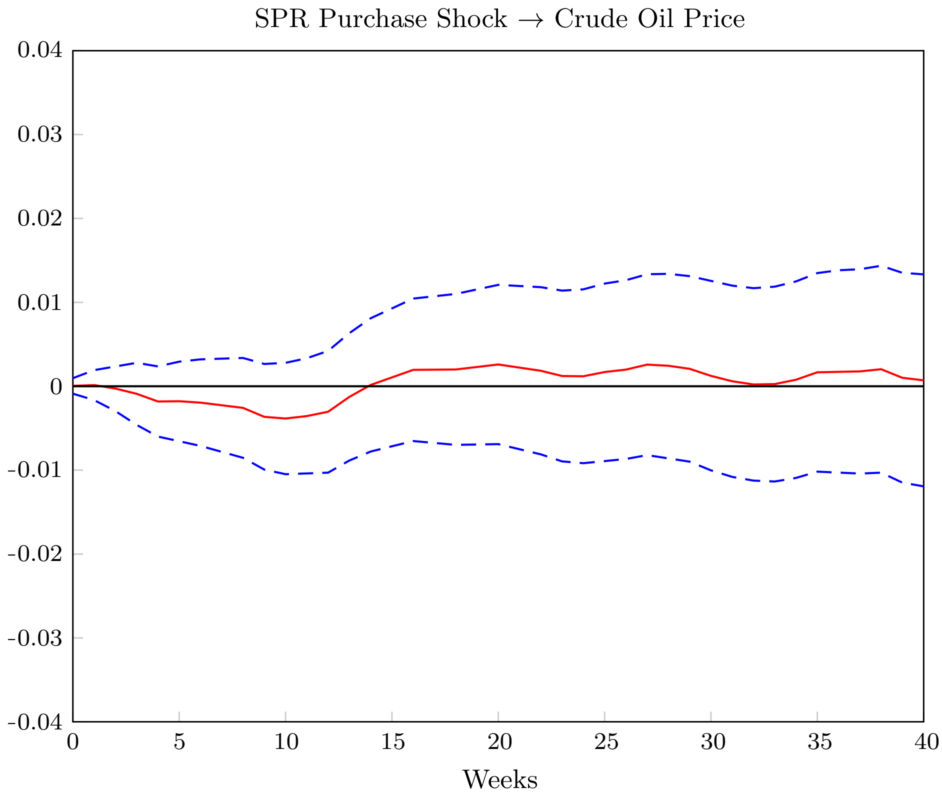

Figure 15: SPR Purchase Impulse Response Function (Low Oil Market Volatility)

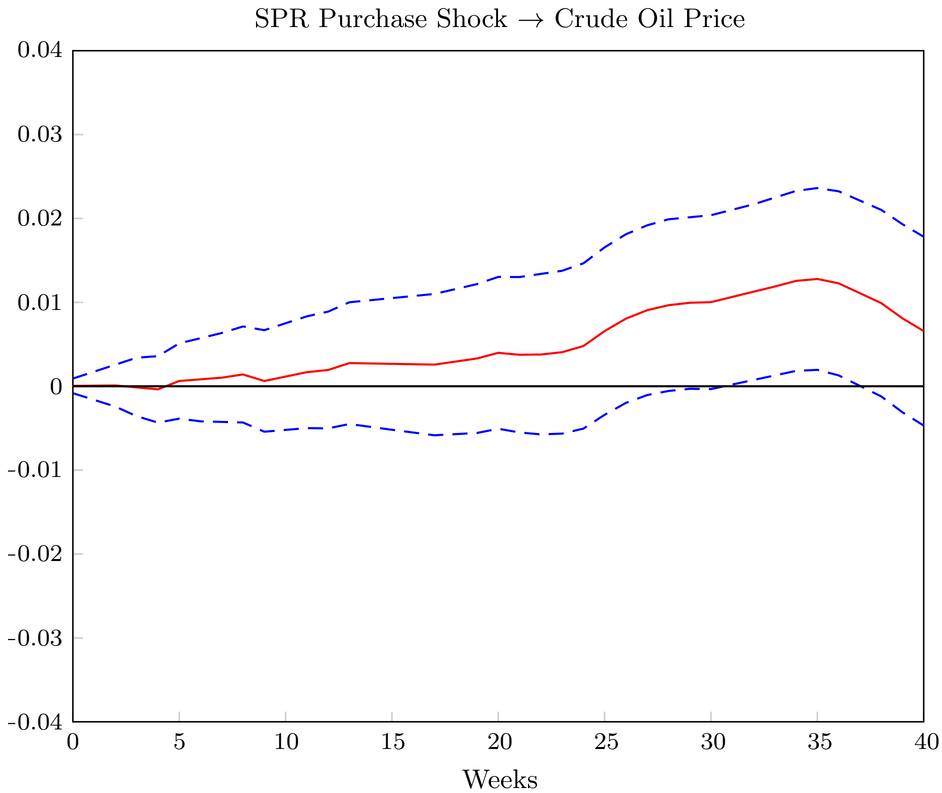

Figure 16: SPR Purchase Impulse Response Function (Low Stock Market Volatility)

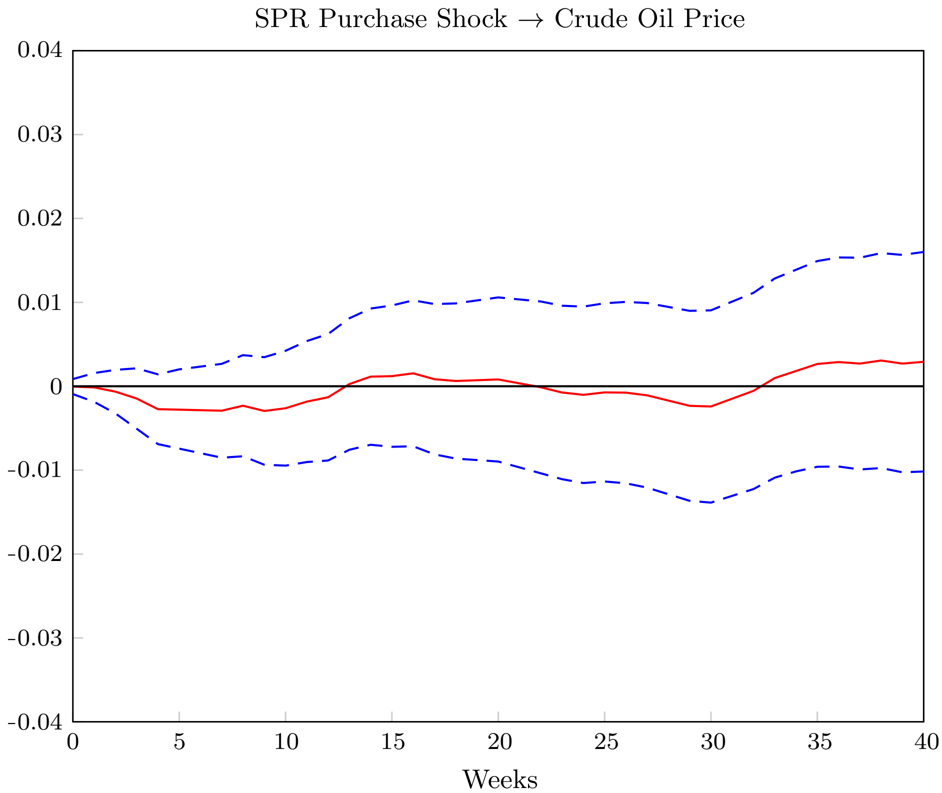

Figure 17: SPR Purchase Impulse Response Function (Low Political Uncertainty)

