ABSTRACT

CHEN, XIAOMEI. Discounting Investments that Save Energy Resources. (Under the direction of Michael Roberts and Laura Taylor.)

In the first chapter, I focus on the relationship between energy futures and the aggregate economy. Wide and frequent variation in the association between energy futures and the overall stock market facilitates a novel test of the conditional Capital Asset Pricing Model (CAPM). Using commodity futures and covariances with the S&P 500, I find excess returns on rollover futures contracts to be strongly associated with market betas, and fail to reject conditional CAPM when weekly beta estimates are corrected for measurement error. Moreover, the time path of estimated betas of crude oil reflects historical shocks to demand and supply that would naturally give rise to betas of different signs. Examples include the Gulf War of 1990-1991 (negative beta), Iraq War of 2003 (negative beta), OPEC oil expansion in 1986 (negative beta), and Recessions of 2000 and 2007(positive betas).

In the second chapter, I examine how uncertainty affects the choice of discount rate especially in the context of energy investments. The expected present value of investments that pay dividends in the form of energy savings is strongly influenced by uncertainty about discount rates. Discount rate uncertainty stems from interest rate uncertainty, risk premium uncertainty and energy-price uncertainty. I develop a state-space model that simultaneously considers all three components of the discount rate and use the model to project uncertainties into the future. From these projections, I estimate certainty-equivalent discount rates for investments with different longevities, including fuel efficient cars (10-15 years), photovoltaic solar panels and windmills (20-30 years). I find certainty-equivalent discount rates are sensitive to model specification but tend to be much lower than discount rates conventionally used in practice.

In the third chapter, I study prices and sales of individual clothes washer models before, during and after a 2007 standard that banned manufacture (but not sale) of low-efficiency units and increased the threshold for Energy Star certification. While quantities sold of washer models banned from manufacture decreased sharply, prices for banned models increased only modestly. At the same time, sales of higher-efficiency units rose markedly while prices for high-efficiency units declined. On average, washer efficiency increased but prices changed little. A simple welfare analysis indicates that consumer welfare loss from banned washers was far outweighed by gains from lower-priced high-efficiency units. While a full cost-benefit analysis is not feasible with the available data, I estimate a lower-bound gain in consumer surplus equal to 3-15 percent of total sales. This result may accord with earlier theoretical research that shows quality standards can increase welfare in monopolistically competitive industries that possess increasing returns to

scale (Ronnen, 1991). Thus, if energy efficiency is a close proxy for quality, energy efficiency standards may increase competition, market efficiency and welfare.

Discounting Investments that Save Energy Resources

by
Xiaomei Chen

A dissertation submitted to the Graduate Faculty of
North Carolina State University
in partial fulfillment of the
requirements for the Degree of
Doctor of Philosophy

Economics

Raleigh, North Carolina

2013

APPROVED BY:

Michael Roberts Co-chair of Advisory Committee	Laura Taylor Co-chair of Advisory Committee
Walter Thurman	Denis Pelletier

Martin Smith

DEDICATION

To my parents, Weiying and Shigui. To my brother, sister-in-law and my lovely niece. I love you all.

BIOGRAPHY

The author came to the United State in 2007, and studied Economics at North Carolina State University.

ACKNOWLEDGEMENTS

I would like to thank my advisers, Michael Roberts and Laura Taylor, for their help and support during my graduate studies. I also want to thank the members of my PhD committee, Professors Denis Pelletier, Walter Thurman, and Martin Smith for their helpful career advice and suggestions in general.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Chapter 1

A Test of Conditional CAPM using High Frequency Energy Futures Prices

1.1 Introduction

The Capital Asset Pricing Model (CAPM) is the benchmark theory for how markets incorporate risk into asset prices. The theory shows that an asset's *beta* connects asset-specific risk premia to the economy-wide market risk premium. *Beta* measures the association of an asset with the market portfolio. The higher an asset's beta, the less of its risk can be diversified away, and thus the higher its risk premium.

While the theory remains a standard part of economics and business school training, empirical evidence in support of the theory remains varied and inconclusive. Early empirical efforts failed to reject the theory, but these tests lacked statistical power, in part because most equities covary together. Without large differences in predicted risk premia, stocks give little basis for comparison. Earlier analysis of cross-sectional data finds a relationship between realized returns and estimated market beta that is too flat relative to CAPM predictions (Black, Jensen and Scholes 1972, Blume and Friend 1973, Fama and MacBeth 1973, and many others). Later work by Fama and French (1992, 1996) devastated CAPM by showing betas explained none of the predictable differences in returns across portfolios sorted by firm size and book-to-market ratios, the largest and most persistent predictable differences in stock returns that have been identified.

Jagannathan and Wang (1996), and similar work by Lettau and Ludvigson (2001), attempted to resurrect the CAPM by showing that both asset betas and overall market risk premium vary with business cycles. Stocks of small companies and "value" stocks with high book-to-market ratios have large betas during recessions when market risk premia are large, and small betas during booms when market risk premia are small. For large companies and "growth" stocks with low book-to-market ratios, the opposite pattern emerges. These relationships show how CAPM can reconcile the relatively higher average returns of small company stocks and value stocks as compared to stocks of large companies and growth stocks. More recently, Lewellen and Nagel (2006) challenged these findings by arguing it would require betas to vary by an implausible amount.

In this paper we consider a general conditional CAPM model in the context of commodity market returns, rather than stocks. Commodity markets provide a particularly interesting area to test the predictions of CAPM because commodity prices do not generally covary systematically with each other or with stocks (Bodie and Rosansky, 1980). That is, different commodities have very different betas, facilitating a larger scope for comparison and greater statistical power.

Commodity betas are also likely to vary over time for clear fundamental reasons that may or may not be related to business cycles. For example, an unexpected discovery of a large new oil deposit would presumably cause oil prices to decline while simultaneously providing good news for the aggregate economy and stock returns. Such an event would be embodied by a *negative* beta and risk premium under CAPM. Conversely, larger-than-expected growth in emerging markets may drive up demand of both stocks and oil prices, giving rise to a positive beta. Precious metals like gold may also be used as a hedge against inflation or financial crises, and thus may have varying and possibly even negative risk premiums. For these reasons, commodity markets provide a wholly different and powerful area for testing CAPM.

In this study, we focus on energy products — crude oil, natural gas, heating oil and gasoline — which are important industrial commodities closely related to the overall economy that have been traded in futures markets for many decades. The well-documented connections between the energy market and the macroeconomy also make these commodities interesting when testing CAPM (Hamilton 1983, Barsky and Kilian 2004, Kilian 2009). This literature indicates that through much of post-World-War II history, supply shocks or anticipated supply shocks have been a major source of the price volatility, particularly for oil. These shocks have mainly followed from conflicts in the Middle East, actions taken the Organization of Petroleum Exporting Countries (OPEC), and discovery of major new oil deposits. These shocks likely played at least a contributing role in many post World War II recessions and often gave rise to a negative association between energy price changes and overall market, hence negative average betas. At other times, and especially recently, energy price changes have been associated with aggregate demand shocks following from global booms, recession and recovery, likely giving rise to positive betas.

The relative frequencies of positive-beta events and negative-beta events may change over time, for many reasons, sometimes with high frequency. For example, recently the world economy experienced large aggregate demand shocks following the 2008 financial crisis, likely giving rise to positive betas. At the same time, however, civil conflict in Libya and parts of the Middle East gave rise to supply shocks and fear of even larger future supply shocks, which could create negative betas. It is therefore easy to see how beta size and sign could change considerably, perhaps even from one week to the next, depending on which kind of events were perceived by the market to be more prevalent.

It requires little introspection to see that, for all assets, betas may change over time quite arbitrarily. Similarly, the risk premium for the market-portfolio may change over time, depending upon perceived uncertainty about aggregate economic growth. The basic asset pricing theory provides no guidance about the structure of these deeper sources of time-varying variances and covariances. More challenging still is that changing expectations about market returns and asset-specific covariances, including the examples for energy-specific assets described above, are unlikely to be quantifiable using objective measures or factors. For this reason, Cochrane (2001) (with attribution to Hansen and Richard (1987)) argues that CAPM, in its most generic form, is fundamentally untestable.

Without some structure to pin down clear comparisons between relatively low-risk-premium assets and high-risk-premium assets, there is little defensible scope for testing the theory. A statistical rejection CAPM may result because the econometrician has assumed conditioning information or some structure that causes assets to have been assigned to incorrect betas and/or risk premiums.

The key contribution of this analysis is to provide a simple structure for time-varying betas that does not assume betas are fixed over time (Dusak 1973; Bodie and Rosansky 1980) or require factor variables to account for changing conditional expectations (Merton 1973, Breeden 1979, Cox, Ingersoll Jr and Ross 1985 among others). Instead of using conditioning factors (which are likely unmeasureable in most cases), we estimate betas using short-window regressions that follow in the spirit of Lewellen and Nagel (2006). We also consider a model that assumes a latent stochastic structure for an underlying beta process, which we are able to estimate using a Kalman filter.

To allow betas to change at a high frequency, we use an approach developed by Lewellen and Nagel (2006) and directly estimate betas and risk premia unconditionally through short-window regressions. This approach is simple and requires fewer structural assumptions than earlier empirical tests of conditional CAPM (Jagannathan and Wang 1996, Lettau and Ludvigson 2001 etc.). Using daily returns in the estimating of the short-window regressions, we use a 2-week period as a baseline window size.

Because bi-weekly betas are estimated using small sample sizes, sample error causes attenuation bias when actual returns are regressed against estimated risk premiums. We correct for the bias in three ways: (1) using the previous period's estimated beta as an instrument for the current-period beta; (2) a Kalman filter assuming the underlying true beta follows a random walk; and (3) using simulation extrapolation (SIMEX) (Cook and Stefanski 1994). SIMEX is simulation based technique that adds a series of additional error vectors to the observed data, with each vector series having a different variance. The series of coefficient estimates are then used to extrapolate what the estimate would be if there were no error in the estimated beta; and the results are asymptotically unbiased and consistent (Cook and Stefanski 1994).

The results show considerable variations in estimated betas over time for all energy commodities. Crude oil betas, for example, have a mean of 0.01 and standard deviation of 1.1. An estimated 62% of the beta variance is due to measurement error, and 38% due to true variation in betas. Crude oil and petroleum products such as gasoline and heating oil are correlated with oil but have slightly different betas.

We find realized return premiums vary significantly with estimated betas over time and across commodities. Where most earlier tests find realized premiums to be smaller than CAPM-estimated premiums, we find realized return premiums increase somewhat more than one for one with estimated premiums.

We test the short-window CAPM by examining (1) the statistical properties of estimated bi-weekly pricing error, *alpha*, and the risk index, *beta*; (2) the economic interpretation of the estimated *beta*; and (3) the comparison between the realized risk premiums and its CAPM predicted counterparts over time. In (1) the pricing error is the deviation of the realized risk premium from its CAPM prediction which can be captured by the intercept in the short-window regressions. The CPAM predicted risk premium can be calculated using the estimated beta, and (3) can be achieved by regressing actual returns on the calculated ones.

1.2 Methods

To test CAPM for commodity trading, we consider speculative trading of commodity futures. Commodity futures are a natural setting for testing CAPM because investors can speculate and trade in these markets with minimal transactions costs and without physically holding or storing the commodity. Specifically, we consider buying and subsequently selling of futures contacts. Contracts are always sold at least 1 week before the final trading date such that the speculator need not ever take physical possession of the commodity. Contracts for the nearest delivery date are bought and sold up to one week before the last trading day, at which time contracts are rolled over to the subsequent delivery date. Arbitrage pricing rules apply to these derivative contracts, and via arbitrage the associated risk premia ought to apply equally to the physical commodities.

Define the futures contract price for delivery in period T_d at time t as $F(t, T_d)$. The nominal return in period t for holding the futures contract from period $t-1$ to period t is therefore:

$$r_t = ln(F(t, T_d)) - ln(F(t-1, T_d))$$

In practice, to buy a futures contract, the buyer must post a Treasury bill as collateral for funds promised on the delivery date. The return on the Treasury bill must therefore be added to r_t. And the daily commodity realized return is:

$$R_t = r_t + R_t^f$$

and R_t^f is the treasury return. We use 3-month t-bill as the risk-free asset, and its daily return is calculated following standard Treasury bill return calculation and formula is specified as follows:

$$R_f = \left[\frac{1}{1 - r_f \cdot (91/360)}\right]^{(days/91)} - 1$$

where r_f is the 3-month Treasury bill rate, calculated as the daily secondary market quote on the most recently auctioned Treasury Bills for 3-month maturity. In case of missing rate, the rate of the nearest neighbor prior will be used. The variable *days* refers to the number of calendar days investor holds the risk-free asset. This equation follows the formula used in "The Dow Jones-UBS Commodity Index^{SM} Handbook" for calculating 3-month t-bill return.

We use S&P 500 as the proxy to the overall-market, and its daily return is defined as:

$$R_t^M = ln(P_{Mt} + D_{Mt}) - ln(P_{Mt-1})$$

where D_{Mt} is the daily dividend calculated from the reported quarterly dividend.

The asset risk premium is defined as the excess return in excess of the risk-free return. Accordingly, the risk premium in the futures market is

$$\begin{aligned} r_t &= R_t - R_t^f \\ &= ln(F(t, T_d)) - ln(F(t-1, T_d)) \end{aligned}$$

Whereas the market risk premium is

$$r_t^M = R_t^M - R_t^f$$

We relate commodity returns to S&P 500 returns by estimating a separate regressions for each commodity and each 2-week period

$$r^i_{t,\tau} = \alpha_{i\tau} + \beta_{i\tau} r^M_{t,\tau} + \epsilon^i_{t,\tau} \tag{1.1}$$
$$t = 1, \cdots, 10$$

where $r^i_{t,\tau}$ and $r^M_{t,\tau}$ are the realized risk premium on commodity i and the market portfolio respectively of the 2-week period τ. $\alpha_{i\tau}$ is the deviation of realized risk premium from the CAPM prediction, or the pricing error. And $\beta_{i\tau}$ represents the exposure of asset i to market risk in the 2-week period τ.

1.2.1 Testing the Short-Window CAPM

The time series of α_τ and β_τ we obtain from estimating equation 1.1 enable several different tests of CAPM. The most straightforward ways are testing whether pricing error, α, is zero, and whether the market risk coefficient, β, reflects the true risk premia.

Further, by comparing the realized risk premiums to the CAPM predictions, which equal to the market risk premiums multiplied by *beta*, we can understand how accurately the model predicts over time. This test can be achieved through a simple regression using estimated bi-weekly *betas* and realized market and asset-specific risk premiums over time for each of the commodities:

$$r^i_\tau = \delta_0 + \delta_1(\hat{\beta}_{i\tau} r^M_\tau) + \mu_{i\tau} \tag{1.2}$$

where $\hat{\beta}_{i\tau}$ is the estimated beta of asset i as of 2-week period τ. r^i_τ and r^M_τ are bi-weekly risk premiums of asset i and market portfolio respectively.

If CAPM predicts the risk premium, the realized risk premium equals its CAPM prediction, i.e. $r^i_\tau = \hat{\beta}_{i\tau} r^M_\tau$, plus a pricing error. It gives us testable hypotheses that $\delta_0 = 0$ and $\delta_1 = 1$.

1.2.2 Error-in-variable Corrections

Bi-weekly betas are estimated in the previous step using small sample sizes, which leads to large estimation error and causes an attenuation bias when they are used directly in estimating equation 1.2. In this section, we propose three ways to correct for the bias: (1) use of an instrumental variable, (2) Kalman filter, and (3) simulation extrapolation.

Although market conditions can change rapidly, it is reasonable to expect that the nature of perceived uncertainties will be similar enough from week to week and month to month that past betas can serve as a viable instrument for current betas. If betas have sufficient serial correlation, which is testable, and expectation errors are independent, as implied by an absence of arbitrage, then exclusionary assumptions are satisfied and instrument has a strong foundation. In the first stage we estimate model 1.2:

$$\hat{\beta}_{i\tau} = \phi_0 + \phi_1 \hat{\beta}_{i\tau-1} + \nu_{i\tau-1} \tag{1.3}$$

We then use the fitted values of *beta* in equation 1.3 for the estimation of model 1.2 in the second stage.

In another attempt to correct for the estimation error, we use the Kalman filter in the state space model. The state space model is a flexible form that can be formulated to model dynamic time series with measurement error. A local level model is used in this study, and the state space representation of short-window betas can be written as:

$$\hat{\beta}_t = \beta_t + \upsilon_t \tag{1.4}$$

$$\beta_t = \beta_{t-1} + \omega_t \tag{1.5}$$

where $\upsilon_t \sim \mathcal{N}(0, V)$ and $\omega_t \sim \mathcal{N}(0, W)$.

Equation 1.4 states that the $\hat{\beta}_t$ estimated using the short window regressions is composed of the true *beta*, β_t, and the estimation error, ϵ_t. Equation 1.5 assumes that the true *beta* follows a random walk. This purpose of this method is to infer relevant properties of the unobserved true β from the observations of $\hat{\beta}$. We modify the model derivation in Harvey, Ruiz and Sentana (1992) and Kim and Nelson (1999), use Kalman Filter recursions to filter out the error term in Equation 1.4, and get the series of underlying true beta through the process. Detailed model modifications and derivation process can be found in Chapter 2.

In the first two methods, we correct for the bias using some econometric assumptions of the time series properties of *betas*. In this part, we introduce a statistical technique, Simulation extrapolation (SIMEX), which is simulation based without making assumptions on *beta*. SIMEX is an error correction technique developed by Cook and Stefanski (1994), where by estimating the regressions with additional errors added, we can establish a trend between estimates and variances of additional error and then extrapolate back to where there is no error. SIMEX works well when measurement error is known or can be well approximated, and since betas are estimated using short-window regressions, we have the estimation error for each estimated *beta* and use the arithmetic average of these errors as a proxy of the measurement error. The relationship between the estimated beta and the true beta is as follows:

$$\hat{\beta}_\tau = \beta_\tau + \sigma Z_\tau \tag{1.6}$$

where Z_τ is a standard normal random variable independent of $\hat{\beta}$ and β; and σ^2 is the measurement error variance.

And we use the following step of SIMEX to correct for this error.

1. Add pseudo errors to the estimated betas

$$\bar{\beta}_\tau = \hat{\beta}_\tau + \lambda^{1/2}\sigma Z_\tau \tag{1.7}$$

where λ is chosen with known increment.

2. Estimate regressions using the constructed betas

$$r_\tau^i = \delta_0^\lambda + \delta_1^\lambda(\bar{\beta}_\tau r_\tau^M) + \mu_{i\tau} \tag{1.8}$$

For each chosen λ, we draw N random numbers from the standard normal distribution, Z. After estimating Equation 1.8 for each Z, we take the arithmetic average of all N estimated δ_0^Z and δ_1^Z. We then get the paired λ and δ.

3. Establish the trend between the estimated average δ_0 and δ_1 with λ. We use different fitting methods for comparison purposes, namely restricted cubic splines, natural splines and polynomial.