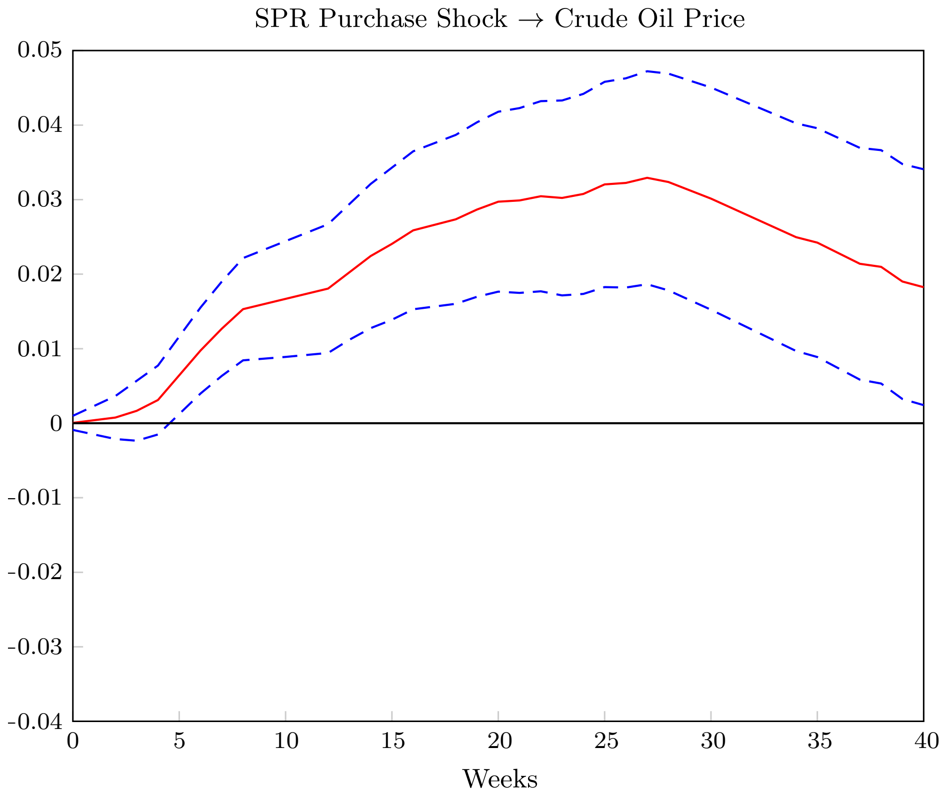

Figure 18: SPR Purchase Impulse Response Function (High Oil Market Volatility)

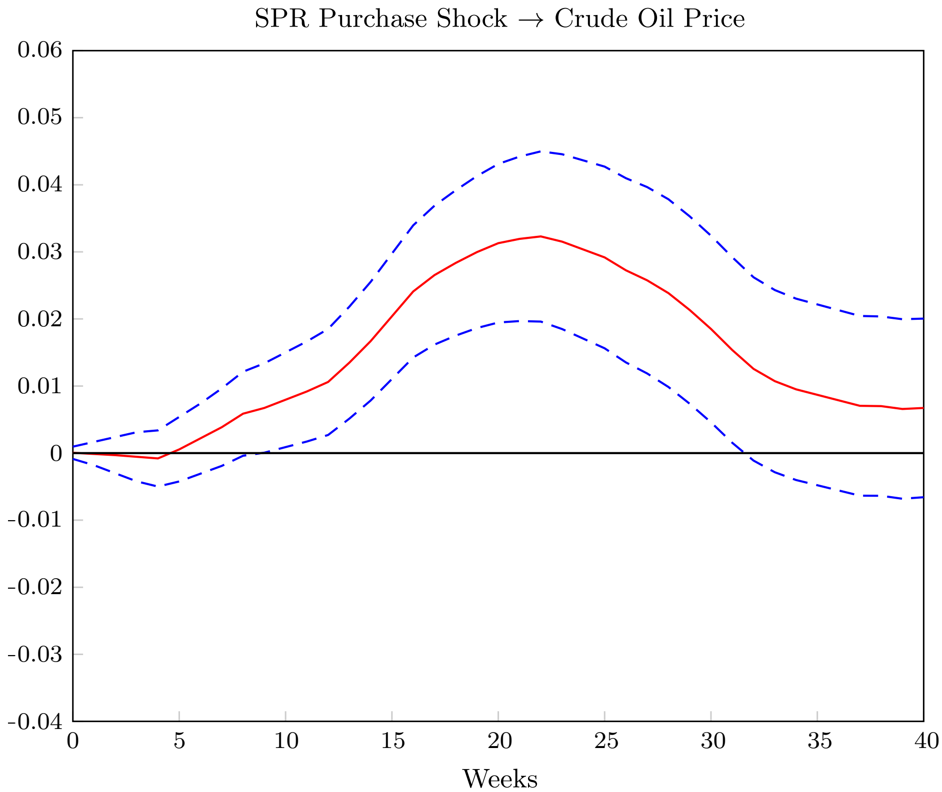

Figure 19: SPR Purchase Impulse Response Function (High Stock Market Volatility)

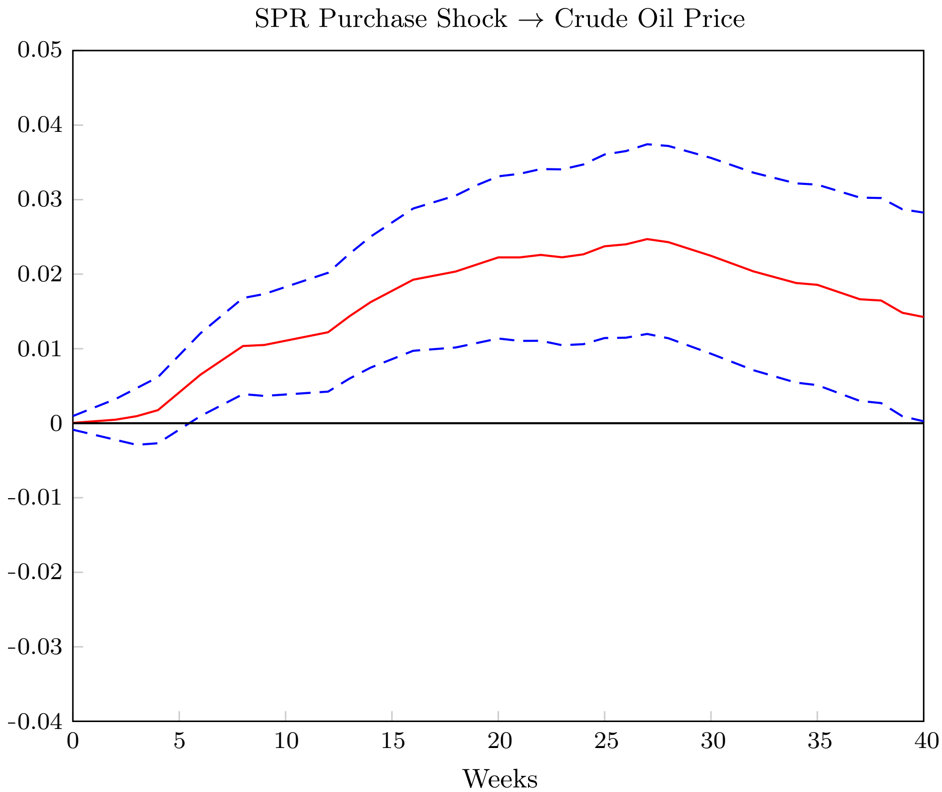

Figure 20: SPR Purchase Impulse Response Function (High Political Uncertainty)

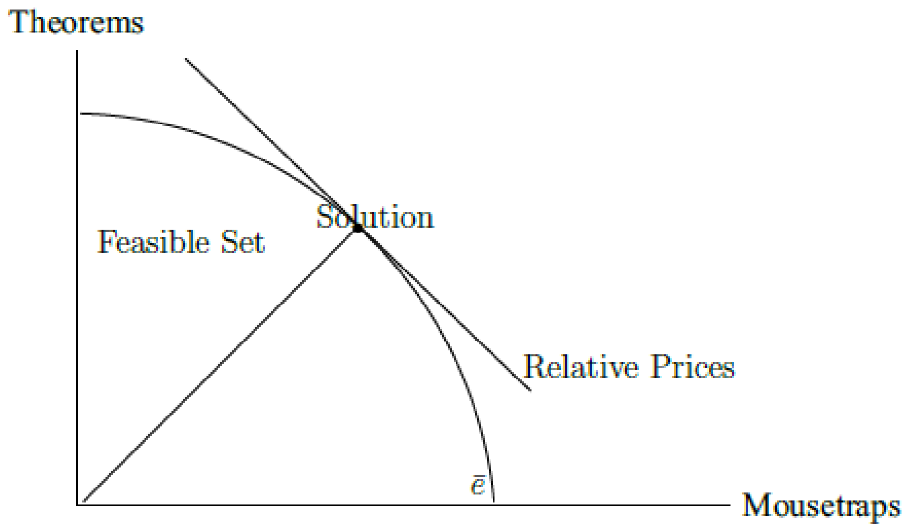

Figure 21: The Production Possibilities Frontier and Expansion Path

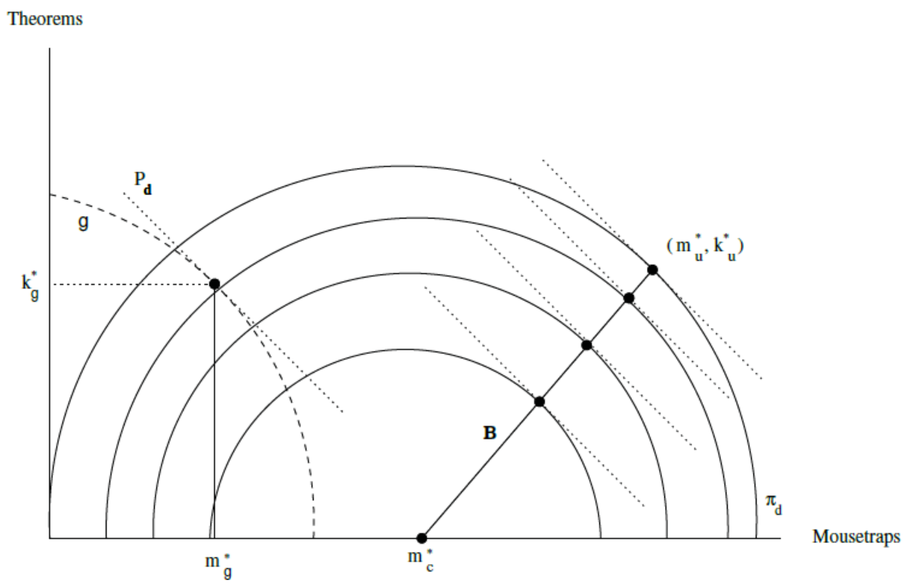

Figure 22: Construction of Threat Points and Bargaining Set, $\mathcal{B}$

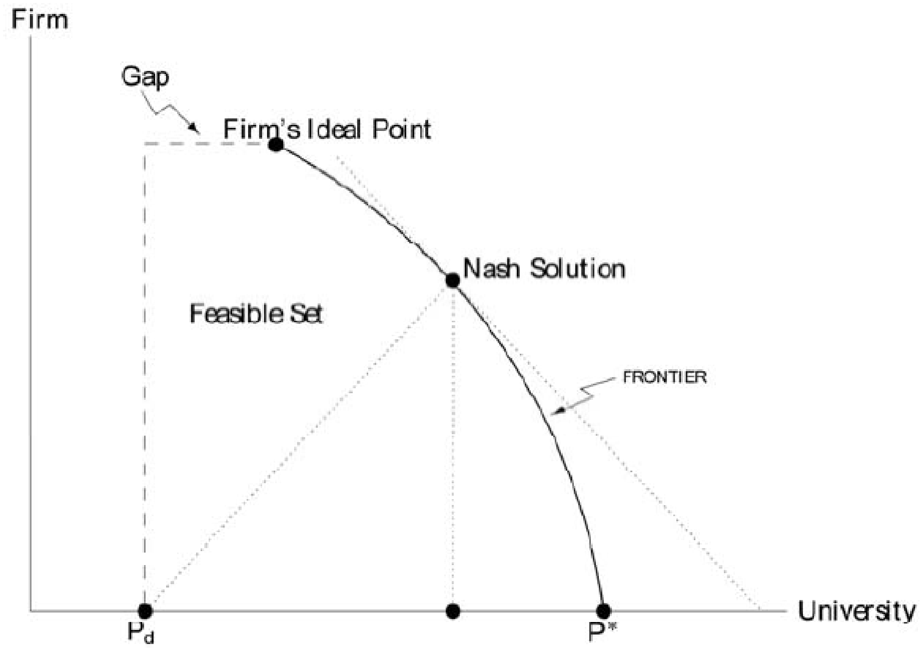

Figure 23: Nash Bargaining Set, Disagreement Point and Solution

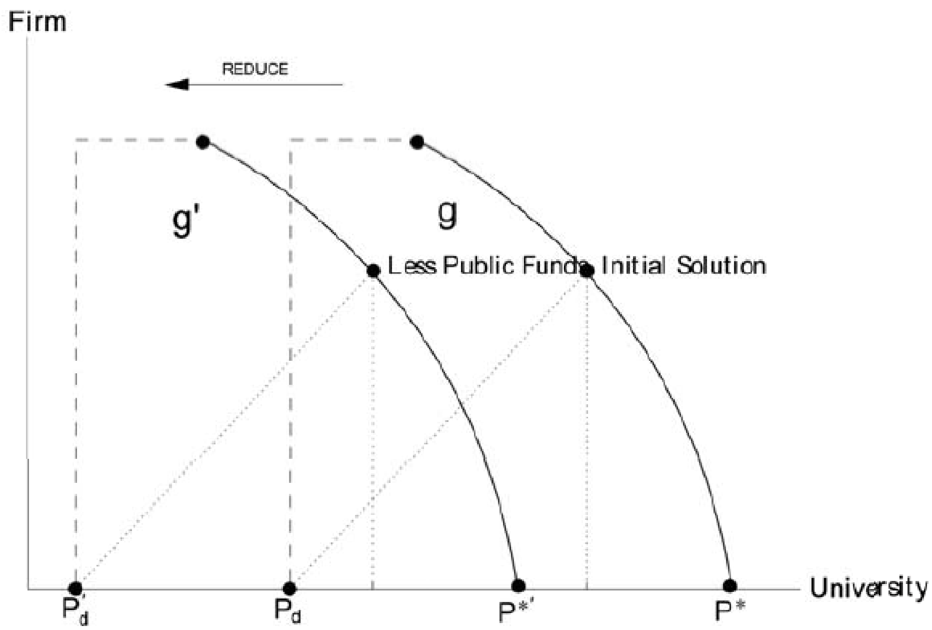

Figure 24: A geometric view of neutrality

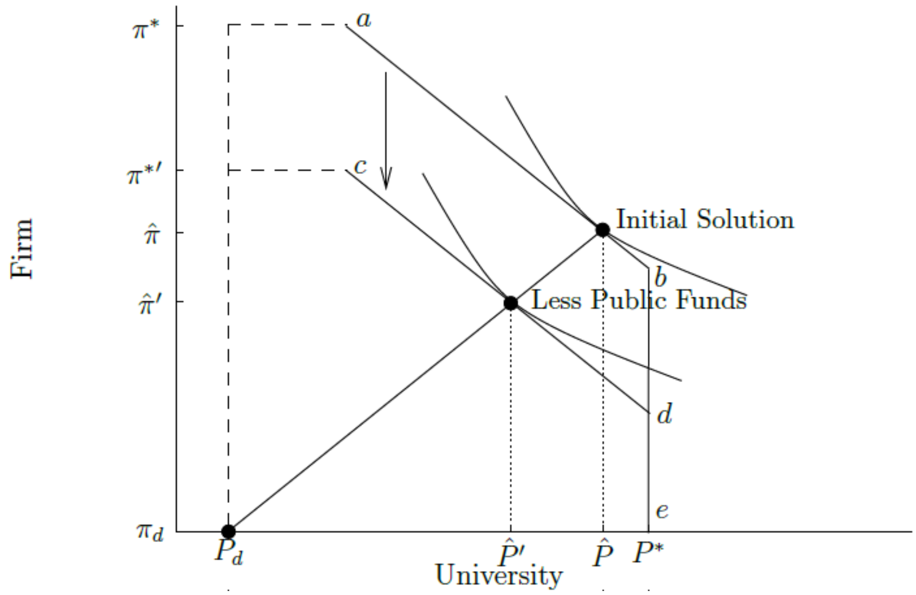

Figure 25: Nonseparable Benefits can lead to crowding-out

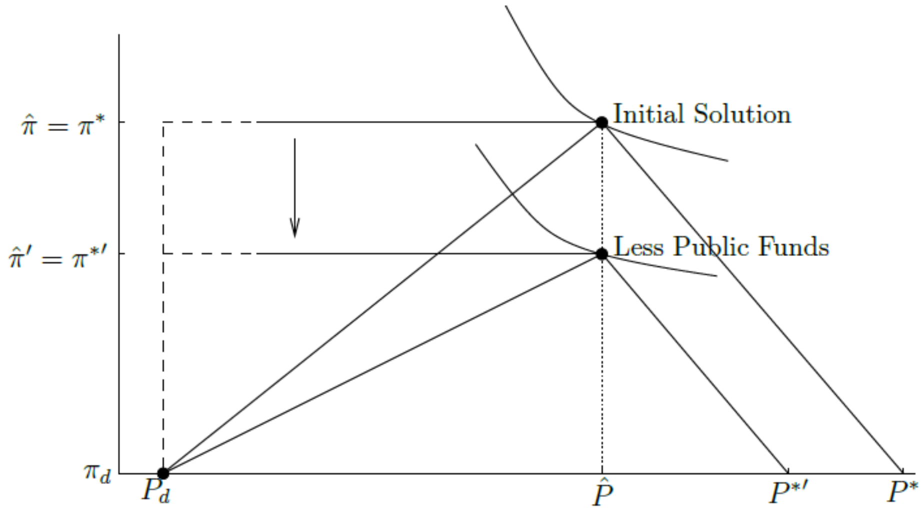

Figure 26: Nonseparable benefits can lead to crowding-in

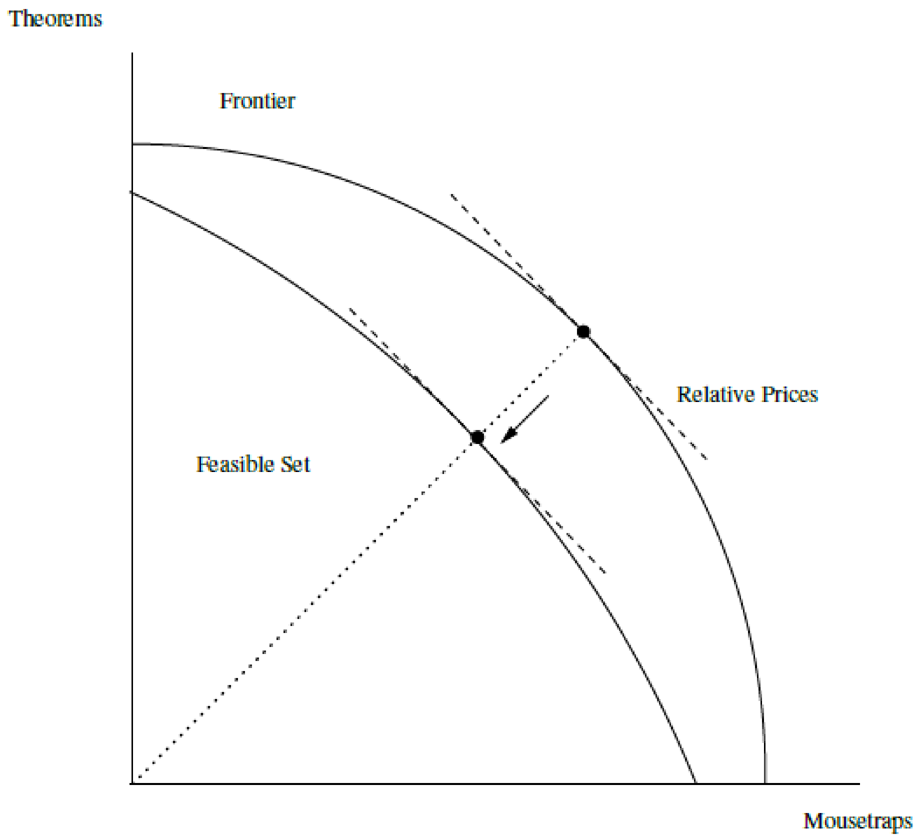

Figure 27: A geometric view of neutrality

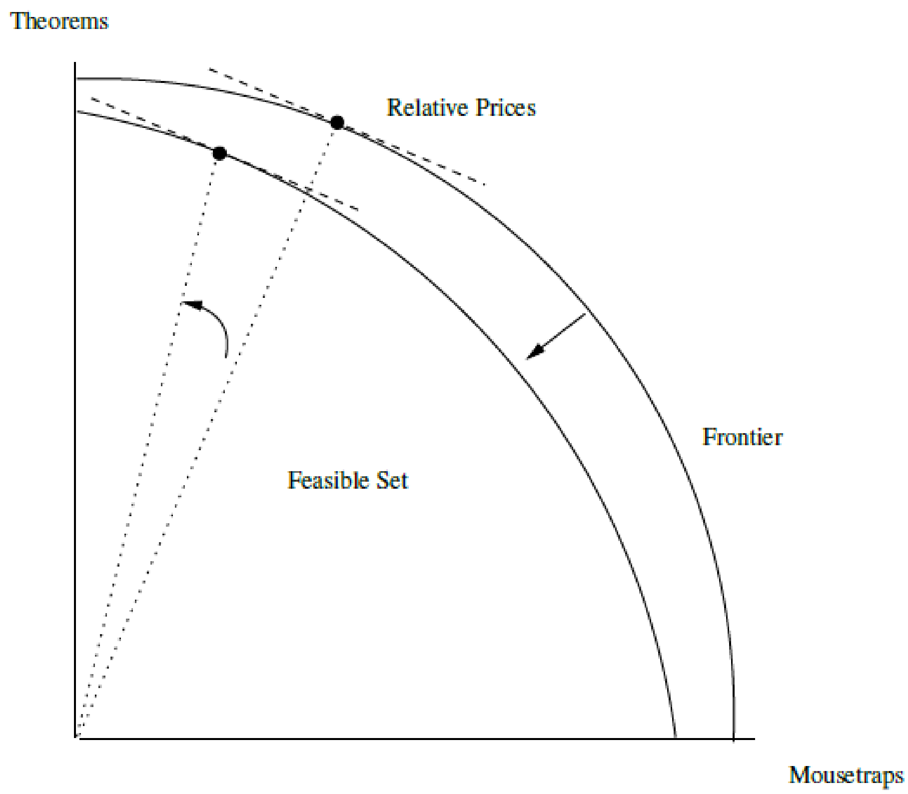

Figure 28: Decrease in governmental funding leads to crowding-in