4. Extrapolate the trend to where there is no error.

We show as follows that the constructed $\bar{\beta}$ has the mean of the true beta in Equation 1.9; and its variance can be reduced to the variance of the true beta when $\lambda = -1$ in Equation 1.10.

$$\begin{aligned} E(\bar{\beta}_\tau) &= E(\hat{\beta}_\tau) + \lambda^{1/2}\sigma E(Z_\tau) \\ &= E(\beta_\tau + \sigma Z_\tau) \\ &= E(\beta_\tau) \end{aligned} \tag{1.9}$$

$$\begin{aligned} Var(\bar{\beta}_\tau) &= Var(\hat{\beta}_\tau) + \lambda\sigma^2 Var(Z_\tau) \\ &= Var(\beta_\tau + \sigma Z_\tau) + \lambda\sigma^2 \\ &= Var(\beta_\tau) + \sigma^2 + \lambda\sigma^2 \end{aligned} \tag{1.10}$$

Therefore, in our extrapolation step, we simply extrapolate back to where $\lambda = -1$, and Cook and Stefanski (1994) show that the SIMEX extrapolated results are unbiased and consistent.

1.3 Data

Futures contracts have different delivery months, and we choose to hold only the nearest contracts and roll contracts into the next contract one week before the last trading day. For example: in September, we hold the crude oil futures contracts that deliver in October, and roll the October contracts into the November contracts one week before the last trading day of the October contract, usually around 15th of the month. Our futures data is from the Commodity Research Bureau (CRB), which covers all commodity futures with daily prices (including "Open", "High", "Low", and "Settle") and volume series (including "Open Interest" and "Volume") for most contracts. We choose the daily "Settle" prices to get daily returns for all the commodities in the study.

Table 1.1 includes the start and end dates of the daily futures price series, and the trading information of each commodity we study in the paper. These futures all have monthly delivery and are traded heavily in the New York Mercantile Exchange, Inc. (NYMEX).

Daily closing S&P 500 index and quarterly dividend data are from the Standard and Poor online database. The 3-month t-bill rate from Fred with the daily return series calculated by assuming it is held for a day, and 3 days on Monday mornings.

Table 1.1: Commodity Trading Information

Commodity	Code	Last Trading Day	Start	End
Light, Sweet Crude Oil	CL	The third business day prior to the 25th calendar day of the month preceding the contract month.	3/30/1983	12/16/2011
Heating Oil	HO	The last business day of the month preceding the contract month.	11/14/1978	12/16/2011
Henry Hub Natural Gas	NG	The third business day prior to the first calendar day of the contract month.	4/4/1990	12/16/2011
Gasoline Blendstock New York Harbor	RB	The last business day of the month preceding the contract month.	1/1/1985	12/16/2011

Note:

1. All commodities are traded in NYMEX
2. All commodities in this study have delivery every month

Table 1.2: Conditional Betas for Crude Oil Futures of Different Sample Periods

	(1983/04-2011/12)	(1983/04-1990/05)	(1990/06-1997/07)	(1997/08-2004/9)	(2004/10-2011/12)
Observations	750	187	187	188	188
Minimum	-8.11	-6.17	-8.11	-4.11	-2.94
25%	-0.47	-0.30	-0.92	-0.56	-0.18
Median	0.03	-0.00	-0.16	-0.03	0.47
Mean	0.01	-0.05	-0.24	-0.10	0.42
75%	0.60	0.30	0.51	0.41	0.98
Maximum	6.79	2.87	5.15	2.81	6.79
Std Deviation	1.14	0.92	1.39	0.95	1.15
Skewness	-0.59	-2.24	-0.83	-0.72	0.79
Kurtosis	7.59	14.54	6.44	2.63	4.68

1.4 Testing with crude oil futures

Historical crude oil prices have fluctuated with shifts in worldwide demand and supply. The real price fell mostly from 1870-1970 in response to discoveries of oil deposits, then rose sharply in 1970s due to conflicts in the Middle East such as the Yom Kippur war and Iranian revolutions. Prices fluctuated a lot during the 80s and 90s, responding to the combination of the effects of global recessions and Middle East conflicts, before it rose steadily due to the high demand induced by strong growth in the emerging markets. A considerable literature (Kilian 2008) shows oil prices have been associated with the overall economy, and we further document and help to clarify these links in this paper using the CAPM.

Table 1.2 show the summary statistics of the bi-weekly Betas estimated by Equation (1.1). Beta is negative on average for crude oil over the periods from 1983 to 2004, suggesting supply shocks are more dominant, and thus negative risk adjustment was applied to the prices. As the supply of crude oil has been disrupted, or threatened to be disrupted, repeatedly over its history, this result provides some supporting evidence on the effect it has on prices and the overall economy: the price of oil and the economy move in opposite directions in the case dominated by supply shocks. As reference, an alternative explanation that supports an exhaustible resource pricing rule, the Hotelling Rule, tells us that the prices of exhaustible resources will grow at the rate of interest over time in the absence of extraction costs (Hotelling 1931). In apparent contradiction to this theory, most natural resources show flat or downward price trends over the long run. However, if the prices of other exhaustible resources also have negative risk adjustment as crude oil does, and Hotelling's interest rate were adjusted for this negative risk, which would make it close to zero over much of history, there would not be an obvious divergence between basic evidence and theory.

Another issue also worth noting is that beta becomes positive on average over the more recent decade, which can be attributed to a combination of fast economic growth in Asia, and the recent global recessions. Much of the huge variations in estimated betas, ranging from -8.11 to 6.79, are likely due to large sampling error. But large variations in the true but unobservable betas are supported by the frequent changes in demand and supply shocks, and the time path of betas seems broadly consistent with historical events.

As one of the major inputs to the global economy, crude oil prices have long been associated with macroeconomic aggregates such as GDP growth, unemployment, etc. (Darby 1982, Hamilton 1983, 1996, and etc.). It is not our intent to establish the correlation of such aggregates with crude oil prices, but it is worth showing how CAPM beta captures main events in oil history. Figure 1.1 shows the yearly average of *Beta* and some main events of the year. When Saudi Arabia decided not to control its production in 1986 as it did during 1981-1985, the price fell sharply from $27/barrel in 1985 to its lowest price, $12/barrel in 1986, which is captured by the model by a negative beta in 1986 on average, indicating supply shocks. Moreover, as OPEC countries produce more than 40% of the world's crude oil, have almost all the spare production capability, and 2/3 of world oil reserves; wars in the Middle East bring about increased possibility of disruption to oil supply, which leads to higher oil prices and a slowdown of economic activities, and thus a negative beta.

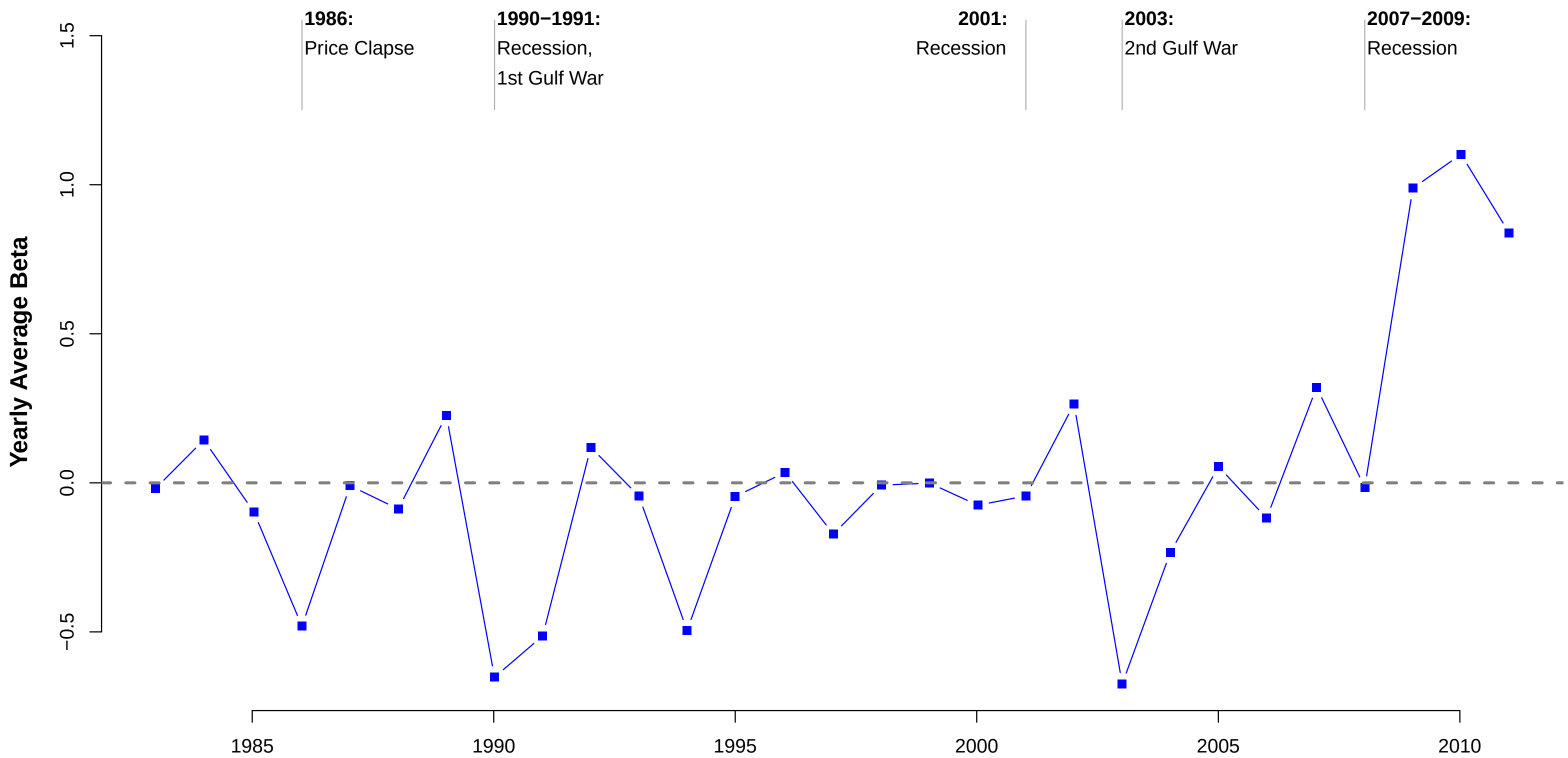

Figure 1.1: Yearly Average Beta and Events of the year

We can also see from figure 1.1 that beta appears to be negative on average during the Gulf War (Aug 2, 1990 - Feb 28, 1991) and 2003 Invasion of Iraq, which is the beginning of the Iraq War (Mar 19 - May 1, 2003). On the other hand, as an industrial commodity, crude oil is also expected to covary positively with the overall economy since strong economic growth will lead to higher demand; therefore, we would see positive betas during recessions and recoveries. This appears to be the case during both the recession in 2001, the most recent recession (Dec 2007 - Jun 2009), and post-recession period after it. In the recession from July 1990 - Mar 1991, however, this is not as apparent. Although the negative beta indicates a supply shock, it can be easily explained by the big supply shock caused by the Gulf War during the same period, which appears to be more dominant. Figure 1.2, where bi-weekly betas are shown instead of yearly average, clearly shows both large positive and negative betas present during that period, providing support for our use of bi-weekly betas, rather than holding beta constant even over a year.

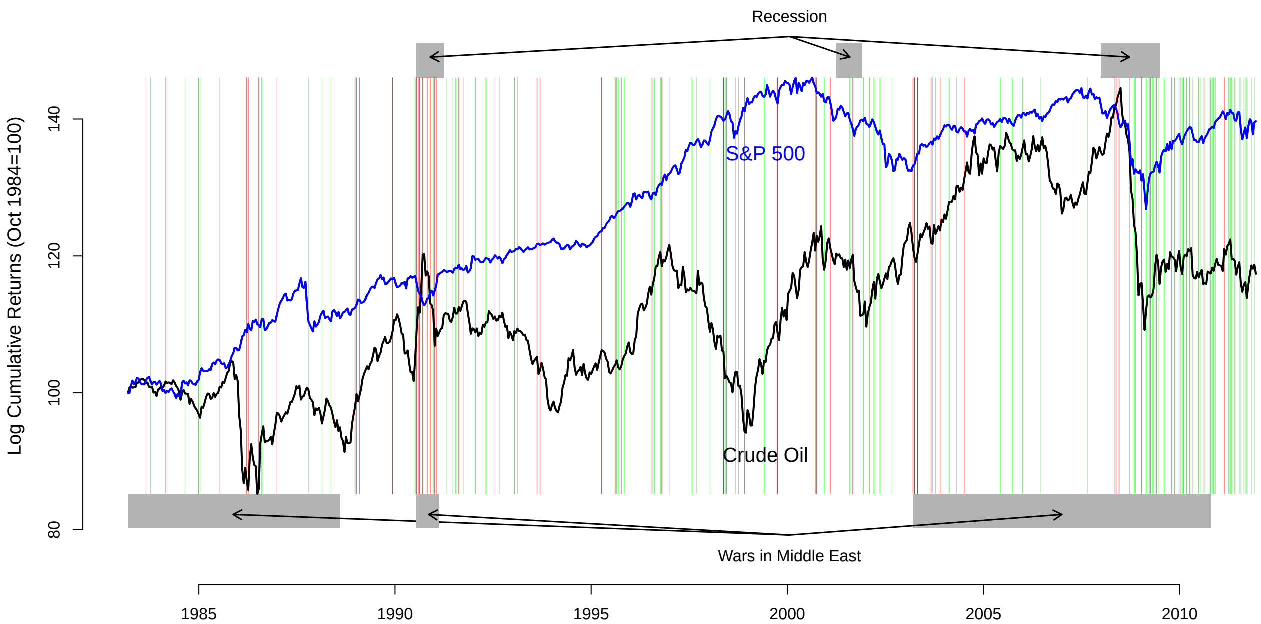

Figure 1.2: Performance of CL and S&P 500 in relation to Beta

Notes: β are shown on graph using vertical lines:

- All betas shown are statistically significant at 90% level
- Red: $\beta \leq -1$; light red: $-1 < \beta < 0$
- Green: $\beta \geq 1$, light green: $0 < \beta < 1$

Figure 1.2 shows the rate of return of Crude Oil futures and the S&P 500, as well as the bi-weekly betas during the period. Market betas included in this figure are all significant at the 10% significance level, i.e. their t values fall in the region of $|t| > 1.833$ [1]. There are 48 negative betas, and 91 positive betas, out of 750 bi-weekly observations. As the graph shows, besides the apparent shocks mentioned above, there are a few other points worth mentioning: 1. Beta changes dramatically from period to period, with a trend toward being positive over the recent years (more green indicating more positive betas). 2. Price rose steadily from 2003 to 2007 due to a combination of the effects of stagnant supply and high demand (Hamilton 2009), which is also supported by the graph where positive and negative betas alternate. It also spiked in mid-2008 (Declining inventory prior to mid-2008), which apparently panicked the market and dragged the economy deeper into recession. 3. The spike of the price of oil in mid-2008 was followed by a big plummet, after which the price began to follow the pace of economic development very closely, as indicated by the positive correlation between oil prices and the economy. Additionally, price return on Oil appears to be higher than the S&P 500 through 2009-2011 where $\beta \geq 1$ for most of the time. [2]

Results shown in Table (1.3) show that the realized excess return of crude oil varies consistently with beta even without error correction; but, we reject the null hypothesis that $\delta_1 = 1$ because of the attenuation bias. When the previous period's estimated beta is used as the instrument, we tend to overestimate the risk , $\delta_1 > 1$, unlike the underestimation of risk in many previous studies. Using a more complex method of state space model, we fail to reject the conditional CAPM.

SIMEX runs with λ chosen to go from 0 to 4 using an increment of 0.04, 100 iterations for each λ. We calculate the sample mean of 100 estimated $\delta' s$ for each chosen λ and form a series of paired δ and λ. Using mean measurement error variance from short-window regression as the proxy of σ^2, which is 0.81 about 62% of the total variation in estimated beta. We choose different fitting methods of the trend and extrapolate to $\lambda = -1$ to get the δ coefficients that are corrected for estimation errors. For fitting methods such as restricted cubic splines and natural splines, linear extrapolation is used given these are locally fitting methods. For polynomial fitting, we use polynomials of the corresponding orders in the extrapolation. In Figure 1.3, the black dots are simulation results for δ_1, and each solid line correspond to one fitting method

[1] 10% critical value of student t with 8 degrees of freedom.

[2]

$$
\begin{aligned}
R_{oil} &= R_f + \beta(R_m - R_f) \\
&\geq R_f + 1(R_m - R_f) \\
&= R_m
\end{aligned}
\tag{1.11}
$$

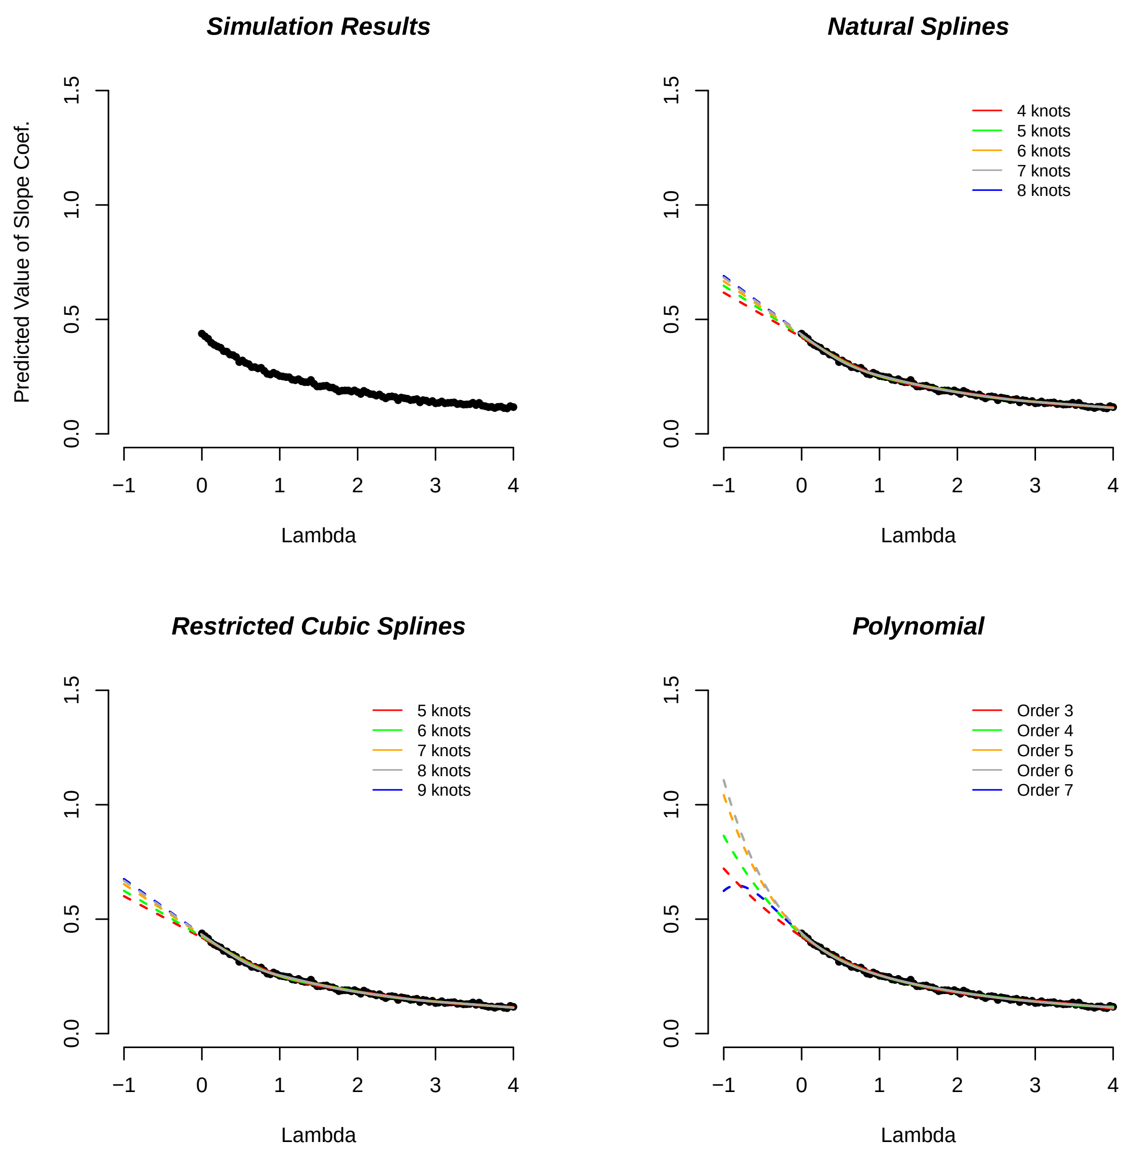

Figure 1.3: SIMEX with Different Fitting Methods

Table 1.3: CAPM Results - Summary

	F-Stat (p-value)	$\delta_0 \times 100$ (s.e $\times 100$)	δ_1 (s.e)
Standard OLS	38.47	0.08	0.44
	(0.00)	(0.23)	(0.07)
IV Corrected	18.11	0.04	1.90
	(0.00)	(0.24)	(0.45)
Kalman Filter	33.40	0.05	1.29
	(0.00)	(0.24)	(0.22)

Notes: The tables shows the regression results of Equation 1.2: $r_{i\tau} = \delta_0 + \delta_1(\hat{\beta}_{i\tau} r_{M\tau}) + \mu_{i\tau}$

1. Standard OLS: βs are estimated results of short-window regressions (Equation 1.1) without error correction
2. IV Corrected: βs are corrected for using its lag as instrument
3. Kalman Filted: βs are filtered using state space model with Kalman filter

with the extrapolation part shown as dotted line. The larger the error is, the more the results are biased towards zero; and there is a uptick at around $\lambda = 0$. The linear extrapolation does not follow the trend as closely as the polynomial; and we don't reject CAPM in most polynomial fitting we tried (shown in Table 1.4).