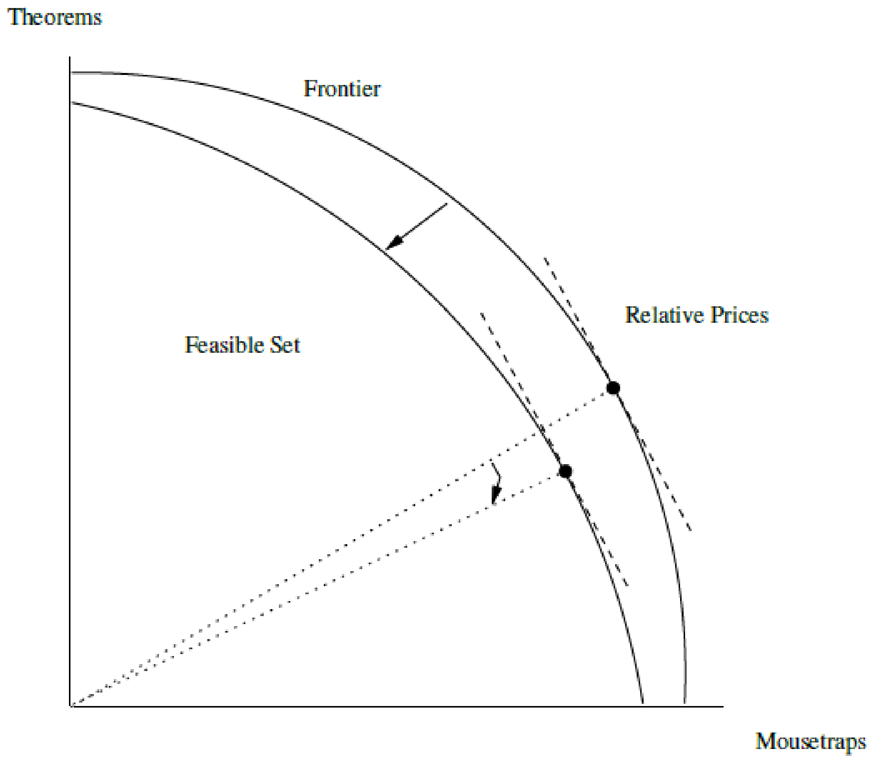

Figure 29: Decrease in governmental funding leads to crowding-out

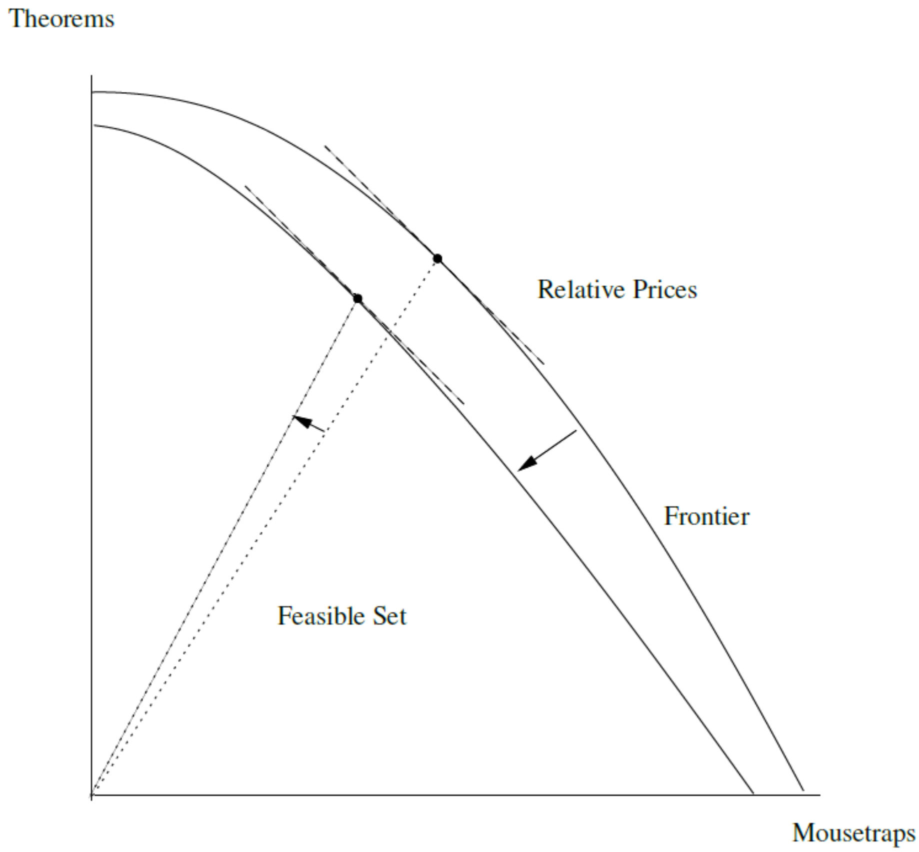

Figure 30: Asymmetric feedback leads to crowding-in

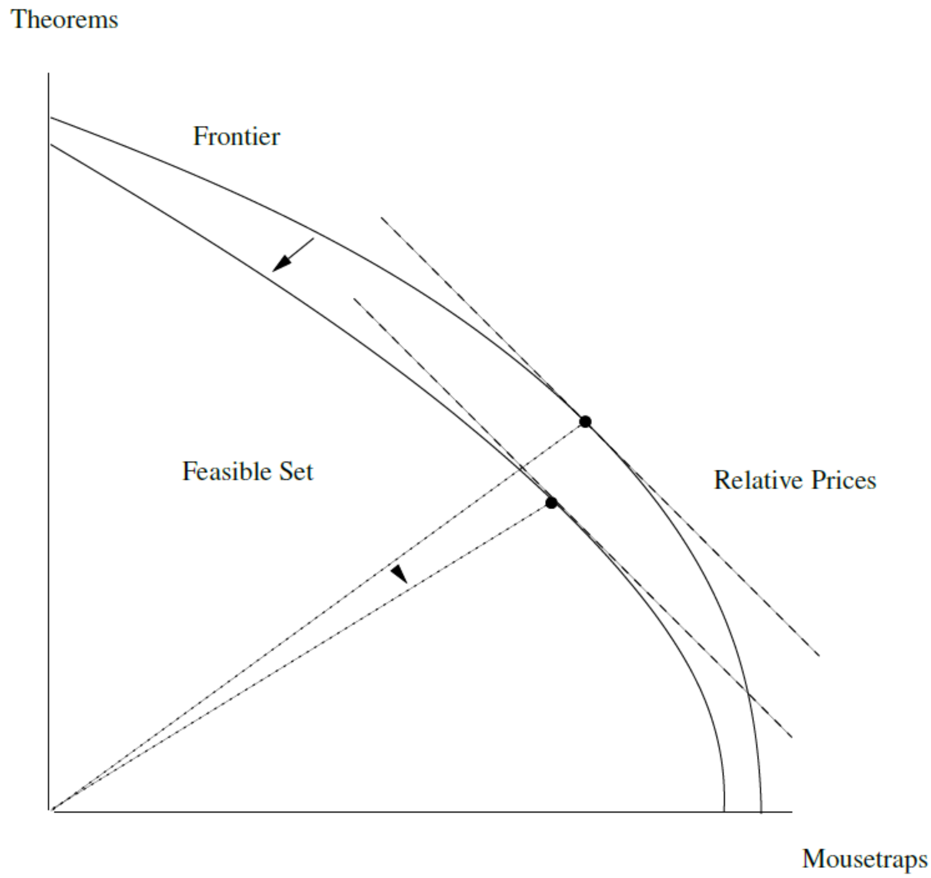

Figure 31: Asymmetric feedback leads to crowding-out

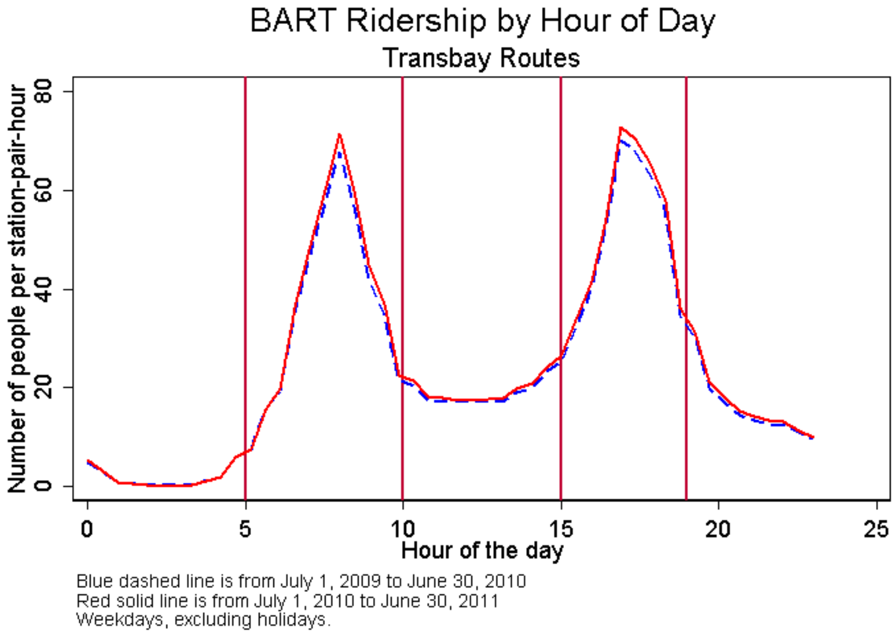

Figure 32: Hourly Transbay BART Ridership

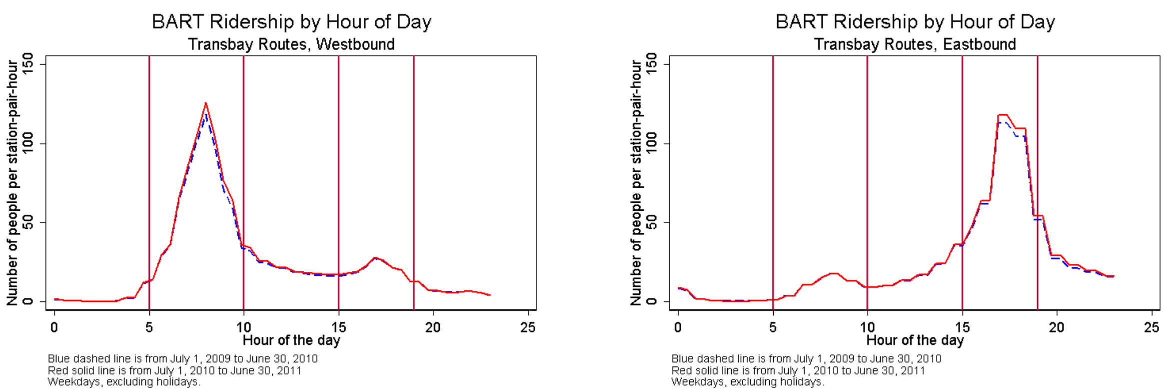

Figure 33: Westbound (left) and Eastbound (right) Hourly BART Ridership

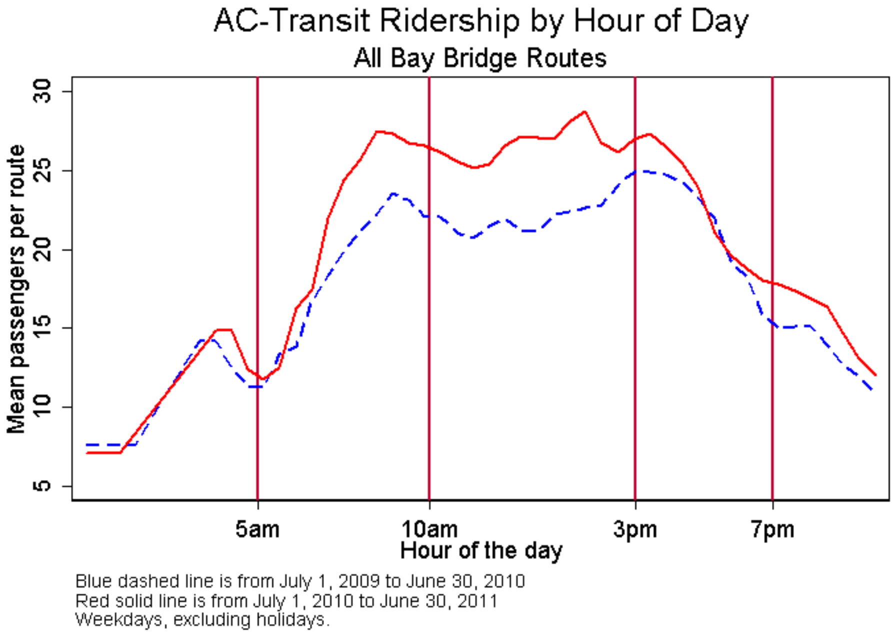

Figure 34: Hourly Transbay AC-Transit Ridership

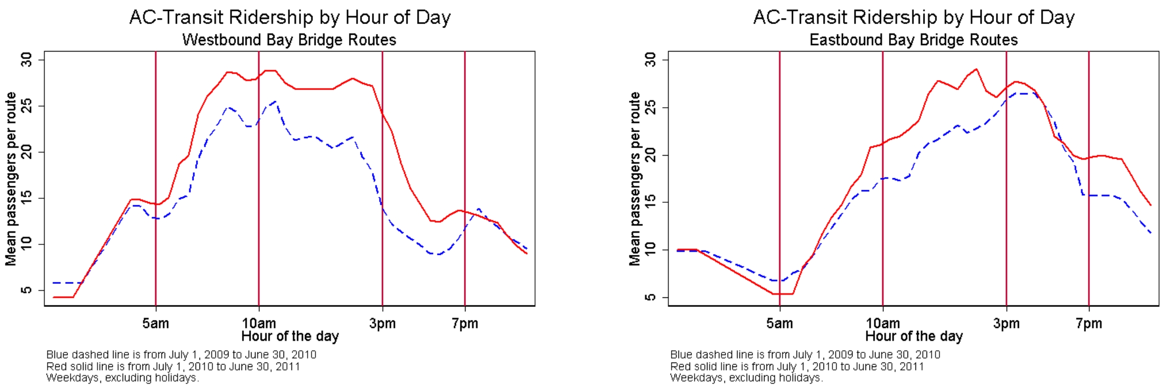

Figure 35: Westbound (left) and Eastbound (right) Hourly AC-Transit Ridership

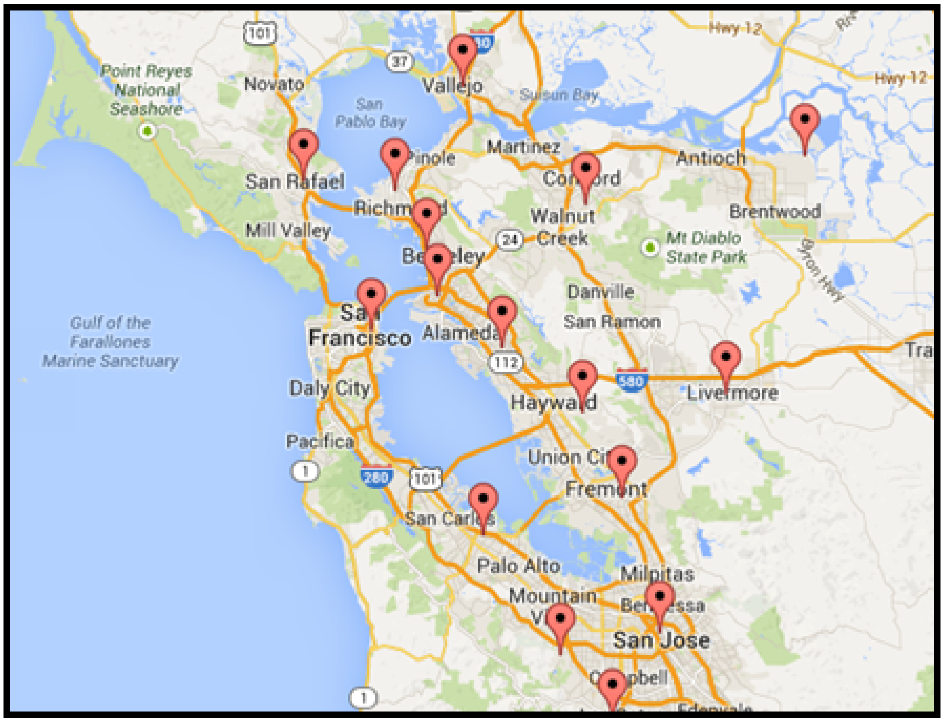

Figure 36: San Francisco Bay Area TAC Monitor Stations