Table 1.4: Simulation Extrapolation Results using Different Fitting Methods

	Natural Spline*		Restricted Cubic Spline**		Polynomial	
df	Intercept(×100) (s.e.×100)	Slope (s.e.)	Intercept(×100) (s.e.×100)	Slope (s.e.)	Intercept(×100) (s.e. ×100)	Slope (s.e.)
3	0.10	0.62	0.09	0.60	0.10	0.72
	(0.23)	(0.09)	(0.23)	(0.09)	(0.23)	(0.11)
4	0.10	0.65	0.10	0.62	0.13	0.86
	(0.23)	(0.10)	(0.23)	(0.09)	(0.23)	(0.13)
5	0.10	0.67	0.10	0.65	0.17	1.04
	(0.23)	(0.10)	(0.23)	(0.10)	(0.23)	(0.15)
6	0.11	0.68	0.10	0.67	0.10	1.11
	(0.23)	(0.11)	(0.23)	(0.10)	(0.23)	(0.19)
7	0.11	0.69	0.10	0.67	0.19	0.62
	(0.23)	(0.11)	(0.23)	(0.11)	(0.26)	(0.28)

Notes:

1. * In Natural Spline, Knots = df +1
2. ** In Restricted Cubic Spline, Knots = df +2
3. Linear Extrapolation is used in Natural Spline and Restricted Cubic Spline for data beyond the original range.

Table 1.5: Summary of Pricing Error (%)

	Crude Oil	Heating Oil	Natural Gas	Gasoline
Observations	750	864	567	706
Minimum	-2.88	-2.84	-3.98	-2.94
25%	-0.38	-0.38	-0.72	-0.42
Median	0.01	-0.01	-0.12	0.07
Mean	0.01	0.02	-0.14	0.03
75%	0.40	0.39	0.48	0.46
Maximum	3.48	3.10	4.07	3.22
Std Deviation	0.70	0.70	1.04	0.73
Skewness	0.08	0.12	-0.19	-0.08
Kurtosis	2.49	2.04	1.51	1.72

1.5 Testing other Energy Resources

In this section, we test CAPM using other energy commodities including unleaded gasoline, heating oil, and natural gas. [3]

Pricing errors for each of these commodities are small and statistically insignificant as suggested by Table 1.5. Since alphas are estimated by using short-window regressions, the significance of each bi-weekly alpha will also show us how accurately CAPM can predict. Figure 1.4 includes the frequency distributions of the t-values for the estimated bi-weekly alphas for each commodity, and the blue line is the student t distribution with 8 degrees of freedom as a reference[4]. If the red histogram bars can fill in the area under the blue line, it indicates that the t-values of bi-weekly alpha follows the t-distribute with zero mean. The figure provides us with evidence that the realized distribution of the pricing error fit the student t distribution well for all commodities, and no commodity has more high/low return than is predicted by the CAPM. [5]

[3]Together with Crude Oil, these four commodities are the energy commodity futures closely monitored by the Commodity Futures Trading Commission (CFTC).

[4]Bi-weekly alphas are estimated using 10 observations in each calendar week.

[5]Jensen (1968) used a similar technique and shows that portfolio managers don't have the extra predictive power to outperform the market prediction. However, he shows the frequency distribution of 115 cross-section portfolio, while we examine a time-series instead.

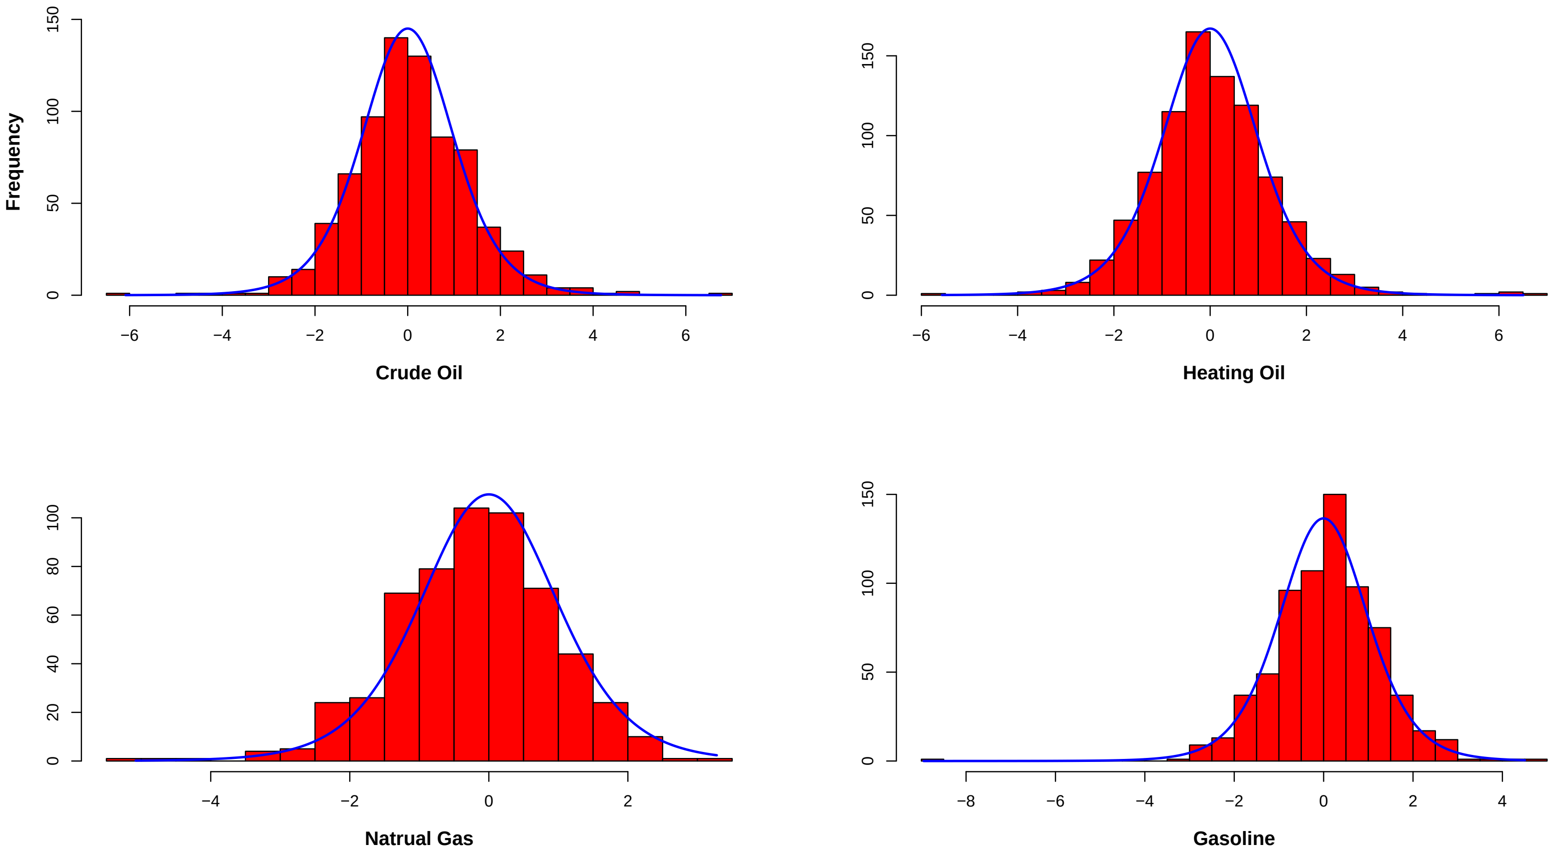

Figure 1.4: T-values of Estimated Pricing Error

Notes:

Blue line is student t distribution with 8 degrees of freedom.

Table 1.6: Summary of Market Betas

	Crude Oil	Heating Oil	Natural Gas	Gasoline
Observations	750	864	567	706
Minimum	-8.11	-8.22	-4.55	-6.97
25%	-0.47	-0.47	-0.66	-0.49
Median	0.03	0.07	0.02	0.09
Mean	0.01	0.01	0.08	0.05
75%	0.60	0.53	0.67	0.62
Maximum	6.79	6.16	12.32	7.29
Std Deviation	1.14	1.10	1.58	1.14
Skewness	-0.59	-0.78	1.55	-0.06
Kurtosis	7.59	9.52	10.20	5.58

Beta (Table 1.6) is on average positive for natural gas (1990 - 2011), gasoline (1984 - 2011) , and heating oil (1978 - 2011), over their respective sampling periods with significant level of variation. The frequency distribution of the t-values of estimated betas is shown in Figure 1.5, and indicates much fatter tails for all the commodities than are suggested by the underlying distribution. The significant number of large positive and negative betas is not only from the tail of the distribution, unlike pricing error as shown earlier, which came largely from the extremes of the distribution. Significantly different from zero, CAPM captures many negative and positive beta events. Had we held beta constant over time, these significantly positive and negative betas would cancel the effects of each other, and the market beta and the risk premium will be much less significant.

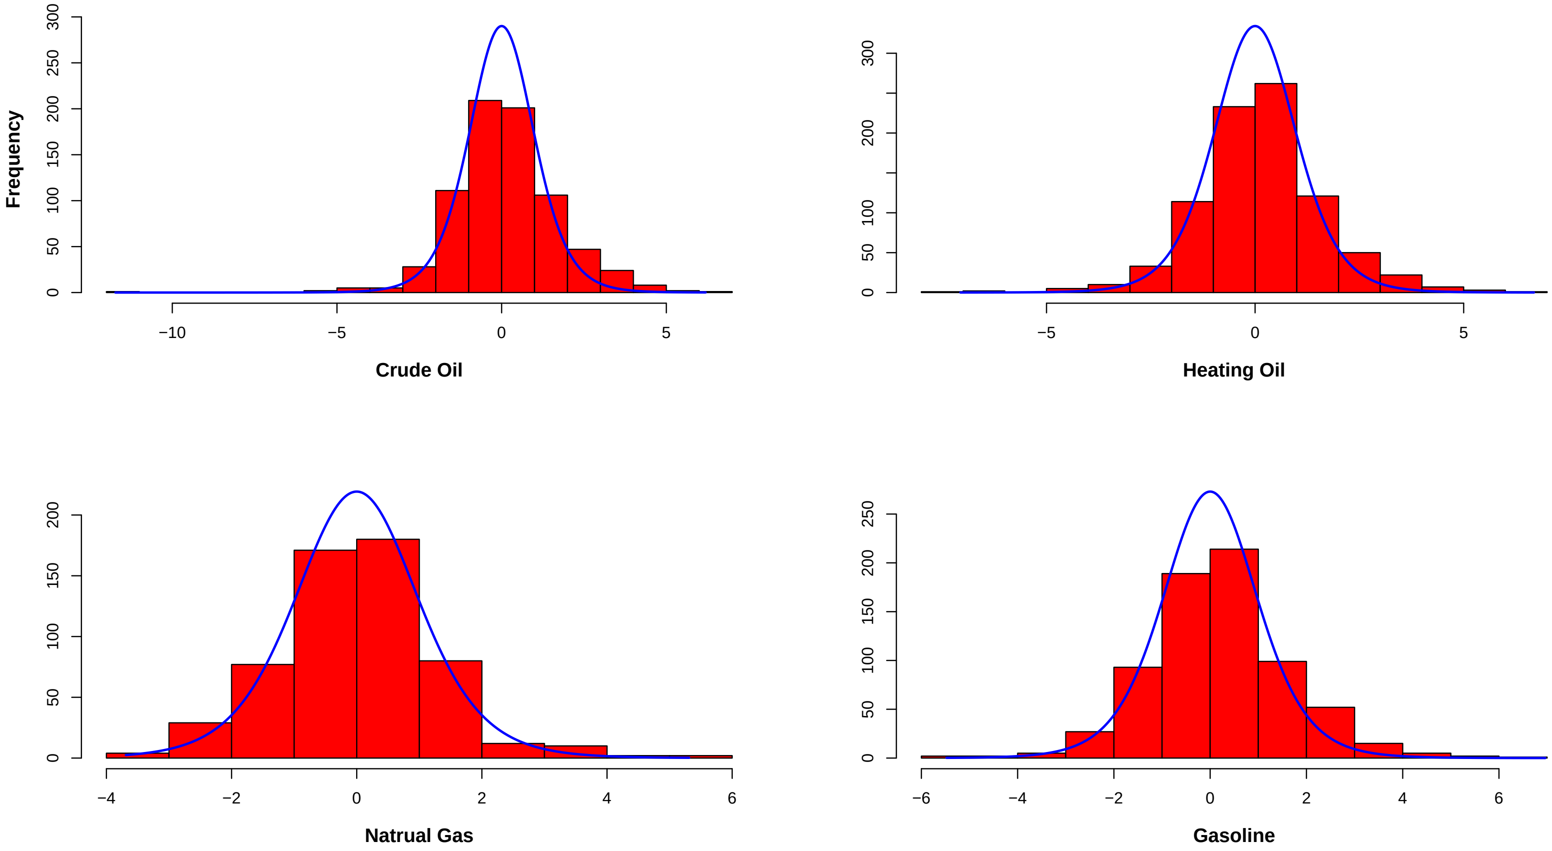

Figure 1.5: T-values of Estimated Market Betas

Notes:

Blue line is student t distribution with 8 degrees of freedom.

As petroleum products, prices of heating oil and gasoline are mostly affected by the price of crude oil, which is determined by its worldwide demand and supply. Figure 1.6 shows the performance of all four energy futures over time; it is clear that prices of heating oil, gasoline and crude oil are highly correlated, while natural gas has a quite different path. Being produced from crude oil, the supply of heating oil and gasoline is greatly affected by oil supply. Supply shocks or storage demand shocks in the crude oil market, such as those resulting from unrest in the Middle East, will also appear as negative beta events for these markets, and similar patterns of negative beta events are shown in Figure 1.7 for heating oil and gasoline. On the other hand, because about 70% of crude oil is used in the transportation sector, demand for crude oil comes in large part from demand for gasoline, which is influenced by growth of the overall economy. Therefore demand shocks such as a recession tend to have bigger effects (positive betas) on gasoline than on crude oil. Despite the frequent supply shocks, gasoline appears to have a positive beta on average as shown in Table 1.2. Heating oil, a heating option in winter, is used by about 7% of Americans and mostly in Northeastern states. The seasonality of the product may reduce demand effects, and although big positive betas can arise with extreme cold weather and recessions, these appear to be outweighed by supply shocks from the oil market; heating oil has a negative beta on average.

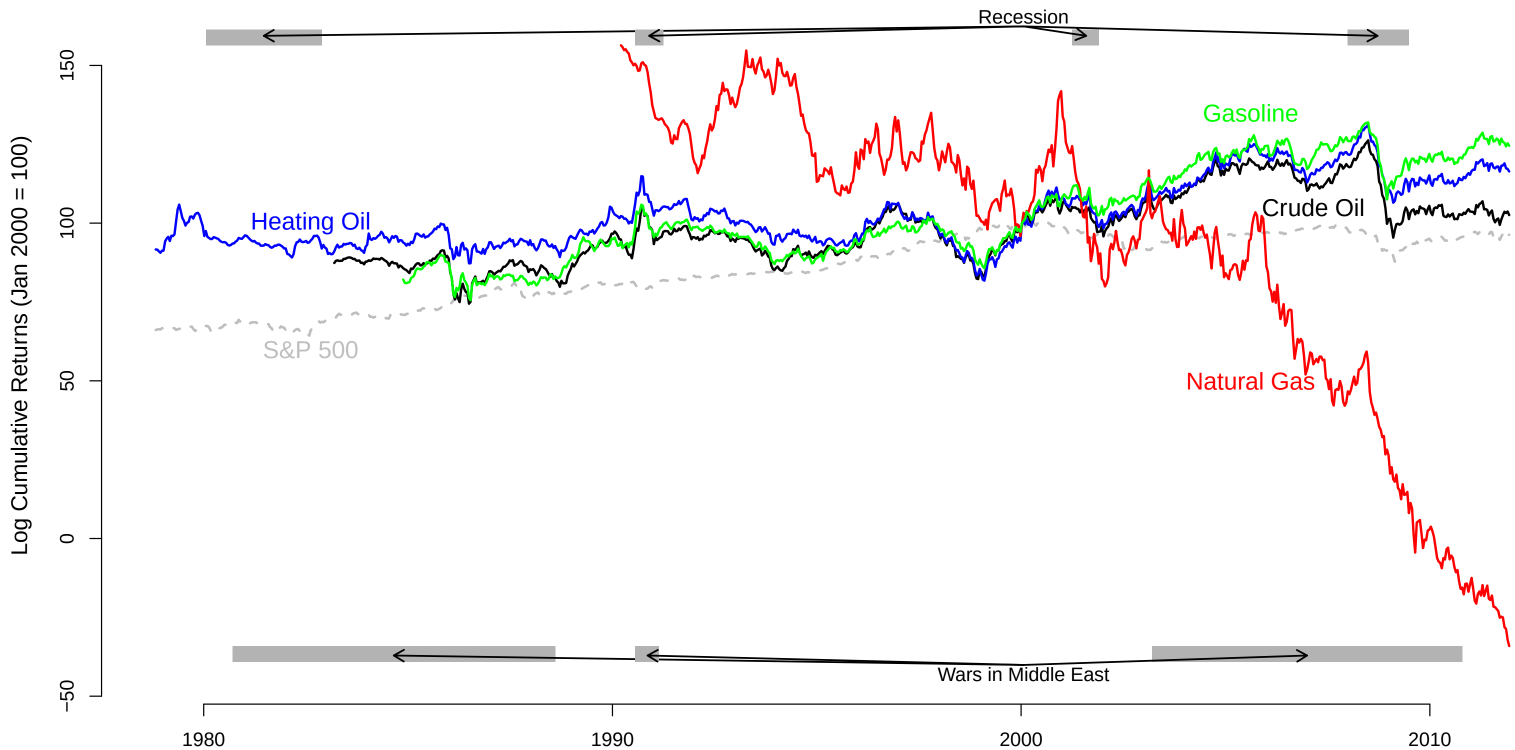

Figure 1.6: Performance of all Commodity Futures and S&P 500

Natural Gas has a different risk profile from crude oil, despite the fact that it is also an exhaustible natural resource like crude oil. Locally supplied natural gas isolates it from disruptions in supply caused by conflicts in the Middle East, but price spikes in crude oil can still lead to higher natural gas prices due to substitution effects, as we can see in the recent price spike in mid-2008. Recent technological advances have enabled us to extract natural gas from shale rock, which greatly affected production. The recent decrease in gas prices from 2007-2011 is likely caused by a combination of recession (demand) and technology (supply) as shown in Figure 1.7, while petroleum products were likely driven mainly by demand fluctuations. Since natural gas is used to heat more than half of the homes in America, short-term demand for natural gas can also be affected greatly by cold temperatures, making both its prices and its market beta fluctuate. Big short-term demand shocks induced by weather lead to a high occurrence of positive betas, and the market beta for natural gas is positive on average.

The test results of conditional CAPM show that realized risk premium varies consistently with beta risk across all commodities of interest, and we fail to reject CAPM when measurement error is corrected for using SIMEX and the Kalman filter. Graphic representation of test results before and after error correction is shown in Figure(1.8), where each commodity has its risk premium grouped into 30 quantiles based on level of CAPM prediction. Realized risk premium varies positively with its CAPM prediction before and after correction; meanwhile, trend is flatter before correction due to attenuation bias. The slope appears to be slightly steeper than perfect CAPM prediction when measurement error is corrected with Kalman filter, indicating that realized risk premium increases more rapidly than one-for-one with CAPM-based risk, contrary to the relatively flat relationship found in earlier CAPM studies.

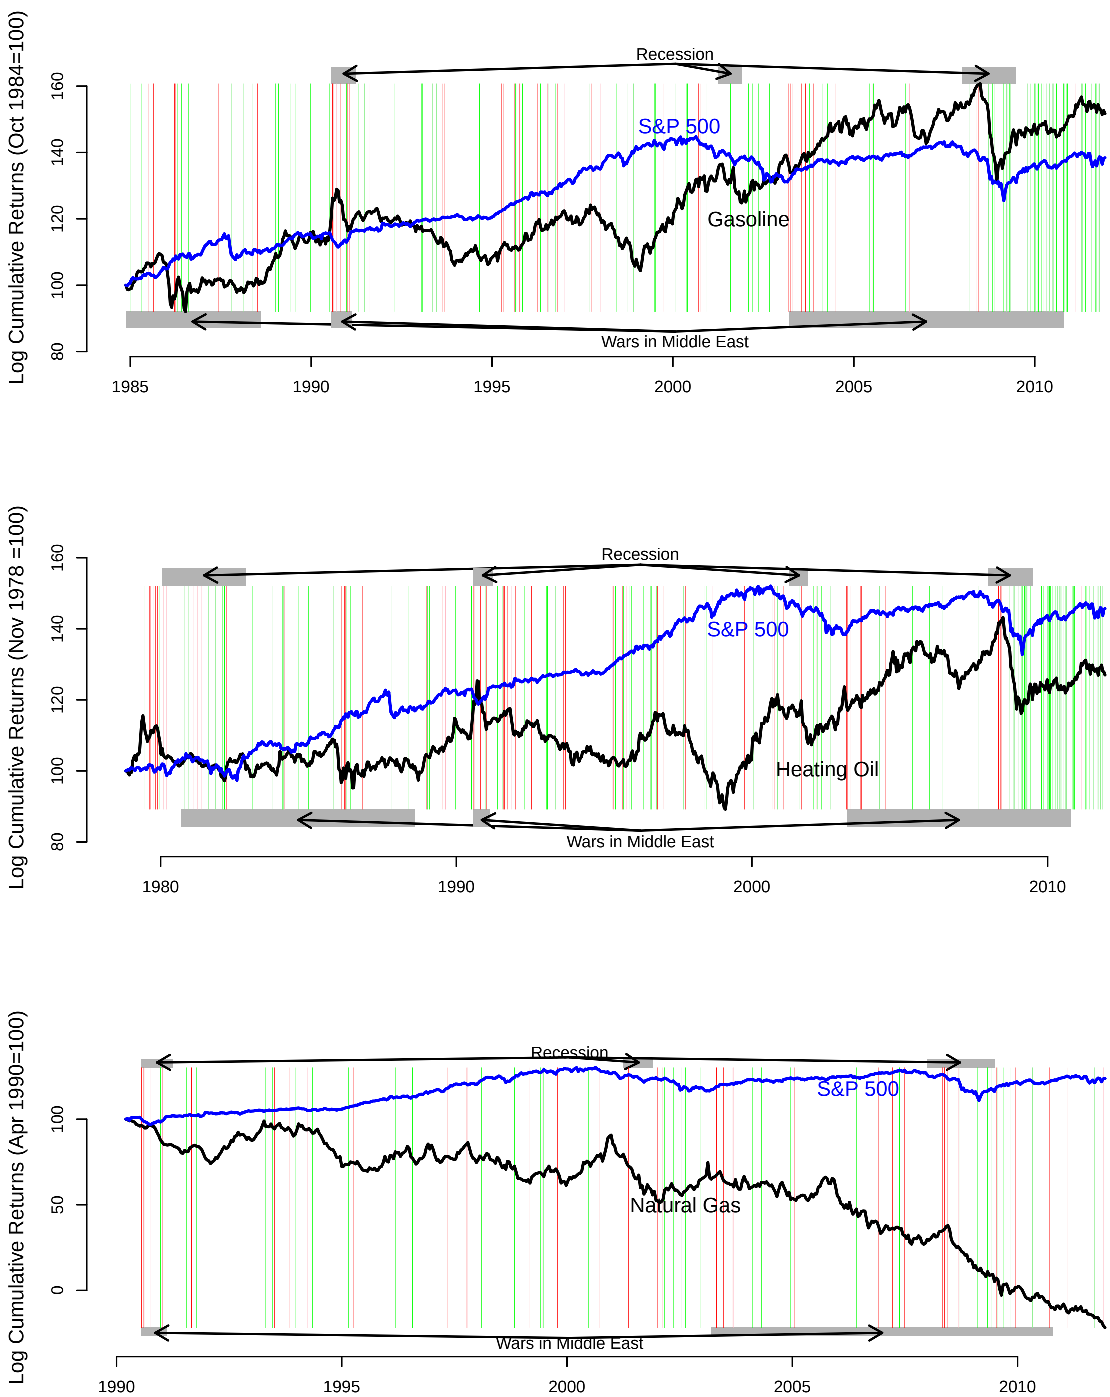

Figure 1.7: Performance of Energy Commodities and their *betas*

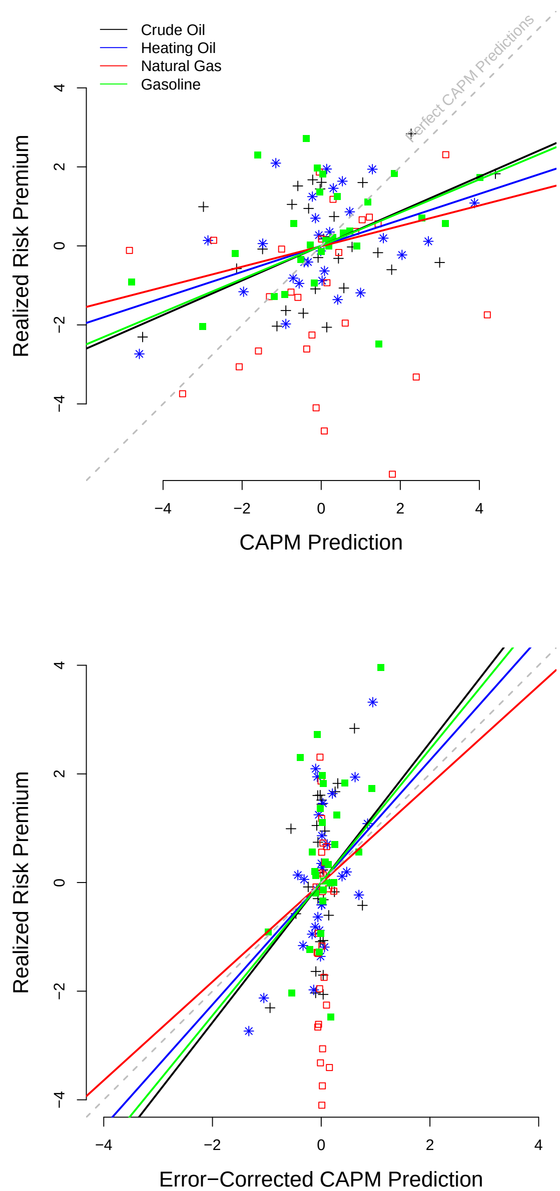

Figure 1.8: Risk Premium (% per 2-week period)

Table 1.7: CAPM Results - Summary

Commodity	**Standard** $\delta_0 \times 100$ (s.e $\times$100)	**Standard** δ_1 (s.e)	**IV Corrected** $\delta_0 \times 100$ (s.e $\times$100)	**IV Corrected** δ_1 (s.e)	**Filtered** $\delta_0 \times 100$ (s.e $\times$100)	**Filtered** δ_1 (s.e)
	CAPM Predictions: $\delta_0 = 0$ **and** $\delta_1 = 1$					
Crude Oil	0.08	0.44	0.04	1.90	0.05	1.29
	(0.23)	(0.07)	(0.24)	(0.45)	(0.24)	(0.22)
Heating Oil	0.11	0.33	0.07	2.04	0.07	1.12
	(0.22)	(0.07)	(0.22)	(0.72)	(0.22)	(0.21)
Natural Gas	-1.08	0.26	-1.05	2.44	-1.04	0.91
	(0.39)	(0.09)	(0.39)	(1.59)	(0.39)	(1.51)
Gasoline	0.30	0.42	0.30	2.84	0.27	1.23
	(0.25)	(0.08)	(0.25)	(0.66)	(0.25)	(0.23)

Notes: The tables shows the regression results of Equation 1.2: $r_{i\tau} = \delta_0 + \delta_1(\hat{\beta}_{i\tau} r_{M\tau}) + \mu_{i\tau}$

1. Standard: βs are estimated results of short-window regressions (Equation 1.1) without error correction
2. IV Corrected: βs are corrected for using its lag as instrument
3. Filtered: βs are filtered using state space model with Kalman filter

Table 1.8: SIMEX results for Energy Commodities

	Natural Spline*		Restricted Cubic Spline**		Polynomial***	
Commodity	Intercept(×100) (s.e.×100)	Slope (s.e.)	Intercept(×100) (s.e.×100)	Slope (s.e.)	Intercept(×100) (s.e. ×100)	Slope (s.e.)
Crude Oil	0.10	0.65	0.10	0.62	0.11	0.85
	(0.23)	(0.10)	(0.23)	(0.09)	(0.24)	(0.13)
Heating Oil	0.13	0.51	0.12	0.49	0.14	0.71
	(0.22)	(0.10)	(0.22)	(0.10)	(0.22)	(0.14)
Natural Gas	-1.12	0.42	-1.11	0.40	-1.17	0.65
	(0.39)	(0.12)	(0.39)	(0.12)	(0.39)	(0.17)
Gasoline	0.29	0.66	0.29	0.63	0.29	0.93
	(0.25)	(0.12)	(0.25)	(0.11)	(0.25)	(0.17)

Notes:

1. * Natural Spline of 5 knots
2. ** Restricted Cubic Spline of 6 knots
3. *** Polynomial of order 4

1.6 Conclusion

Using high frequency energy futures data, we estimate a series of bi-weekly betas using CAPM and document a strong relationship between energy futures prices and the aggregate economy. Unlike previous studies, where betas for stocks are mostly close to 1 and zero for commodities, our model shows great variability in the betas for a single commodity over time that vary from strongly positive to strongly negative, and significant differences in betas across different commodities, reflecting fundamental differences in the nature of supply and demand and connection to the macro economy.

The market beta displays a significant amount of price variation in the energy market and we fail to reject conditional CAPM, correcting for measurement error. Realized risk premium, however, appears to increase more rapidly than one-to-one with CAPM-based risk, after measurement error is corrected with IV. An area for further investigation is to explain why there appears to be too much covariance risk pricing, rather than too little as previous studies indicate.

For crude oil, we are able to disentangle the different shocks in the market and show how oil prices and the aggregate economy affect each other; similar findings are presented in Kilian (2009) who studies this issue using a structural VAR model. Beta and risk premium in the oil market appear to be large and highly uncertain, which gives rise to a highly uncertain discount rate. Weitzman (2001) proposed gamma discounting where a smaller rate should be used in the distant future in the presence of uncertainty, and Newell and Pizer (2003) quantifies this effect using US interest rate. While risk premium is not part of the discount rate in Newell and Pizer's analysis due to lack of data, it would be nice to extend and test their findings by incorporating the risk premium that we estimated in the energy market.

Chapter 2

Discounting Investments that Save Energy Resources

2.1 Introduction

Relative to traditional technologies, investments in renewable energy and energy-conserving technologies involve higher up-front costs in exchange for a lower stream of future operating costs. The discount rate for future expected operating costs is therefore central to the investment decision.

Recent research has emphasized the importance of uncertainty about the discount rate, especially when investments are long-lived. Because present values are convex in the discount rate, uncertainty about the rate combined with Jensen's inequality implies that optimal investment decisions are governed by a relatively low certainty-equivalent rate that declines with the investment's duration (Weitzman 1998, Weitzman 2001, Newell and Pizer 2003). This past research has emphasized discounting public investments to curb global warming, given the especially long time horizon and vigorous debate (and thus uncertainty) about which discount rate should be used (Stern 2006, Tol 2006, Nordhaus 2007, Dasgupta and Maskin 2005, Weitzman 2007 etc). But uncertainty about future discount rates could be important to many kinds of investments, public and private. And while more tangible investments in renewable energy or energy-conserving technologies surely have shorter durations than broader public policies to curb global warming, the horizons are long enough for uncertainty to matter.

This article considers empirical measurement of the uncertainty-adjusted discount rate for application to energy-saving investments. Examples of such investments include windmills, solar energy infrastructure, hybrid and electric cars, and more energy efficient buildings. While such investments embody many tradeoffs besides up-front costs versus future operating costs, this tradeoff is arguably the most essential. To our knowledge, the market price of this tradeoff has not been systematically measured accounting for stochastic variations in interest rates, risk premiums and the operating costs themselves. Our central contention is that careful account of these uncertainties substantially influences optimal investment decisions.

Modeling discount rate uncertainty requires systematic decomposition of the size and stochastic evolution of its components. A natural starting point for the analysis follows Newell and Pizer (2003), who estimated long-run discount rate uncertainty using time-series analysis of historic real interest rates. In addition to stochastic variations in interest rates, in this paper we simultaneously account for (a) uncertainty about energy prices; (b) uncertainty about overall-market risk premiums; and (c) uncertain and time-varying risk premia of energy savings.

Risk premiums for energy-related investments are particularly important given connections between oil prices and the macroeconomy (Hamilton 1983, Barsky and Kilian 2004, Kilian 2009). Historically, price spikes in oil and other energy resources have been associated with supply shocks or with speculative demand driven by fear of future supply shocks. Such shocks are classic examples of negative aggregate supply shocks for the macroeconomy, and should thus have negative risk premiums. But not all oil prices shocks are alike (Kilian 2009). More recent spikes in oil prices can be connected to positive aggregate demand shocks, particularly growth in Asia, which would logically have positive risk premiums (Hamilton and Wu 2011, Kilian 2009). Thus, events that drive energy price variations suggest that underlying risk premiums may be large, positive or negative, and time varying, giving rise to considerable uncertainty about future risk-adjusted discount rates of energy-conserving investments. In quantifying all of these uncertainties simultaneously, we essentially attempt to rigorously quantify the value of so-called "energy security."

Following the Capital Asset Pricing Model (CAPM), the risk premium equals an overall-market risk premium multiplied by the investment-specific *beta*. The *beta* measures the covariance between returns of a specific investment and returns of the overall market. Whereas the overall-market risk premium is always positive, its size is also time varying, particularly with macroeconomic business cycles. *Beta*, per the discussion above, can also be time-varying, but may be positive or negative. For example, an unexpected discovery of a large new oil deposit would cause oil prices to decline while simultaneously providing good news for the aggregate economy and stock returns. Such an event would be embodied by a negative *beta* and risk premium. Conversely, larger-than-expected growth in emerging markets may drive up aggregate demand, increasing both overall-market returns and oil prices, giving rise to a positive *beta* and risk premium. The discount rate is comprised of the investment-specific risk premium plus a risk-free (or zero-*beta*) interest rate. The risk-free rate can also vary over time, and will be uncertain over long horizons, as modeled by Newell and Pizer.

Our objective here is to measure all of these components of the stochastic discount rate simultaneously, and use that model to derive inferences about future risk-adjusted rates and uncertainty about those rates. Given the stochastic nature of the discount rate, we derive constant-equivalent discount rates for different time horizons for comparisons and discussions. The constant-equivalent discount rate is defined in this paper as the constant discount rate that will produce the same expected present value of the future net benefit assuming energy prices and consumption stays at current level. These constant-equivalent rates may serve as useful guides for energy-related investment decisions.

We begin by modeling a simple stylized model of a consumer facing a choice between two cars with identical features except fuel efficiency. This simplified example helps set up the connection between the expected present value of the future net benefits and the constant-equivalent discount rate. In the absence of any private value attached to the social benefits of fuel economy on the environment, the consumer chooses between the higher up-front cost of a fuel-efficient car and the higher stream of energy cost of the standard car. The present value of the total savings on future energy spending is central to this choice, which depends on both energy prices and discount rates. For this simplified example, we assume all other attributes of the car, like performance, utility, and aesthetic appeal, are held constant.