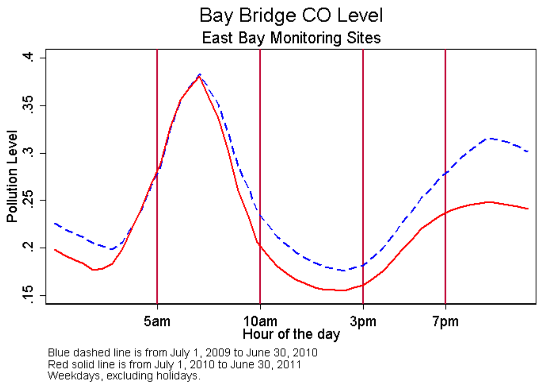

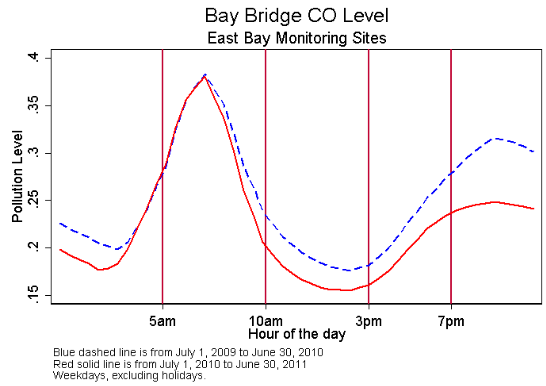

Figure 37: Average Hourly CO Levels

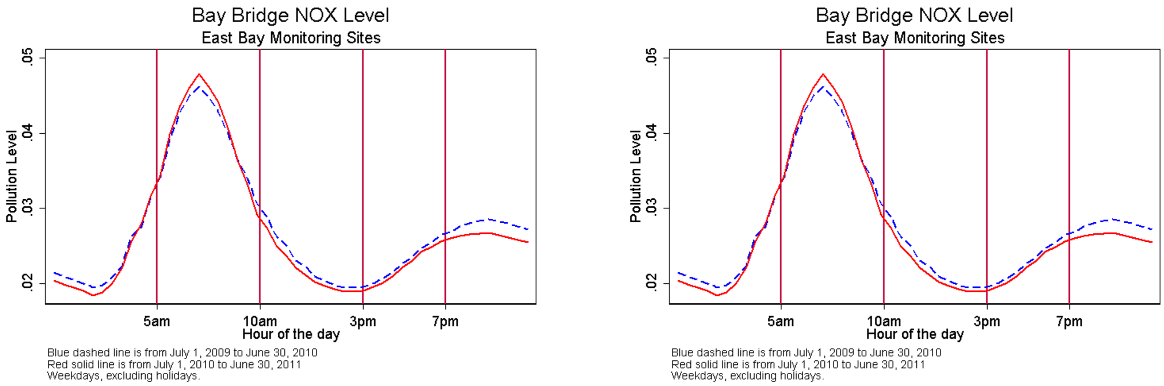

Figure 38: Average Hourly NOX Levels

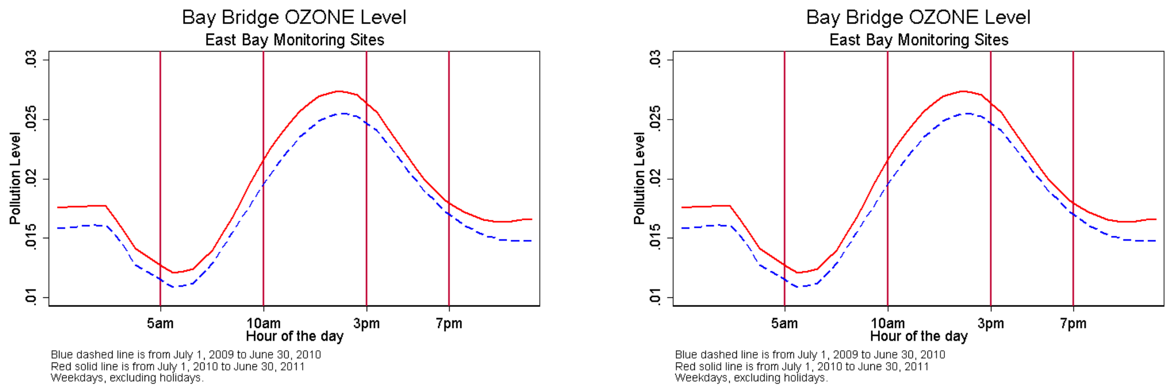

Figure 39: Average Hourly Ozone Levels

Dependent Variable:	BART		AC-Transit	
Riders per hour	(1)	(2)	(1)	(2)
After Congestion Pricing Implemented	0.99***	0.363***	2.23***	-0.47
	(0.05)	(0.06)	(0.16)	(0.87)
Peak Hours		37.58***		14.62***
		(0.09)		(0.77)
Off-Peak Hours		7.96***		9.34***
		(0.04)		(0.81)
Peak Hours x After		1.43***		2.87***
		(0.14)		(0.88)
Off-Peak Hours x After		0.16**		3.24***
		(0.06)		(0.94)
Contant	23.57	4.83***	21.33***	7.58***
	(0.04)	(0.03)	(0.12)	(0.93)
Observations	2,305,152	2,305,152	27362	27362
F Statistic	342	66,262	207	523