We first estimate series of bi-weekly energy-specific *betas* estimated using a time-varying CAPM. Because estimated betas contain large sample errors, we then use a state-space model to filter the error and estimate the stochastic evolution of the underlying true *beta*. We model the energy prices, the market risk-premium and the risk-free rate with a vector autoregressive model (VAR) to allow the variables to covary with each other and account for heteroskedasticity using a generalized autoregressive conditional heteroskedasticity (GARCH) model. After estimating these models using past data of energy prices, we forecast future values and uncertainties using Monte Carlo simulation. We use elasticities from the existing literature to account for adjustments in vehicle miles traveled and associated surplus changes in response to projected energy prices. To estimate the expected present value of the net savings on future energy costs, we integrate the distribution of risk-adjusted savings by simulating the estimated model 1,000 times, calculating the realized present value of savings in each iteration, and taking the average.

We find that the conditional CAPM has strong predictive power over the risk premia in the oil market and fail to reject the model when estimation errors in *beta* are corrected for using the Kalman Filter. By allowing *beta* to vary at relatively high frequency, the time path of filtered *beta* captures many of the historical shocks in the oil market. Over the time period from 1983-2011, crude oil has had an average -0.37% annual risk premium, assuming an overall market risk premium of 5%, but it varied significantly from positive to negative over time in response to the market changes (such as recessions and booms), as well as shifts in oil supply (such as those caused by unrest in the Middle East).

Different model specifications have different levels of uncertainty, where models that assume energy *betas* follow a random walk have larger variations in predicted values than those that assume *betas* follow a stationary autoregressive (AR) process. We find that the greater the uncertainty, the higher the present value of total energy saving, and thus the lower the certainty-equivalent discount rate. Few investments actually save oil, but rather gasoline or electricity or some composite of energy resources, although prices of substitute energy resources tend to be correlated. Further, consumption behavior would presumably change in response to prices. People will drive less if gasoline prices rise. Taking these into account for a vehicle, our estimate of an implicit discount rate ranges from 18% in the most conservative model setup, to -20% in the most uncertain for the 5-year time horizon.

When we expand the calculations to cover the lifetimes of other, longer-lived energy-related investments, we find that the more distant into the future, the lower the discount rate for each model specification due to the rising uncertainty with time. With the most conservative method, where *beta* follows an AR(1) and other variables follow VAR(1) with no trend, the annualized discount rate falls from 18% in 5 years to 6% in 25 years, 3.7% in 50 years, and 2% in 100 years. Whereas adding a linear trend to the VAR model and keeping other specifications unchanged, the annualized discount rate becomes 16% in 5 years, 3% in 25 years to -3.62% in 100 years. To further illustrate how high levels of uncertainty may affect the constant-equivalent discount rate, we include models with quadratic time trend which have the higher level of uncertainty; and its predicted discount rates are much lower for all different lifetimes of projects -14%, -19% and -39% respectively. The qualitative results fit the declining pattern both Weitzman (1998) and Newell and Pizer (2003) found, where Newell and Pizer conclude that certainty-equivalent discount rate have declined from 4% to 2% after 100 years, using interest rate as the single component in the discount rate and ignoring the risk component.

2.2 A Simple Numerical Example

In order to understand how energy prices, energy consumption and project durations would affect the choice of the discount rate, we set up a simple numerical example to simplify and focus on the connection between the present value of the net benefit and the constant-equivalent discount rate. We start by considering a consumer choosing between a Honda Civic Hybrid and standard Honda Civic. Assume that these two cars have identical features, except that the hybrid version has a higher price, but uses less gasoline. Suppose also that the consumer cares only about the expected present value of expenditures, and has no intrinsic preferences regarding minimizing the environmental consequences of fuel consumption. The decision then boils down to a comparison between the hybrid's extra up-front cost and the expected present value of its gasoline savings.

The present value of the total energy savings over a N-period horizon is equal to

$$S = P_1 Q_1^s exp\left(-\delta_1\right) + P_2 Q_2^s exp\left(-\delta_1 - \delta_2\right) + ... + P_N Q_N^s exp\left(-\sum_{t=1}^{N} \delta_t\right). \tag{2.1}$$

where P_t is the price of energy, Q_t^s is the amount of fuel that consumer can save by driving the hybrid rather than the standard, and δ_t is the risk adjusted discount rate, defined as

$$\delta_t = r_t + \gamma_t, \tag{2.2}$$

with r_t denoting risk-free interest rate and γ_t the risk premium.

An essential consideration is the fact that prices P_t and future discount rates δ_t are not known to the consumer in advance. Further, the consumer's expectations about P_t and δ_t may well be correlated. The law of one price rules that the price of an asset equals to its expected discounted payoffs.

$$p_t = E_t\left[\Delta_{t+1} x_{t+1}\right]$$

where Δ_{t+1} is the intertemporal marginal rate of substitution, or the stochastic discount factor, and x_{t+1} is the asset's payoff.

In the consumption-based asset pricing models, the stochastic discount factor is often defined as

$$\Delta_{t+1} = \rho \frac{u'(c_{t+1})}{u'(c_t)}$$

where ρ is the subjective rate of time preference, and $u'(c)$ is the marginal utility of consumption.

To further illustrate how prices and discount rate may be correlated, from the first equation above we can get:

$$\begin{aligned} 1 &= E_t\left(\frac{x_{t+1}}{p_t}\Delta_{t+1}\right) \\ &= E_t\left[(R_{t+1})\Delta_{t+1}\right] \\ &= E_t(R_{t+1})E_t(\Delta_{t+1}) + Cov_t(R_{t+1}, \Delta_{t+1}) \end{aligned} \tag{2.3}$$

where R_{t+1} is the gross return of the asset. Equation (2.3) can be rewritten as:

$$E_t(R_{t+1}) = \frac{1 - Cov_t(R_{t+1}, \Delta_{t+1})}{E_t(\Delta_{t+1})} \tag{2.4}$$

An asset has a higher expected return, when the covariance between its return and the stochastic discount factor is low. For example, the price of an asset that is positively correlated with consumption would be negatively correlated with the intertemporal marginal rate of consumption substitution, i.e. the stochastic discount factor, and thus earn a positive risk premium and have a higher expected return. Such an asset typically has a low return when investors' consumption levels are low (marginal utility is high) and the asset is perceived as risky because it does not provide wealth when it is needed most. Therefore, investors will demand higher premia to compensate for the risk they bear holding such assets.

Our general form of the net present value of total energy saving (Equation 2.1) can fully embody such correlation and provide us with better understanding of how savings would change given uncertain prices, consumption levels, discount rates and their potential correlations. The next step is to model how these variables have evolved historically using time series econometric models. We will then forecast from the estimated models to quantify trends and uncertainty in P_t and δ_t.

2.3 Estimating Risk Premium using the Capital Asset Pricing Model

The Capital Asset Pricing Model (CAPM) ties the scholastic discount factor to the return on the wealth/market portfolio and is most often expressed in its beta-representation. It states that the risk premium equals an overall-market risk premium multiplied by the investment-specific *beta*, which measures the covariance between the returns of a specific investment and returns of the overall market. Similar to the covariance between prices and the marginal rate of consumption substitution in Equation (2.4), *beta* is also a good measure of risk and an important component of the risk premium. For example, an asset that co-varies negatively with the overall market would become a good hedge against bad economic times protecting consumers like an insurance. Thus, such asset would be perceived as less risky and warrant a smaller risk premium.

Following the CAPM, we can rewrite Equation (2.2) as:

$$\delta_t = r_t + \beta_t \gamma_t^M, \tag{2.5}$$

where β_t is the asset-specific beta and γ_t^M is the market risk premium. The overall-market risk premium is always positive and its size generally varies with time, particularly with macroeconomic factors such as unemployment and growth. Markets become riskier in bad economic times such as recessions, therefore investors demand higher risk premia to compensate for the extra risk they bear. And the poorer the state of the economy is, the higher is the risk premium demanded by the investors.[1]

The asset-specific *beta* can also vary over time, but can be either positive or negative, especially for energy-saving investments. For example, an unexpected shift in supply, like discovery

[1]See Chen (1991), Keim and Stambaugh (1986), Chen, Roll and Ross (1986)

of a large new oil deposit, would cause oil prices to decline, while simultaneously providing good news for the aggregate economy and stock returns. Such an event would be embodied by a negative beta and risk premium. Conversely, the repeatedly faster-than-expected growth in Asia may drive up aggregate demand, increasing both overall-market returns and oil prices, giving rise to a positive beta and risk premium. Similarly in the recent work by Kilian (2009), he attributes the recent rise in oil price mostly to the positive demand shock. The relative frequencies of such events may change over time, sometimes quite frequently, causing large variations in *beta* that cannot be captured by the static CAPM many previous studies have adopted.[2]

To allow betas to change with high frequency, we follow Lewellen and Nagel (2006) and directly estimate betas and risk premia unconditionally through short-window (bi-weekly) regressions to generate an estimate for *beta* in each two-week period t:

$$\gamma_i = \alpha + \beta\gamma_i^M + \nu_i \qquad (2.6)$$
$$i = 1, \cdots, 10$$

where γ_i and γ_i^M are asset risk premia and market risk premium on the i^{th} day within each time period, respectively.

Short-window regressions give unbiased estimates of the average *beta* within each window. A short window therefore allows for relatively unconstrained changes in *beta* over time. The tradeoff is that the small sample size within each window results in large estimation error. Below, we will consider state-space methods for filtering this error to uncover the underlying process for the true *beta*.

The state space model is a flexible form that can be formulated to model dynamic time series with measurement errors. The state space representation of short-window betas can be written as:

$$\hat{\beta}_t = \beta_t + \epsilon_t \qquad (2.7)$$

$$\beta_t = \phi_0 + \phi_1\beta_{t-1} + \upsilon_t \qquad (2.8)$$

where $\epsilon_t \sim \mathcal{N}(0, V)$ and $\upsilon_t \sim \mathcal{N}(0, W)$.

Equation (2.7) states that the $\hat{\beta}_t$ estimated using the short window regressions is composed of the true *beta*, β_t, and the estimation error, ϵ_t. Despite the fact that *beta* may change significantly over time, the nature of uncertainty could likely be similar enough from week to week and there is some level of persistence. The transition equation (2.8) aims to capture the persistence of *beta* over time and assumes that the true *beta* follows an autoregressive process.

In the presence of heteroskedastic disturbances, the statistical inferences on estimates from

[2]See Black, Jensen and Scholes (1972), Blume and Friend (1973), Fama and MacBeth (1973), and many others

the basic space state model will be inefficient. In this paper, we adopt the method derived in Harvey, Ruiz and Sentana (1992) and Kim and Nelson (1999), and combine the state space model with the Generalized Autoregressive Conditional Heteroskedasticity (GARCH) process to model the disturbances. That is, if the distributions of both ϵ_t and υ_t depend on the past information up to $t-1$, they can be modeled by the GARCH process.

In order to do this, we define the conditional distributions of disturbances as $\epsilon_t|\psi_{t-1} \sim \mathcal{N}(0, h_{1,t})$ and $\upsilon_t|\psi_{t-1} \sim \mathcal{N}(0, h_{2,t})$, where ψ_{t-1} refers to the information set up to $t-1$. $h_{1,t}$ and $h_{2,t}$ follow this GARCH(1,1) process:

$$\begin{aligned} h_{1,t} &= a_0 + a_1\epsilon_{t-1}^2 + a_2 h_{1,t-1} \\ h_{2,t} &= b_0 + b_1\upsilon_{t-1}^2 + b_2 h_{2,t-1} \end{aligned} \tag{2.9}$$

Rewriting the model into the canonical state space form, we have:

$$\hat{\beta}_t = \begin{bmatrix} 1 & 1 & 0 \end{bmatrix} \begin{bmatrix} \beta_t \\ \epsilon_t \\ \upsilon_t \end{bmatrix}$$

$$(\hat{\beta}_t = \quad H \quad \beta_t^*)$$

$$\begin{bmatrix} \beta_t \\ \epsilon_t \\ \upsilon_t \end{bmatrix} = \begin{bmatrix} \phi_0 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} \phi_1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \beta_{t-1} \\ \epsilon_{t-1} \\ \upsilon_{t-1} \end{bmatrix} + \begin{bmatrix} 0 & 1 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \epsilon_t \\ \upsilon_t \end{bmatrix}$$

$$(\beta_t^* \quad = \quad \mu^* \quad + \quad F \quad \beta_{t-1}^* \quad + \quad G \quad \upsilon_t^*)$$

where

$$\begin{aligned} E(\upsilon_t^* \upsilon_t^{*'}|\psi_{t-1}) &= \begin{bmatrix} h_{1,t} & 0 \\ 0 & h_{2,t} \end{bmatrix} \\ &= R_t \end{aligned}$$

We then modify the model derivation in Harvey, Ruiz and Sentana (1992) and Kim and Nelson (1999), and the resulting Kalman Filter recursions are as follows:

Prediction equations:

$$
\begin{aligned}
\beta^*_{t|t-1} &= E(\beta^*_t|\psi_{t-1}) \\
&= \mu^* + F\beta^*_{t-1|t-1} \\
P^*_{t|t-1} &= E\left[(\beta^*_t - \beta^*_{t|t-1})(\beta^*_t - \beta^*_{t|t-1})'\right] \\
&= FP^*_{t-1|t-1}F' + GR_tG' \\
\eta_{t|t-1} &= \hat{\beta}_t - E(\hat{\beta}_t|\psi_{t-1}) \\
&= \hat{\beta}_t - H\beta^*_{t|t-1} \\
f_{t|t-1} &= E(\eta^2_{t|t-1}) \\
&= HP^*_{t|t-1}H' \qquad (2.10)
\end{aligned}
$$

Updating equations:

$$
\begin{aligned}
\beta^*_{t|t} &= E(\beta^*_t|\psi_t) \\
&= \beta^*_{t|t-1} + P^*_{t|t-1}H'f^{-1}_{t|t-1}\eta_{t|t-1} \\
P^*_{t|t} &= E\left[(\beta^*_t - \beta^*_{t|t})(\beta^*_t - \beta^*_{t|t})'\right] \\
&= P^*_{t|t-1} - P^*_{t|t-1}H'f^{-1}_{t|t-1}\eta_{t|t-1} \qquad (2.11)
\end{aligned}
$$

In Equation (2.9) ϵ^2_{t-1} and υ^2_{t-1} are unobserved. To account for this unobservability, we adopt the same remedy as in Harvey, Ruiz and Sentana (1992), and instead estimate with their conditional expectations:

$$
\begin{aligned}
h_{1,t} &= a_0 + a_1E(\epsilon^2_{t-1}|\psi_{t-1}) + a_2h_{1,t-1} \\
h_{2,t} &= b_0 + b_1E(\upsilon^2_{t-1}|\psi_{t-1}) + b_2h_{2,t-1} \qquad (2.12)
\end{aligned}
$$

Also,

$$
\begin{aligned}
E(\epsilon^2_{t-1}|\psi_{t-1}) &= E(\epsilon_{t-1}|\psi_{t-1})^2 + E\left[(\epsilon_{t-1} - E(\epsilon_{t-1}|\psi_{t-1}))^2\right] \\
E(\upsilon^2_{t-1}|\psi_{t-1}) &= E(\upsilon_{t-1}|\psi_{t-1})^2 + E\left[(\upsilon_{t-1} - E(\upsilon_{t-1}|\psi_{t-1}))^2\right] \qquad (2.13)
\end{aligned}
$$

where $E(\epsilon_{t-1}|\psi_{t-1})$ and $E(\upsilon_{t-1}|\psi_{t-1})$ are the last 2 elements of $\beta^*_{t-1|t-1}$; and $E\left[(\epsilon_{t-1} - E(\epsilon_{t-1}|\psi_{t-1}))^2\right]$ and $E\left[(\upsilon_{t-1} - E(\upsilon_{t-1}|\psi_{t-1}))^2\right]$ are the last two diagonal elements of $P^*_{t-1|t-1}$.

We estimate the model using the maximum likelihood with prediction error, $\eta_{t|t-1}$, and its variance, $f_{t|t-1}$, which we obtained using a Kalman filter. The log likelihood function can be written as:

$$ln(L) = -0.5\sum_{t=1}^{T} ln\left[2\pi f_{t|t-1}^{-1}\right] - 0.5\sum_{t=1}^{T} \eta'_{t|t-1} f_{t|t-1}^{-1} \eta_{t|t-1} \tag{2.14}$$

2.4 Vector Autoregressive Models

The risk-free return, market return and commodity prices are modeled using Vector Autoregressive Process of order p (VAR(p)). To increase the predictive power, we also include in the VAR model the Cyclically Adjusted Price Earnings Ratio (CAPE), whose explanatory power over the long run market return was successfully demonstrated in the book, *Irrational Exuberance*, by Shiller (2005). To ensure the non-negativity of prices and the CAPE, both variables are modeled in their natural logarithm form, resulting in:

$$\mathbf{y}_t = \mathbf{A}_0 + \mathbf{A}_1\mathbf{y}_{t-1} + ... + \mathbf{A}_p\mathbf{y}_{t-p} + \mathbf{u}_t, \tag{2.15}$$

where

$$\mathbf{y}_t = \begin{bmatrix} ln(P_t) \\ r_t^M \\ r_t \\ ln(cape_t) \end{bmatrix} \quad and \quad \mathbf{u}_t = \begin{bmatrix} u_{1,t} \\ u_{2,t} \\ u_{3,t} \\ u_{4,t} \end{bmatrix}$$

The conditional distributions of disturbances are defined as $u_{i,t}|\psi_{t-1} \sim \mathcal{N}(0, g_{i,t})\ \forall i = 1 \cdots 4$; and $g_{i,t}$ follows a GARCH(1,1) process.

$$g_{i,t} = c_{i,0} + c_{i,1}u_{i,t-1}^2 + c_{i,2}g_{i,t-1}, i = 1, 2, 3, 4$$

To compare how model specification would affect the savings and discount rates, we also included cases of VAR(p) with a linear trend, quadratic trend, and price trend interaction. There is a prominent downward trend in risk-free rate of return, predictions of risk-free return will be dominated by the extrapolation of time trend and lead to spurious results; therefore risk-free returns are modeled without time trend for all different model specifications.

The price of energy is frequently assumed to follow a random walk with no drift; Hamilton (2008) tested the statistical behavior of the real price of oil and concluded that it is consistently a random walk. Therefore, for each of the model specifications we have price following a random walk as the special case, and test how the discount rate will change in response to this specification.

2.5 Fuel consumption estimation

In Equation (2.1), the amount of fuel saved every period depends on the total consumption and relative fuel efficiency. According to the US department of Energy, the 2012 Honda Civic has an average EPA MPG of 32 for the standard version, and 44 MPG for the hybrid version.[3] This means that a consumer saves about 25% of the total gas consumption by driving a hybrid rather than standard Civic.

To simplify the model, we estimate the energy consumption using the short-run price elasticities of demand and prices forecasts.

$$Q_t^s = 0.25Q_t$$

where Q_t is the gasoline consumption when driving a standard Civic, at time t.

The price elasticity of demand is defined as:

$$e = \frac{lnQ_{t+1} - lnQ_t}{lnP_{t+1} - lnP_t} \tag{2.16}$$

Using the two equations above, we are able to see that consumption savings change in response to price change and chosen elasticity, given by:

$$Q_{t+1}^s = 0.25exp\left\{[lnQ_t + e(lnP_{t+1} - lnP_t)]\right\} \tag{2.17}$$

2.6 Path-wise Simulations for β_t, r_t, r_t^M, P_t and Q_t^s

The prediction equations for each of the variables used to create this model are:

$$\begin{aligned}
\beta_{T+t} &= \phi_0 + \phi_1\beta_{T+t-1} + \upsilon_{T+t} \\
\mathbf{y}_{T+t} &= \mathbf{A}_0 + \mathbf{A}_1\mathbf{y}_{T+t-1} + ... + \mathbf{A}_p\mathbf{y}_{T+t-p} + \mathbf{u}_{T+t} \\
Q_{T+t}^s &= 0.25exp\left\{[lnQ_{T+t-1} + e(lnP_{T+t} - lnP_{T+t-1})]\right\} \\
h_{2,T+t} &= b_0 + b_1\upsilon_{T+t-1}^2 + b_2h_{2,T+t-1} \\
g_{i,T+t} &= c_{i,0} + c_{i,1}u_{i,T+t-1}^2 + c_{i,2}g_{i,T+t-1}, i = 1, 2, 3, 4
\end{aligned} \tag{2.18}$$

where T refers to the end period of the observed data. Due to the serial dependence of the time series data, we calculate the future values using their lags and instantaneous random shocks. This proceeds as follows:

[3]Source: $http://www.fueleconomy.gov/feg/bymodel/2012_{H}onda_{C}ivic.shtml$

1. At time T, randomly draw 1,000 sets of β_T, υ_T and $\mathbf{u}_T$ from their estimated distributions. To simplify calculations, we normalize $ln(Q_T)$ to 1;

2. At time $T+t$, calculate the variances of disturbances, $h_{2,T+t}$ and $g_{i,T+t}$, using past realizations of the error terms and lag variances. Draw instantaneous disturbances υ_{T+t} and $u_{i,T}$ using the updated distributions;

3. Calculate β_{T+t}, $\mathbf{y}_{T+t}$ and Q^s_{T+t} following the equations listed above using their lag values and the random disturbances drawn in step 2.

4. Repeat step 2 and 3 $\forall t = 1, \cdots, N$. And the distributions for time horizons, N, are selected according to typical lifetimes of prominent energy-related investments, like vehicles, household appliances, photo-voltaic solar panels, and windmills.

2.7 Data and Estimation Results

We choose crude oil for the study because gasoline prices are closely related to the price of crude oil, and more importantly, crude oil prices have a bigger influence over energy policies, given its market share (37% of the U.S. energy consumption[4]), the turbulent market conditions and its well-documented connections with the overall economy (Hamilton 1983, Barsky and Kilian 2004, Kilian 2009).

The daily crude oil spot prices are obtained from the Commodity Research Bureau (CRB). The excess returns are defined as:

$\gamma_t = ln(P_t) - ln(P_{t-1}) - r_{ft}$

S&P 500 indices are used as the proxy for the market portfolio, the indices, dividends, and earnings data are obtained from the Standard and Poor's online database.[5] Its price returns, used in the vector autogressive model in Section 2.4, is defined as:

$r_t^M = ln(SP_t + Dividend_t) - ln(SP_{t-1})$.

where daily dividend is calculated from the quarterly ones, and the excess market return is:

$\gamma_t^M = r_t^M - r_{ft}$

We use 3-month Treasury bill as our risk-free asset, and follow the standard yield calculation on the Treasury bill and the return is:

$r_{ft} = \left[\frac{1}{1-DTB3_t\cdot(91/360)}\right]^{(days/91)} - 1$

where $DTB3_t$ is the daily quotation in the secondary market on the most recently auctioned Treasury Bills for 3-month maturity. The daily rate is from FRED.®[6] The rate will be used

[4]Source: *http : //www.eia.gov/totalenergy/data/annual/pdf/sec1₉.pdf*, percentage is calculated using share of Btu.

[5]Source: $http://www.standardandpoors.com/indices/sp-500/en/us/?indexId = spusa-500-usduf--p-us-l--$.

[6]Source: *http : //research.stlouisfed.org/fred2/series/DTB3.*

until new rate becomes available. *days* is the number of calendar days investors hold t-bill, and it is 1 for weekdays, and 3 over the weekend except for holidays.

We adopt CAPE data developed by Robert Shiller and linearly interpolate to bi-weekly frequency for our analysis in Section 2.4.[7]

Our daily data begins in the early 1980s and extended to December 2011. The simulated data predictions span from 2012-2112, covering the lifetimes of most energy-related investments. To predict the energy savings, we also adopt the short-run price elasticity of gasoline demand estimated by Hughes, Knittel and Sperling (2008). In their paper, they use aggregate monthly data from January 1974 to March 2006, and find that the price elasticities range from -0.034 to -0.077 during the period from 2001-2006, and -0.21 to -0.34 in 1975-1980. Given the time window being studied, 1983-2011, we use the price elasticities of -0.034 and -0.077 that reflect the similar time period in our model, and have a special case of perfectly inelastic demand, i.e. price elasticity is equal to zero, as a comparison.

After filtering out the estimation error using the State Space Model, the bi-weekly beta is estimated to follow with the standard errors in the parentheses:

$$\beta_t = \underset{(0.0094)}{-0.0057} + \underset{(0.0324)}{0.9234}\ \beta_{t-1} + v_t$$

Beta is strongly persistent with a long-run average of -0.0744 which implies an average -0.37% risk premium in the oil market assuming a 5% market risk premium. Oil covaries negatively with the overall market, on average providing an insurance against the market. Figure (2.1) shows how the bi-weekly *beta* evolved over time, as well as its 90% confidence interval when beta follows AR(1). The shaded areas at the top of the figures show the major recessions during this time period, where the sluggish economy lowered the demand for crude oil, causing the price of oil to fall, giving rise to positive beta coefficients. The shaded areas at the bottom of the figure point out the time periods with major wars and unrests in the Middle East that interrupted the oil supply in those areas and caused oil price to rise; rising oil prices then, in turn, slowed economic activities, and created negative betas in these cases. The trends of *beta* shown in Figure (2.1) coincide well with the important historical events in oil history. When Saudi Arabia stopped being the "Swing Producer" and raised its oil production, the oil price collapsed to $12/barrel by mid-1986, and beta was significantly negative in 1986. During the first Gulf War in 1990-1991, we also have a significantly negative beta. Although there was a recession around that same time, the negative demand shock caused by the recession was overwhelmed by the strong negative supply shock due to the war. The recent oil price spike was caused by the high demand of oil due to persistent strong growth in Asian countries, and

[7]Source: *http : //www.econ.yale.edu/ shiller/data.htm*

low spare production capacity. Despite the fact that it happened during the recent recession, beta for this period appears to be negative. The recession became the leading force in the oil market after the spike and the positive beta trend prevailed, until the very recent unrest in the Middle East created a dip in the trend of beta.

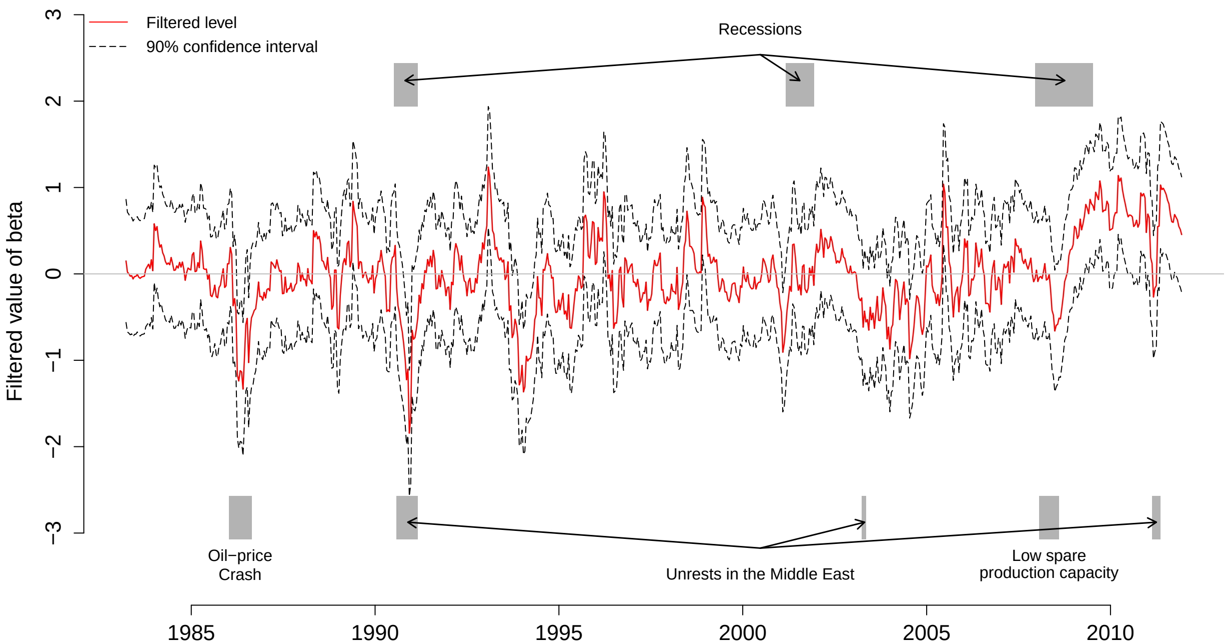

Figure 2.1: Filtered Beta (AR(1)) with 90% Confidence Interval

Calculating the risk premium in the oil market using the CAPM relies on the validity of the model which can be easily tested using the bi-weekly *beta* and the risk premia in the oil market as follows:

$$\gamma_t = \xi_0 + \xi_1(\beta_t \gamma_t^M) + \mu_t$$

where β_t is the estimated beta of bi-week t. γ_t and γ_t^M are bi-weekly risk premium of crude oil and the market portfolio respectively. When $\xi_0 = 0$ and $\xi_1 = 1$, the realized risk premium equals the CAPM prediction plus a pricing error.

The following equation shows the result testing the conditional CAPM using the filtered *betas* where the estimation errors in the estimated *betas* from the short-window regressions are corrected for with the State Space model.

$$\gamma_t = \underset{(0.2455)}{-0.0239} + \underset{(0.1781)}{0.9274} \; (\beta_t \gamma_t^M) + \mu_t$$

We fail to reject the conditional CAPM and there is significant pricing of risk with the model in the oil market, where the rejections of the model in earlier literature can very well be the result of the large variations in the *betas* that they fail to account for. Figure (2.2) shows the same test results where the risk premia are grouped into 50 quantiles based on levels of the CAPM predictions. Comparing the actual fitted line in black, and the perfect CAPM predictions, the gray dotted line, we can conclude that the conditional CAPM captures the oil market rather accurately and it is, to our acknowledge, the most compelling support to CAPM in the literature.

VAR(1) suffices for our estimation purposes, and the following is the estimation results of VAR(1) model with a linear intercept (* indicates that estimates are significant at 5% level).

$$\begin{bmatrix} ln(\hat{P_t}) \\ \hat{r_t^M} \\ \hat{r_t} \\ ln(\hat{cape_t}) \end{bmatrix} = \begin{bmatrix} 0.124^* \\ 9.25^* \\ 0.05 \\ 0.10^* \end{bmatrix} + \begin{bmatrix} 5e-07 \\ 2e-03^* \\ 0 \\ 2e-05^* \end{bmatrix} t + \begin{bmatrix} 0.98^* & 0.00 & -0.15^* & -0.01 \\ -1.28^* & 0.03 & 0.12 & -1.65^* \\ -0.00 & -0.00^* & 0.85^* & -0.01^* \\ -0.01^* & 0.00 & -0.00 & 0.98^* \end{bmatrix} \times \begin{bmatrix} ln(P_{t-1}) \\ r_{t-1}^M \\ r_{t-1} \\ ln(cape_{t-1}) \end{bmatrix}$$

The results of the simulation for *beta*, the oil prices, the amount of fuel saving, the market return, and the risk-free return over the next 5 years are shown in Figure (2.3) - (2.7). In each figure, the solid lines indicate the mean values of the 1,000 iterations, and the shaded areas represent the range of the 25% and 75% quartiles of the 1,000 iterations.[8]

Under the assumption that *beta* follows random walk with no drift, we introduce more

[8] We sort the1,000 simulated paths at each time period, and choose the 25% and 75% quartiles.

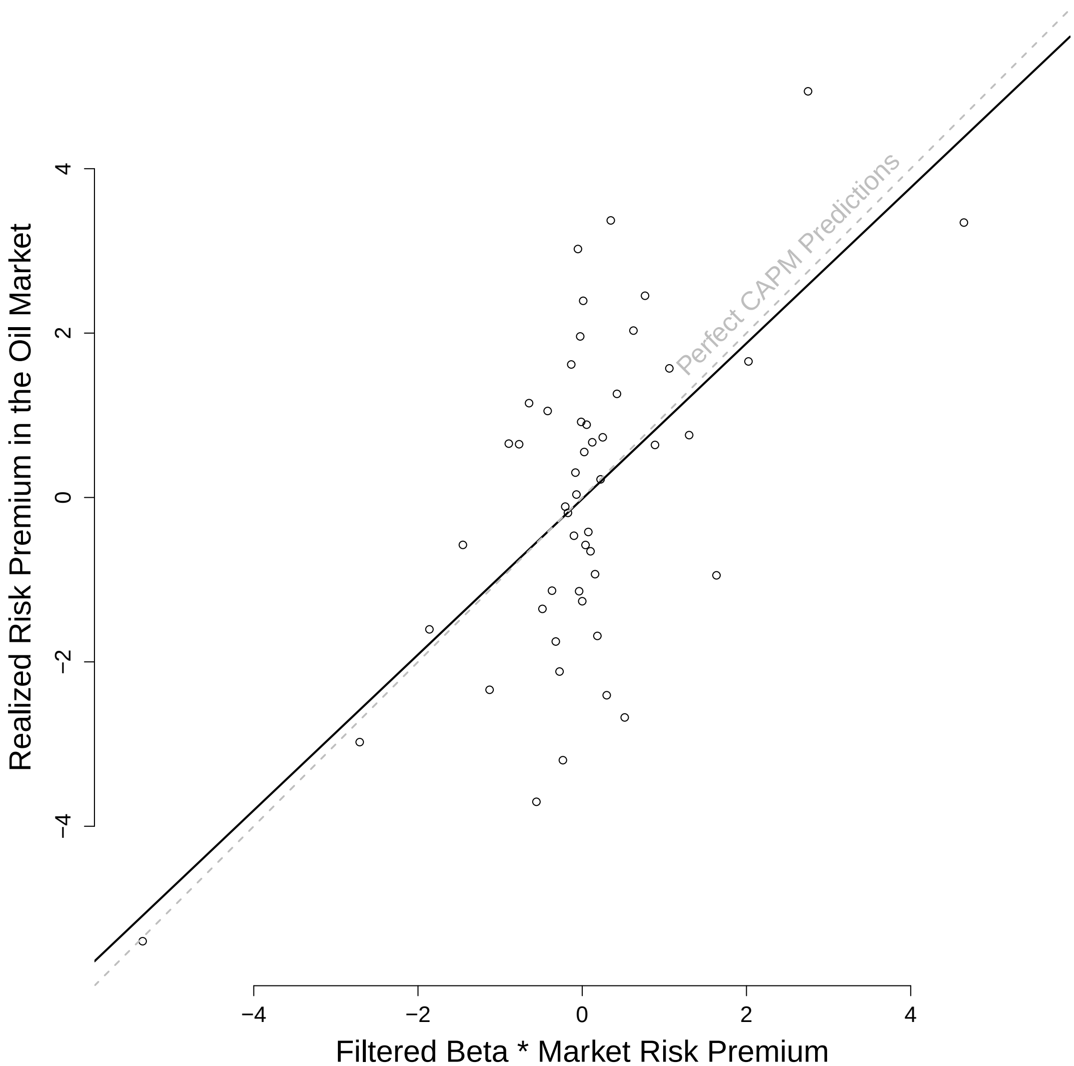

Figure 2.2: The CAPM predicted risk premium (%/bi-week)

uncertainty into the model. In Figure(2.3), *beta* forecasts have a much wider range under the assumption of random walk than under an AR(1), and its uncertainty is increasing at a much faster rate into the future as well. The price forecasts in Figure (2.4) are increasing at a faster rate when price follows random walk than VAR process. Figure (2.5) shows how the amount of fuel saved each period changes with the price elasticities; the more responsive consumers are to fuel prices, the more uncertainty there is in the model. When consumers are perfectly inelastic, the amount of fuel saved each period stays constant at the current level with no variations. Figure (2.6) and (2.7) show the market and risk-free return forecasts under 4 different specifications of the VAR model (Equation 2.15); in general, there is the most uncertainty in the model when there is a quadratic trend, and the least when there is no trend. Models with linear trend and price trend interaction appear to have similar levels of uncertainty when considering a 5 year window, and into the future as time horizon increases.