Notes: Heteroskedasticity-robust standard errors reported. *** denotes statistical significance at the 1 percent level and ** denotes statistical significance at the 5 percent level.

Figure 40: Public Transit Ridership Regressions

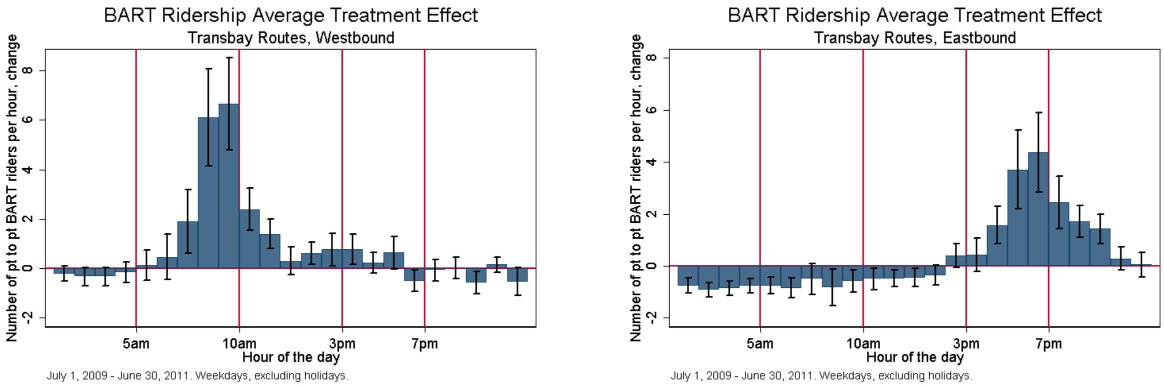

Figure 41: Westbound (left) and Eastbound (right) BART Ridership Estimates

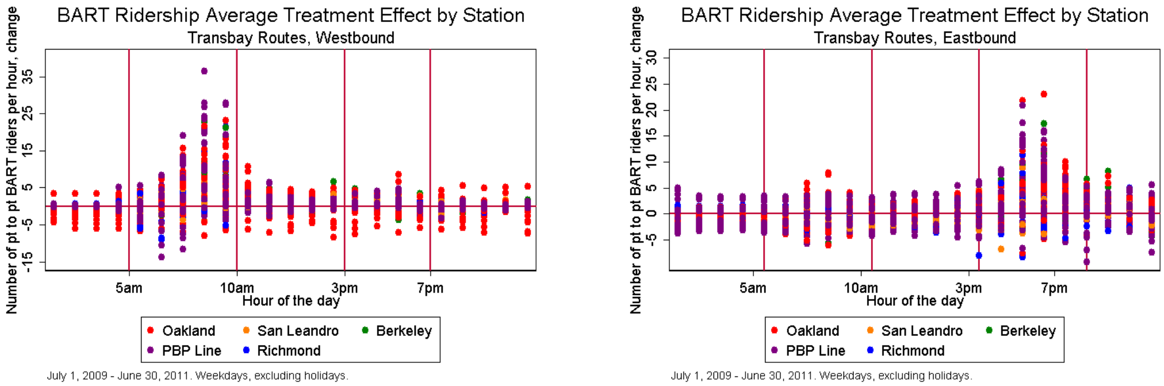

Figure 42: Westbound (left) and Eastbound (right) Station-level BART Ridership Estimates

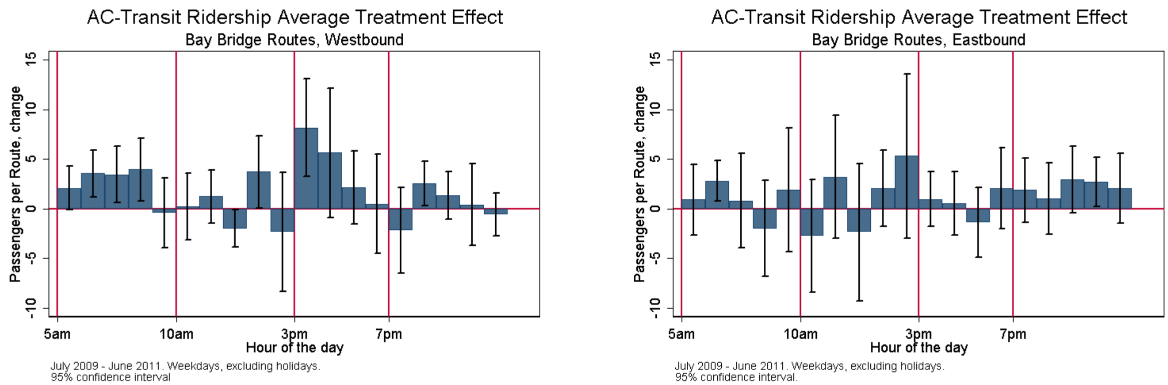

Figure 43: Westbound (left) and Eastbound (right) AC-Transit Ridership Estimates

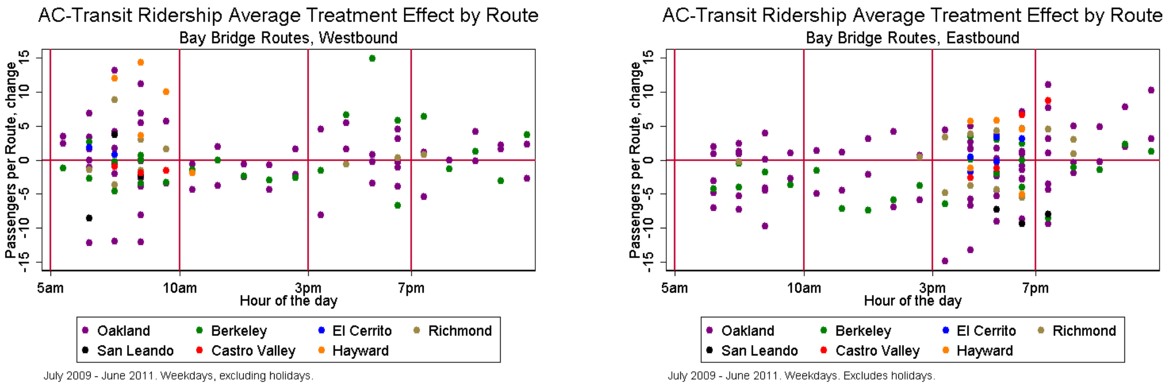

Figure 44: Westbound (left) and Eastbound (right) Route-level AC-Transit Ridership Estimates

Dependent Variable:	CO		NOX		Ozone	
Riders per hour	(1)	(2)	(1)	(2)	(1)	(2)
After Congestion Pricing Implemented	-0.01**	-0.007	0.001**	-0.0011	0.001***	0.0014**
	(0.002)	(0.011)	(0.0003)	(0.0012)	0.00013	(0.0006)
Peak Hours		0.055***		0.0097***		0.0003***
		(0.008)		(0.0009)		(0.0005)
Off-Peak Hours		0.010		0.0018**		0.0023***
		(0.008)		(0.0008)		(0.00046)
Peak Hours x After		0.010		0.0043**		-0.0002
		(0.012)		(0.0013)		(0.0007)
Off-Peak Hours x After		-0.006		0.00042		0.0001
		(0.01)		(0.0012)		(0.00067)
Contant	0.25***	0.22***	0.027***	0.022***	0.019***	0.017***
	(0.002)	(0.008)	(0.0002)	(0.0008)	0.0001	(0.0004)
Observations	41671	41671	42490	42490	34720	34720
F-Statistic	6	87	6	171	108	77

Notes: Heteroskedasticity-robust standard errors reported. *** denotes statistical significance at the 1 percent level and ** denotes statistical significance at the 5 percent level.