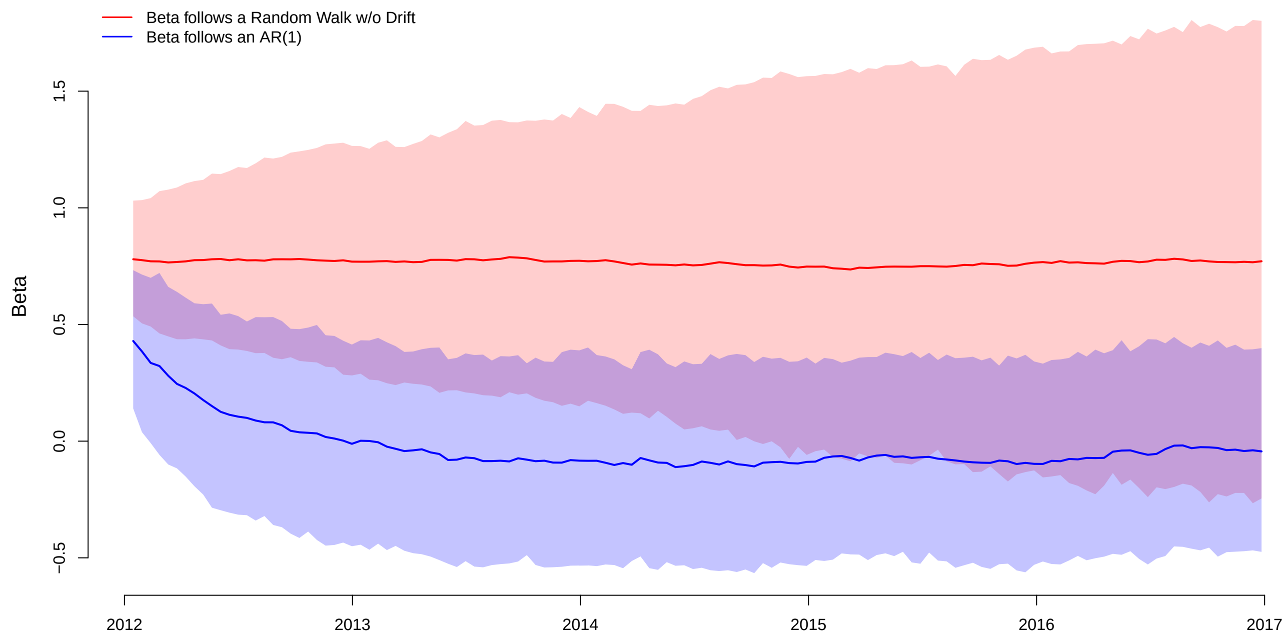

Figure 2.3: Simulated predictions of *beta* over the next 5 years

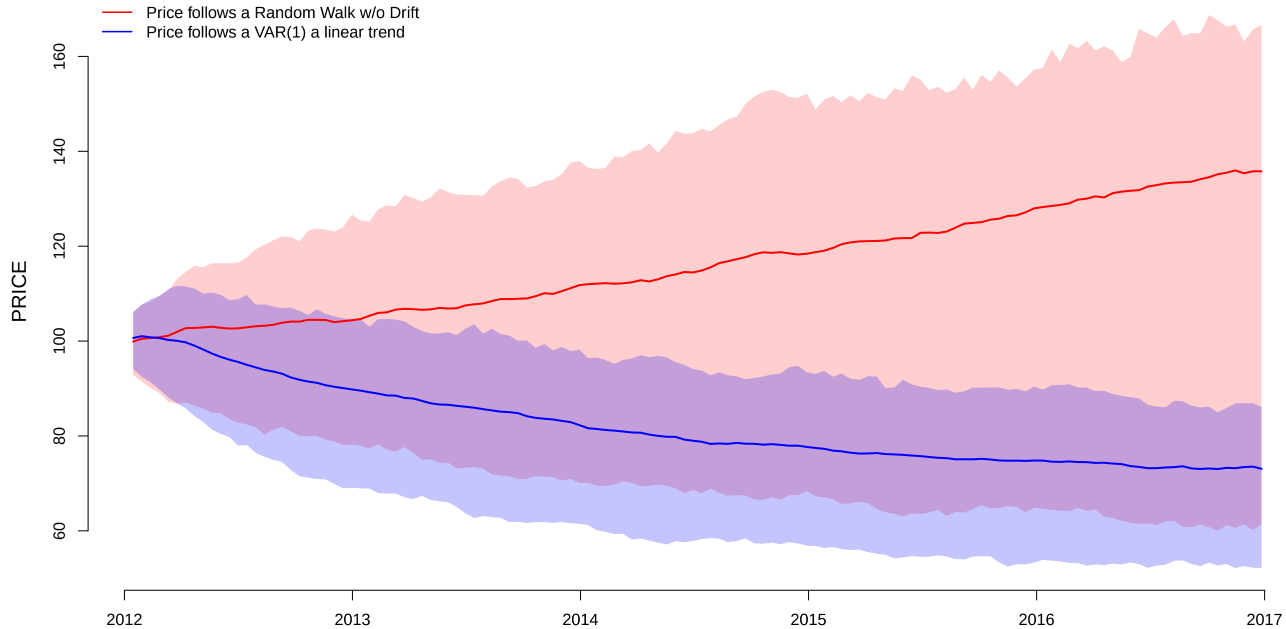

Figure 2.4: Simulated predictions of oil price over the next 5 years

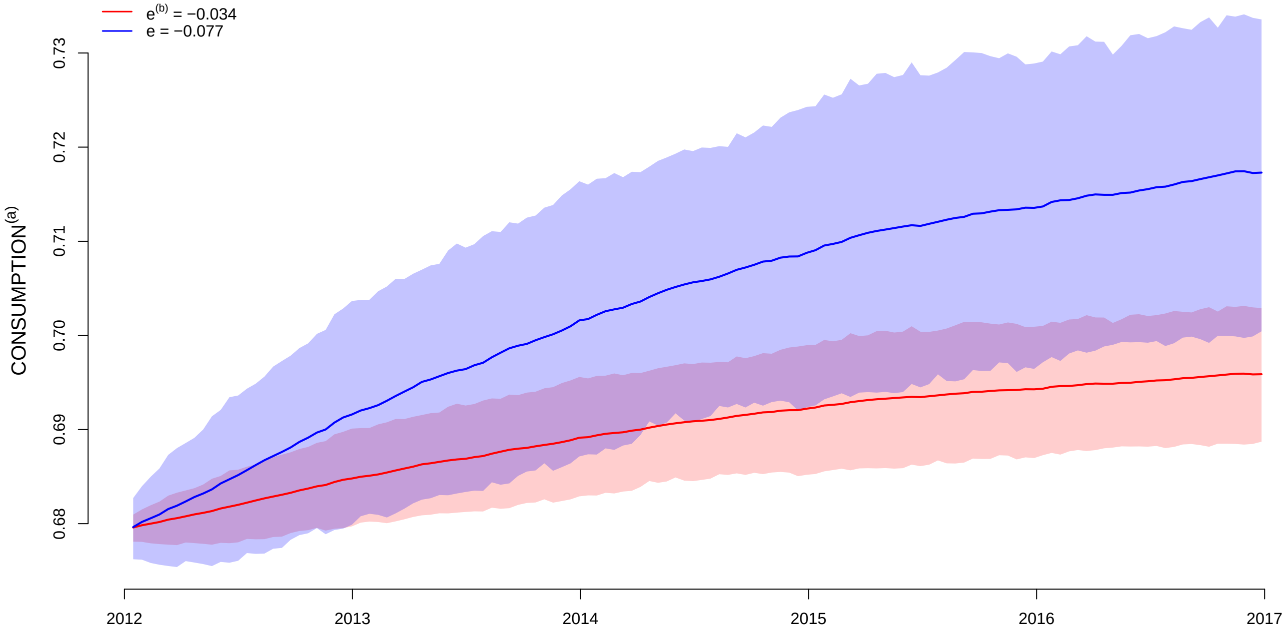

Figure 2.5: Simulated predictions of energy savings over the next 5 years

Note:

1. Quantity of fuel savings assumes that price follows a VAR(1) with no trend
2. e is the price elasticity of demand

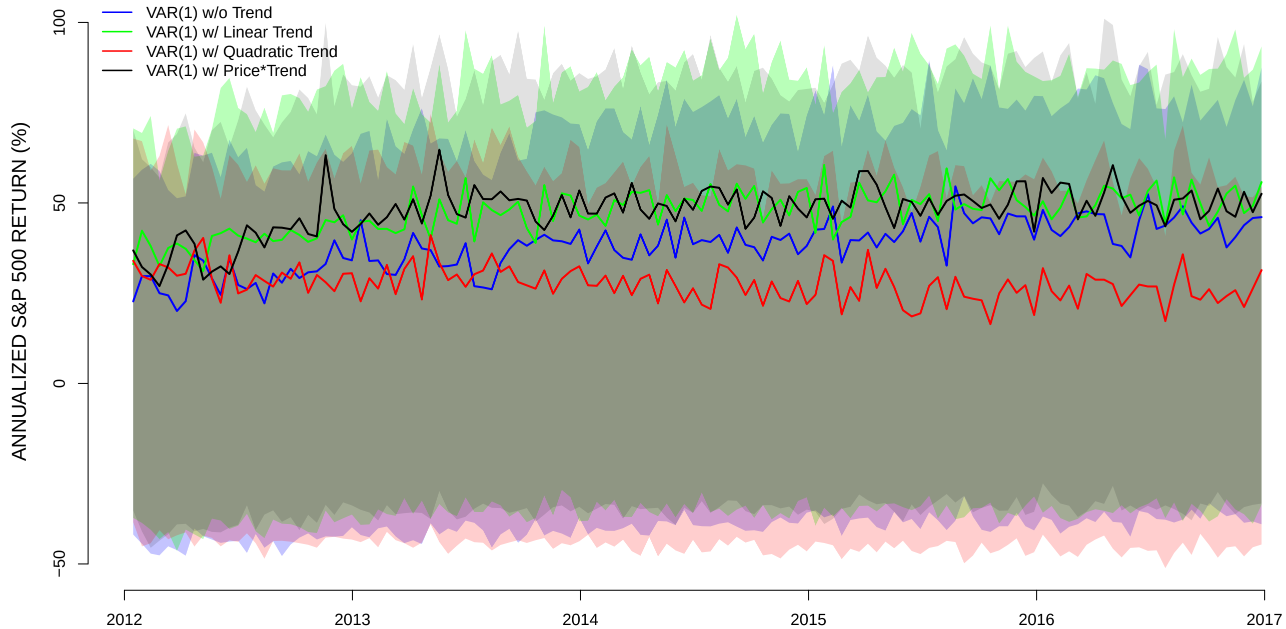

Figure 2.6: Simulated predictions of market return over the next 5 years

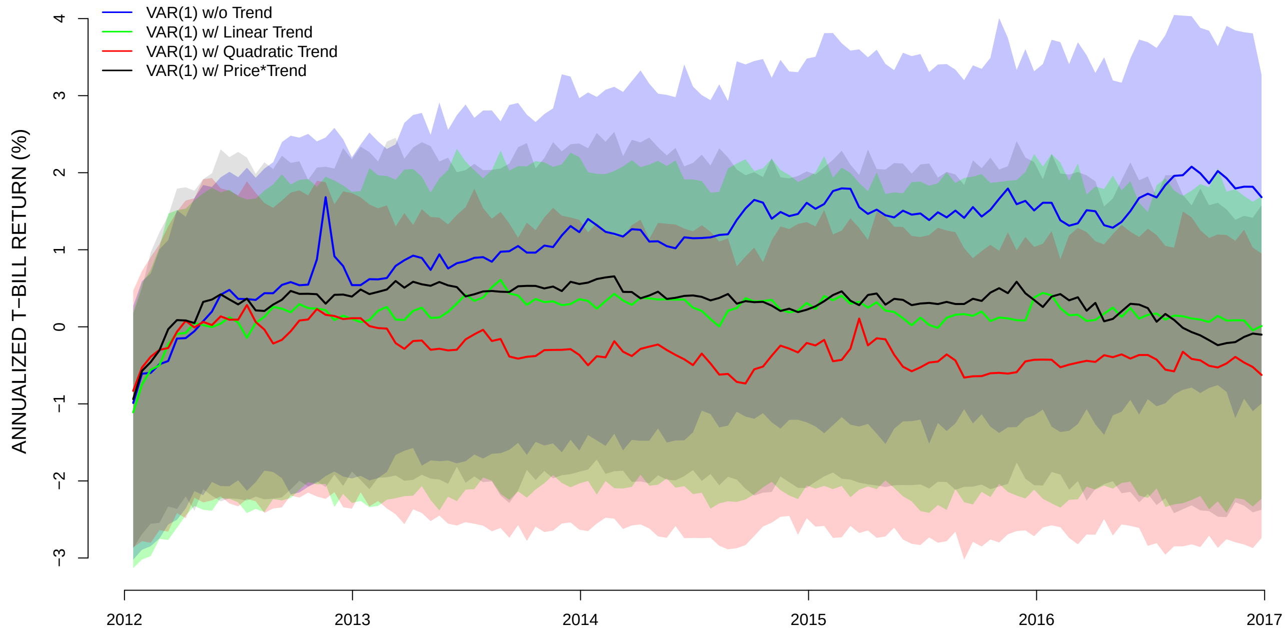

Figure 2.7: Simulated predictions of risk free return over the next 5 years

To put discounting under uncertainty and risk into a tangible perspective, we define a constant-equivalent discount rate that gives the same expected present value as the stochastic case. The constant-equivalent rate also assumes prices are constant. Some argue that consumers in the case of normal economic activities use the current price as the prediction of future prices (Anderson, Kellogg and Sallee 2011) and discount (if they do) at a constant rate instead of different instantaneous rates at different time periods. This facilitates a clear comparison between our analysis and previous ones, as well as illustrates how the discount rate would change in response to different kinds of uncertainty. Specifically, we define the constant-equivalent rates as follows:

$$S = P_0 Q_0^s exp\left(-\delta\right) + P_0 Q_0^s exp\left(-2\delta\right) + ... + P_0 Q_0^s exp\left(-N\delta\right). \tag{2.19}$$

where P_0 and Q_0^s are the current price and fuel savings; δ is the discount rate. If δ yields the same present value energy savings over the same N periods as our simulated results, δ is the constant-equivalent discount rate.

We calculate the total present value of fuel savings and constant equivalent discount rates under different model specifications over the lifetime of a vehicle assuming a 10-year lifetime. In Table (2.1), discount rates are similar across different elasticities within the same model specification.

On the other hand, with similar specifications, models with a random walk setup have a lower discount rate than the models without. As described above, the presence of a random walk introduces more uncertainty into the model. Therefore, the more uncertainty there is, the lower the constant-equivalent discount rate.

Table 2.1: The Annualized Constant Equivalent Discount Rate (%) for Projects with 10 Years Lifetime

	Price follow VAR(1)			Price follow Random Walk		
	$e^{(b)} = 0$	e = -0.034	e = -0.077	e = 0	e = -0.034	e = -0.077
Beta follows AR(1) Others VAR(1) w/o trend	12.20	11.84	11.37	-4.62	-4.26	-3.81
Beta follows AR(1) Others VAR(1) w/ linear trend	10.05	9.71	9.27	-5.83	-5.46	-5.01
Beta follows AR(1) Others VAR(1) w/ price*trend	9.76	9.40	8.94	-6.89	-6.51	-6.05
Beta follows AR(1) Others VAR(1) w/ quadratic trend	-15.02	-14.48	-13.79	-4.55	-4.23	-3.84
Beta follows RW Others VAR(1) w/o trend	9.46	9.12	8.69	-56.35	-56.09	-55.77
Beta follows RW Others VAR(1) w/ linear trend	8.95	8.62	8.20	-85.38	-85.29	-85.16
Beta follows RW Others VAR(1) w/ price*trend	8.75	8.40	7.96	-92.13	-92.06	-91.97
Beta follows RW Others VAR(1) w/ quadratic trend	-36.54	-36.04	-35.40	-54.19	-54.10	-54.00

Notes:

1. e: price elasticity of demand
2. RW: Random Walk with no drift

Figure (2.8) shows the discount rate trend when the project's lifetime increases. Uncertainty increases with the lifetime of almost all projects, which leads to a decreasing discount rate. In Table (2.2), with the more conservative model in which *beta* follows an AR(1) and price and other variables follow a VAR(1) with no trend, the annualized discount rate decreases from 17% with a 5-year lifetime to 2% when considering 100-year lifetime. The pattern is consistent with Newell and Pizer (2003), who found that with interest rate as the single source of uncertainty, the discount rate falls from 4%, to 2% in 100 years, and further to 1% in 200 years. But their results consider a much longer time horizon while ours consider the uncertainty in the near and intermediate future. And without assuming non-negative discount rate, some of our models with higher level of uncertainty give negative discount rate weighing future more than the present.

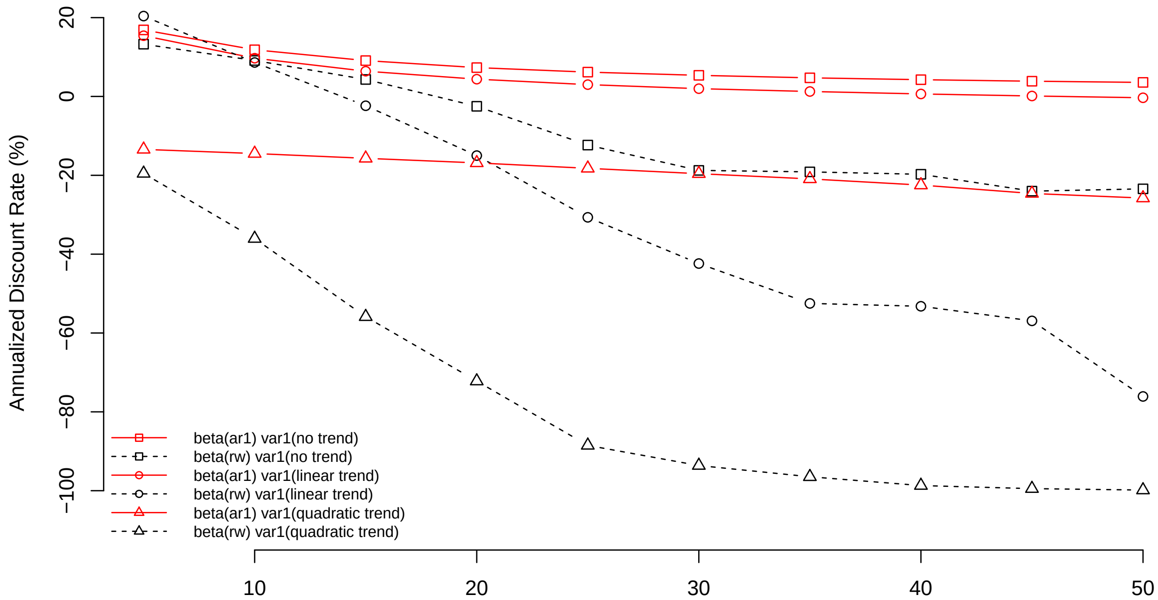

Figure 2.8: Annualized Constant equivalent Discount Rate (%)

Note:

1. The constant equivalent discount rates in this figure are calculated assuming price elasticity = -0.034, and that price and other variables follow a VAR(1)
2. Models with price trend interaction generate results very similar to models with linear trend thus not shown in this graph

When comparing the results using different elasticities over various durations of projects, we will see that the more elastic the demand is, the lower the absolute values of the discount rate are. For projects with longer duration, the demand will be more elastic given that consumers have opportunities to make adjustments and therefore a smaller constant-equivalent discount rate. Figure 2.9 shows how the constant-equivalent discount rate would change with different choices of elasticities.