Figure 45: Air Pollution Regressions

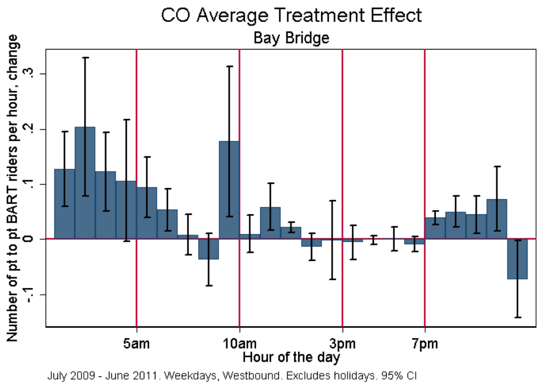

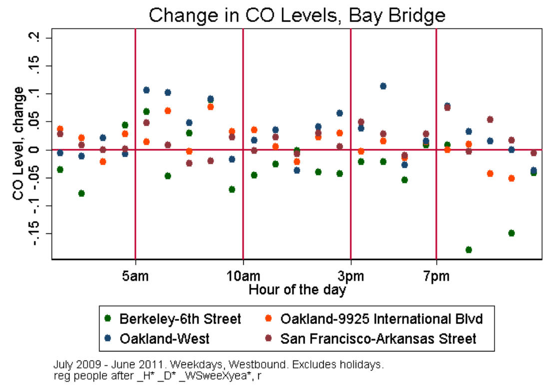

Figure 46: Overall (left) and TAC Station-level (right) CO Estimates

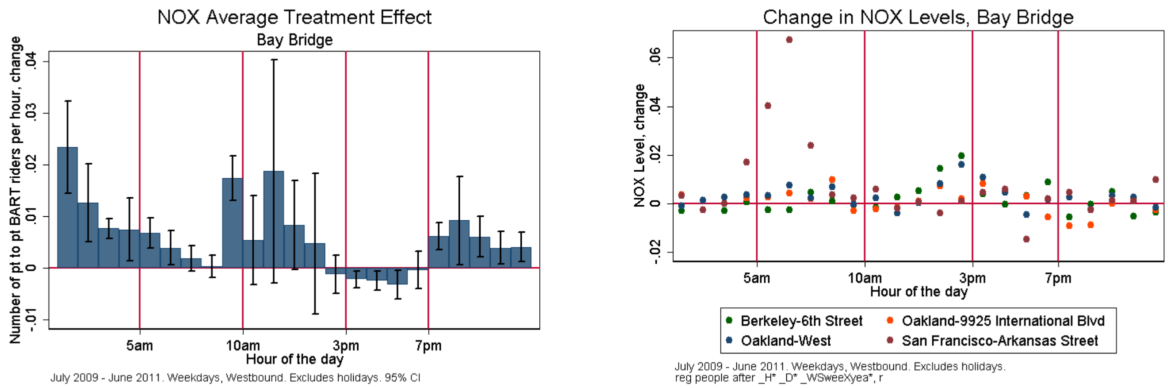

Figure 47: Overall (left) and TAC Station-level (right) NOX Estimates

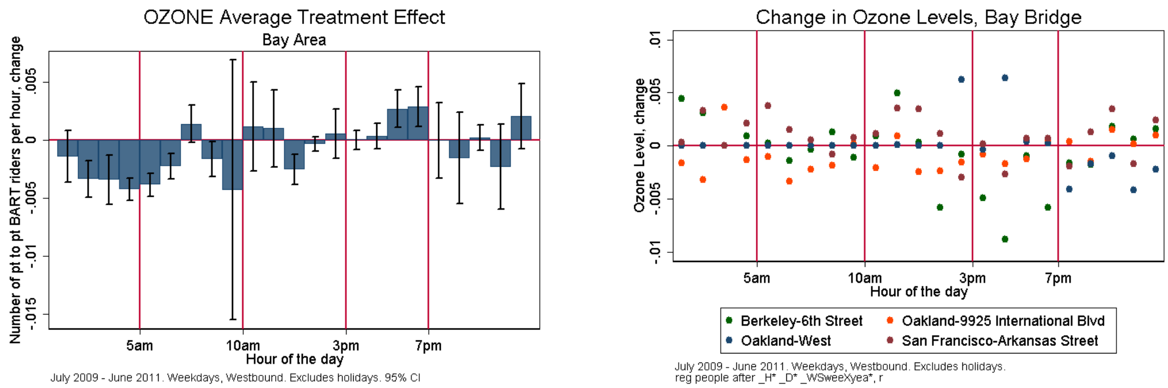

Figure 48: Overall (left) and TAC Station-level (right) Ozone Estimates

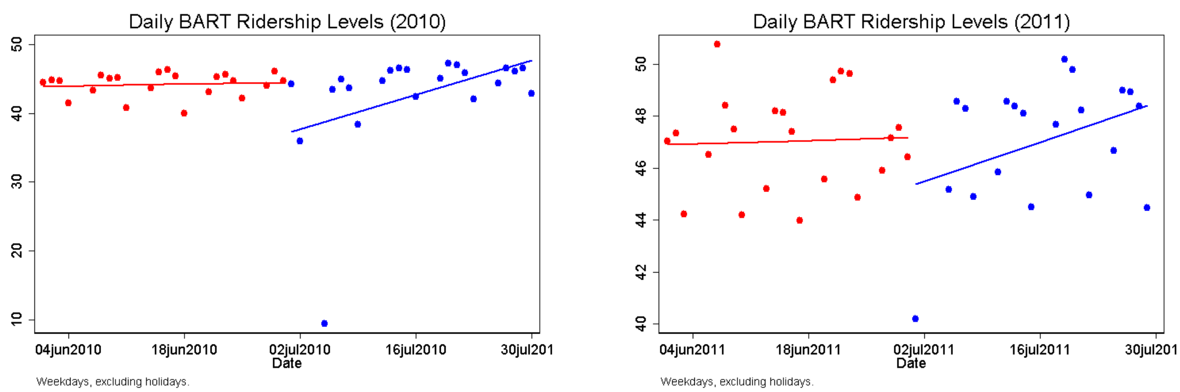

Figure 49: BART Regression Discontinuity Graphs

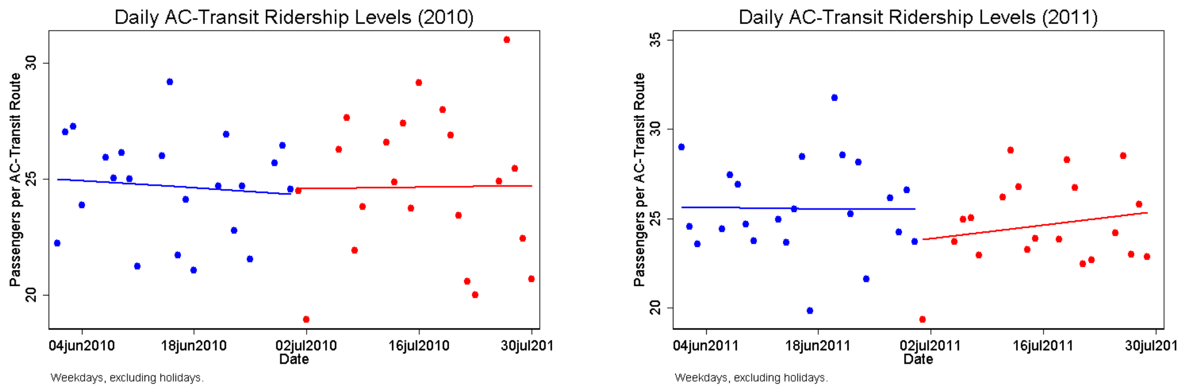

Figure 50: AC-Transit Regression Discontinuity Graphs

	1983-1989	1990-1994	1995-2001	2002-2008	2009-2014
SPR Purchase	-0.02	-0.07	-0.01	0.06	0.00
SPR Release	0.00	0.02	0.06	-0.05	0.1

Table 1: Average Structural Policy Residuals for Subperiods