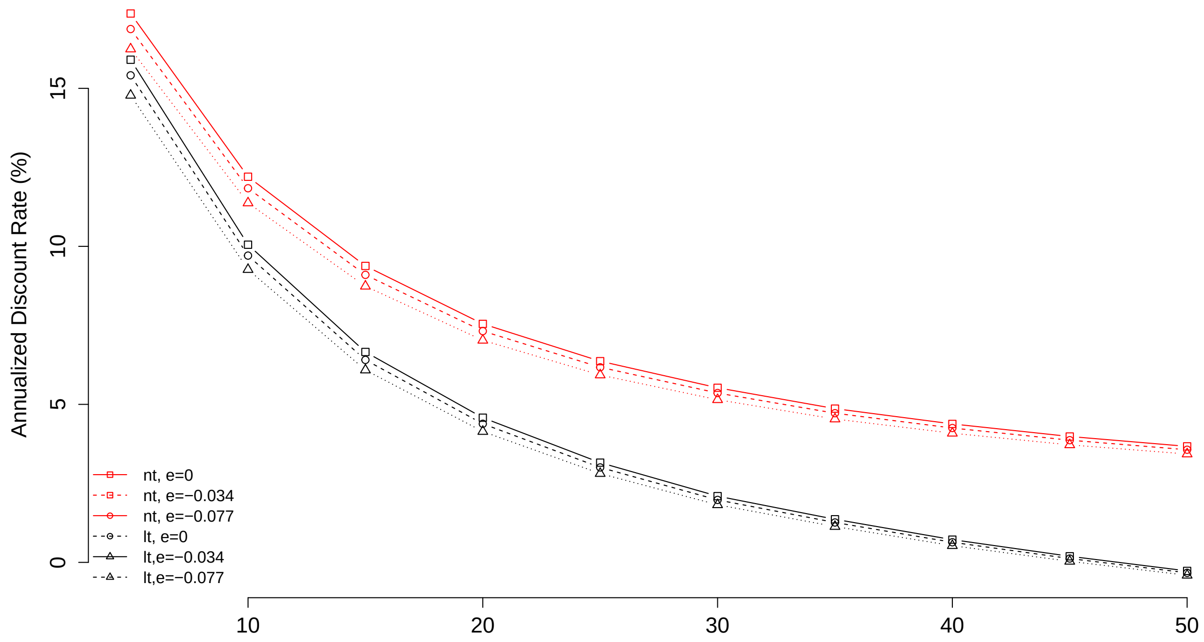

Figure 2.9: Constant-Equivalent Discount Rate by Elasticity of Demand

Note:

1. e: elasticity of demand
2. nt: Vector Autoregressive model with no trend
3. lt: Vector Autoregressive model with linear trend

Table 2.2: The Annualized Constant Equivalent Discount Rate (%) for Projects with Different Lifetimes

	Oil Price Follows VAR(1)				Oil Price Follows Random Walk with no Drift			
	5 Years	25 Years	50 Years	100 Years	5 Years	25 Years	50 Years	100 Years
Beta follows AR(1) Others VAR(1) w/o trend	16.88	6.17	3.56	2.13	-4.39	-4.45	-9.19	-5.93
Beta follows AR(1) Others VAR(1) w/ linear trend	15.41	3.00	-0.33	-3.64	-4.88	-7.73	-19.94	-21.20
Beta follows AR(1) Others VAR(1) w/ price*trend	16.93	2.85	-0.15	-3.31	-4.86	-13.43	-35.44	-57.22
Beta follows AR(1) Others VAR(1) w/ quadratic trend	-13.42	-18.23	-25.77	-38.00	-4.24	-13.55	-56.47	-86.97
Beta follows RW Others VAR(1) w/o trend	13.26	-12.33	-23.44	-45.28	-18.25	-92.35	-99.81	NA
Beta follows RW Others VAR(1) w/ linear trend	20.39	-30.66	-76.09	-95.81	-24.59	-99.58	-100.00	NA
Beta follows RW Others VAR(1) w/ price*trend	22.40	-36.50	-55.20	-91.77	-35.71	-99.97	NA	NA
Beta follows RW Others VAR(1) w/ quadratic trend	-19.52	-88.51	-99.84	NA	-12.27	-99.99	NA	NA

Notes:

1. The table shows results assuming price elasticity e = -0.034
2. RW: Random Walk with no drift

Because the projected time path of prices and discount rates depends on the initial conditions, we consider different starting points for the simulation process, in 1997, 2002, 2007 and 2012 respectively. The constant-equivalent discount rates, shown in 2.3, vary rather significantly across different starting points. The value of the constant-equivalent discount rate is affected by the level of uncertainty: (1) the discount rate is smaller for random walk models than autoregressive; (2) the discount rate falls with the duration of projects. The rate is also affected largely by the predicted paths of energy prices, beta, market return and risk-free return. With higher energy prices and large positive co-movement of the energy prices and overall market (thus positive risk premium), the constant-equivalent discount rate tend to be much higher in recent periods than in earlier periods.

Table 2.3: The Annualized Constant Equivalent Discount Rate for Simulations using Different Starting Points

	Simulation Starting Time	Project Lifetime 5 years	10 years	15 years	20 years	25 years	50 years
Beta follow RW Others VAR(1) w/o Trend	12/9/2012	13.71	9.46	4.55	-2.33	-12.23	-23.39
	2/9/2007	6.29	2.44	-3.87	-11.43	-21.25	-29.08
	8/2/2002	-12.12	-8.48	-9.66	-13.46	-19.28	-32.84
	7/11/1997	-15.00	-11.40	-10.25	-11.80	-20.41	-35.17
Beta follow AR(1) Others VAR(1) w/o Trend	12/9/2012	17.37	12.20	9.38	7.55	6.37	3.67
	2/9/2007	6.56	4.80	3.76	3.07	2.61	1.66
	8/2/2002	-11.29	-6.18	-4.22	-3.13	-2.44	-1.02
	7/11/1997	-14.49	-8.97	-6.33	-4.80	-3.84	-1.77

Notes:

1. The table shows results assuming price elasticity e = 0
2. RW: Random Walk with no drift

Table 2.4: Start Values for Each Point in Time

	Energy Beta	Energy Price	Market Return (%)	Risk-free Return (%)
12/9/2012	0.78	99.56	8.01	-0.05
2/9/2007	0.04	66.63	0.94	-0.03
8/2/2002	0.12	33.82	1.86	-0.06
7/11/1997	-0.26	27.18	3.21	0.10

Notes:

1. Energy prices are real prices in 2012 $
2. Market return and Risk-free Return are inflation adjusted real return

2.8 Discussion

Our study aims to estimate the proper market rate of discount for energy-related investments under uncertainty. Using 30 years of daily data on crude oil prices, the 3-month treasury bill rates, and S&P 500 returns, we estimated several VAR models of how these variables evolve, as well as how the market *beta* of crude oil changes in response to changes in fundamental market conditions over time.

We find that the market *beta* is highly persistent with a long-term average of -0.07. It also varied from negative to positive when responding to negative beta events such as unrests in the Middle East, and positive beta events like recessions. Our VAR models estimates suggest that both price of crude oil and risk-free return are persistent; although market return is not, the extra variable CAPE adds a lot to it. Newell and Pizer (2003) found that the certainty-equivalent discount rate will decrease to its lower boundary when the lifetime increases under uncertainty with persistence. We have similar findings in our simulation-based method: for the model with the least amount of uncertainty, our certainty-equivalent discount rate decreases from 18% in 5 years to 2% in 100 years. Comparing across different model specifications, the more uncertain the model is, the lower the discount rate. With high levels of uncertainty, the discount rate sometimes even becomes negative.

Our estimates show a consistent pattern across the models that the discount rate decreases when the project lifetime increases. The low discount rates implied by this analysis greatly increases the value of projected future energy savings as compared to conventional analyses, even with the most conservative model. With any negative discount rate, as some of the model suggested, even at a modest level, the future is weighed more than the present. In this case, should we go all out and invest in more of these energy conservation projects for our future's benefit? The short answer is no. CAPM assumes complete markets and no transaction costs. Individuals and firms making decisions about energy-saving investments are likely to be constrained in many ways that would cause them to value energy savings differently than frictionless financial markets would. On the one hand this would seem to give rise to arbitrage opportunities. On the other hand, presumably there are considerable transactions costs and information problems that arise between investors and individuals making energy investment and consumption decisions, and these tensions prevent efficient trade and investment. As a result, it is impossible to know whether greater investments in energy saving technologies could be efficiently made without consideration of the mechanism that would give rise to such investments.

APPENDIX

2.9 Appendix

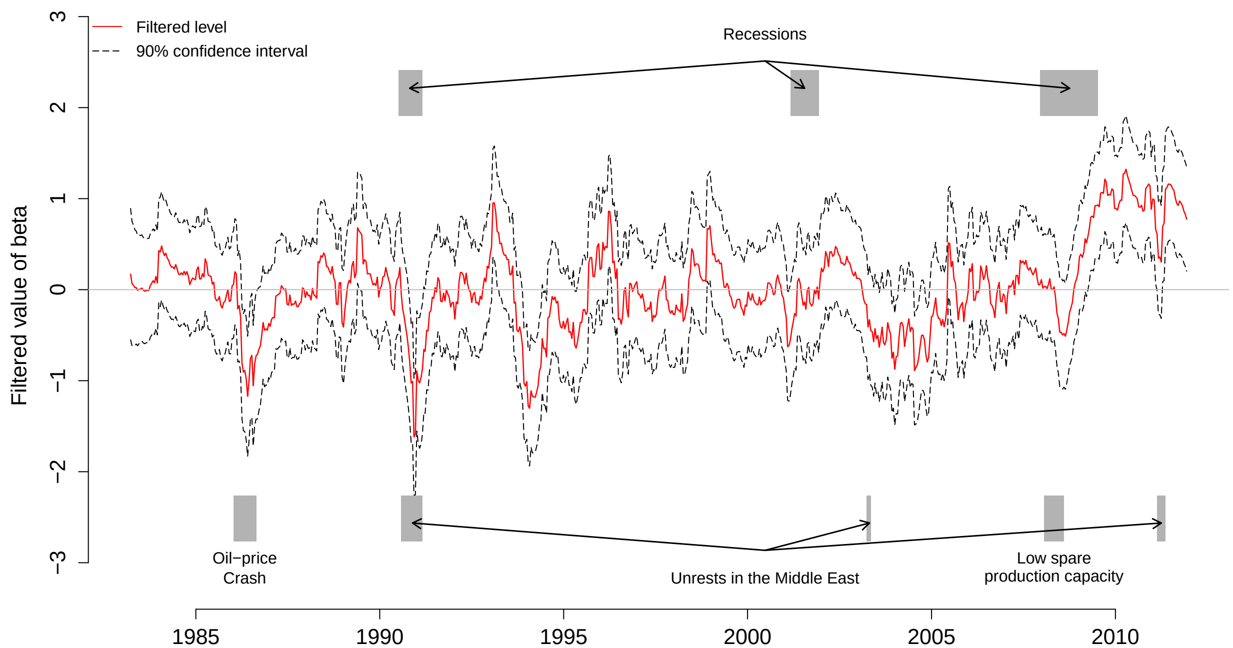

Figure 2.10: Filtered Beta (Random Walk) with 90% Confidence Interval

Table 2.5: The Annualized Equivalent Discount Rate (%) for Projects with 25 Years Lifetime

	Price follow VAR(1)			Price follow Random Walk		
	e$^{(b)}$ = 0	e = -0.034	e = -0.077	e = 0	e = -0.034	e = -0.077
Beta follows AR(1) Others VAR(1) w/o trend	6.37	6.17	5.93	-4.82	-4.45	-4.00
Beta follows AR(1) Others VAR(1) w/ linear trend	3.16	3.00	2.81	-8.11	-7.73	-7.27
Beta follows AR(1) Others VAR(1) w/ price*trend	3.01	2.85	2.66	-13.72	-13.43	-13.10
Beta follows AR(1) Others VAR(1) w/ quadratic trend	-18.94	-18.23	-17.32	-13.46	-13.55	-13.67
Beta follows RW Others VAR(1) w/o trend	-12.23	-12.33	-12.46	-92.38	-92.35	-92.32
Beta follows RW Others VAR(1) w/ linear trend	-30.61	-30.66	-30.71	-99.58	-99.58	-99.58
Beta follows RW Others VAR(1) w/ price*trend	-36.46	-36.50	-36.56	-99.97	-99.97	-99.97
Beta follows RW Others VAR(1) w/ quadratic trend	-88.62	-88.51	-88.36	-99.99	-99.99	-99.99

Notes:

1. e: price elasticity of demand
2. RW: Random Walk with no drift

Table 2.6: The Annualized Equivalent Discount Rate (%) for Projects with 50 Years Lifetime

	Price follow VAR(1)			Price follow Random Walk		
	$e^{(b)} = 0$	e = -0.034	e = -0.077	e = 0	e = -0.034	e = -0.077
Beta follows AR(1) Others VAR(1) w/o trend	3.67	3.56	3.43	-9.73	-9.19	-8.51
Beta follows AR(1) Others VAR(1) w/ linear trend	-0.27	-0.33	-0.41	-20.33	-19.94	-19.45
Beta follows AR(1) Others VAR(1) w/ price*trend	-0.09	-0.15	-0.24	-35.72	-35.44	-35.07
Beta follows AR(1) Others VAR(1) w/ quadratic trend	-26.71	-25.77	-24.57	-56.38	-56.47	-56.59
Beta follows RW Others VAR(1) w/o trend	-23.39	-23.44	-23.49	-99.82	-99.81	-99.81
Beta follows RW Others VAR(1) w/ linear trend	-76.08	-76.09	-76.11	-100.00	-100.00	-100.00
Beta follows RW Others VAR(1) w/ price*trend	-55.18	-55.20	-55.21	NA	NA	NA
Beta follows RW Others VAR(1) w/ quadratic trend	-99.84	-99.84	-99.84	NA	NA	NA

Notes:

1. e: price elasticity of demand
2. RW: Random Walk with no drift

Table 2.7: Annualized Constant Equivalent Discount Rate (%) for Projects with Different Lifetimes

	Oil Price Follows VAR(1)				Oil Price Follows Random Walk with no Drift			
	5 Years	25 Years	50 Years	100 Years	5 Years	25 Years	50 Years	100 Years
Beta follows AR(1) Others VAR(1) w/o trend	17.37	6.37	3.67	2.19	-4.76	-4.82	-9.73	-6.07
Beta follows AR(1) Others VAR(1) w/ linear trend	15.91	3.16	-0.27	-3.62	-5.24	-8.11	-20.33	-21.04
Beta follows AR(1) Others VAR(1) w/ price*trend	17.53	3.01	-0.09	-3.29	-5.22	-13.72	-35.72	-57.10
Beta follows AR(1) Others VAR(1) w/ quadratic trend	-13.93	-18.94	-26.71	-39.30	-4.60	-13.46	-56.38	-86.95
Beta follows RW Others VAR(1) w/o trend	13.71	-12.23	-23.39	-45.28	-18.70	-92.38	-99.82	NA
Beta follows RW Others VAR(1) w/ linear trend	20.87	-30.61	-76.08	-95.81	-25.20	-99.58	-100.00	NA
Beta follows RW Others VAR(1) w/ price*trend	23.00	-36.46	-55.18	-91.77	-36.27	-99.97	NA	NA
Beta follows RW Others VAR(1) w/ quadratic trend	-20.02	-88.62	-99.84	NA	-12.81	-99.99	NA	NA

Notes:

1. The table shows results assuming inelastic energy demand, i.e. e = 0
2. RW: Random Walk with no drift

Chapter 3

Can Standards Increase Consumer Welfare? Evidence from a Change in Clothes Washer Energy Efficiency Requirements

3.1 Introduction

Minimum quality standards are a prevalent feature of regulation, particularly with regard to energy efficiency and safety of buildings, cars, and appliances. In the conventional model of a competitive economy with perfect information, it is not clear how such standards might improve social welfare. Rather, minimum quality standards are normally justified using models of asymmetric information (Leland, 1979), imperfect competition (Ronnen, 1991) or perception bias (Allcott, 2011). Government's regulatory analyses typically justify standards by explicitly evaluating the costs and benefits of specific product attributes that affect efficiency and safety. There are at least two problems with this approach: it generally requires strong and often unverifiable assumptions about intangible costs and preferences and it cannot account for the influence of standards on market and industry structure.

In this study we take a new and simple approach to empirically evaluate a change in standards. We exploit a 2007 policy change in energy efficiency standards for clothes washers to obtain direct evidence on the consumer welfare impact of the more stringent energy efficiency standard. This policy change increased both the minimum threshold of energy efficiency and the threshold for Energy Star certification. The minimum threshold increased from a Modified Energy Factor (MEF) of 1.04 to 1.26; energy star certification increased from 1.42 to 1.72. (MEF measures the ratio of the capacity of the washer to the energy used in one cycle.) We exploit this policy change as an exogenous intervention in the washer market and examine how prices and quantities sold of different washer types changed with the change in policy. This analysis allows for partial identification of consumer welfare impacts, including a lower bound on the net change in consumer surplus. The results indicate a positive change in consumer welfare and indicate a pattern of response that aligns with Ronnen's model of imperfect competition and economies of scale (Ronnen, 1991).

3.2 Consumer Welfare Analysis of a Change in Standards

The 2007 change in clothes washer standards could have had a wide range of economic impacts on consumers. To evaluate these effects, it helps to treat the washer market as a mix of different types of washer, which vary in quality, functions, size, efficiency levels, and manufacturer reputation. The policy change effectively phased out one portion of the washers in the market, by banning the manufacture of the least efficient washer units and changed Energy Star labeling criteria for another portion of the washer market. This change in policy only affected the manufacture of washer units, not unit sales, so the supply of unqualified units was not completely eliminated from the market when the standards were first enacted.

Under standard assumptions of perfect competition and decreasing returns to scale in washer unit production, the policy changes would have two main effects: (1) supply of low-efficiency washers would decrease sharply, and the inward-shifting supply curve would result in a decrease in quantity and an increase in price for less efficient washers (Figure 3.1); (2) demand for high-efficiency washers would increase significantly, and the outward-shifting demand curve would lead to an increase in both quantity and price for more efficient washers (Figure 3.2, Panel A). If, however, washer manufacturers have pricing power, or there are increasing returns to scale in washer production (cost per unit declines with volume), or both, prices for higher efficiency units might fall (Figure 3.2, Panel B). In this case, the downward-sloping supply curve portrays increasing returns to scale in clothes washer production.

Ideally, an estimate of the policy's net effect on consumer welfare would account for demand and competition across all types of washers. Although the ban of low-efficiency washers would generally increase demand for washers not banned, depending on the shape of aggregate preferences, demand for any particular washer may increase or decrease. Further, if there is monopolistic competition, the ban of inefficient washers may cause the market for efficient washers to become more competitive, which could make type-specific demands more elastic. The change in Energy Star threshold could influence consumers' information about relative washer quality, which could further shift individual washer demands.

To account for these many different kinds of possible demand shifts without making strong assumptions about preferences or market structure, we use a rich data set on quantities and prices of every individual model sold, and attempt to estimate the causal effect of the policy change on each one. Although we do not observe costs or washer-specific demand elasticities, changes in prices and quantities can, by themselves, reveal a minimum change in consumers' surplus, regardless of the underlying market structure, as we explain below.

New standards are usually published a couple of years in advance of the effective date; therefore, washer manufacturers may have adjusted their production plan and pricing strategy well before the standards actually take effect. Producers and consumers may rush to make or buy washers about to be banned from manufacture. The transition period could be long because washers are durable goods and washers manufactured prior to the ban may still be sold.

Putting aside issues related to market transition, the net welfare change resulting from a change in standards cannot be fully identified from ex-ante and ex-post prices and quantities of each washer. One key problem is that we cannot identify the slope of demand for units not banned from manufacture. Taking estimated changes in price and quantity as given, we are able to derive the loss in consumer surplus in the banned low-efficiency market, shown as the shaded area of Figure 3.1. This value is identifiable because changes in price and quantity occur along the demand curve, and the slope of the demand curve for low-efficiency washers is revealed by the observed changes in price and quantity. If the price increase is large, it implies the demand curve is relatively steep (more inelastic) and the loss in consumer surplus is expected to be large as well (Figure 3.1, Panel A). On the other hand, if the price increase is small, it implies the demand curve is relatively flat (more elastic), and the loss in consumer surplus is small too (Figure 3.1, Panel B).

What we cannot estimate is the gain in consumer surplus due to the policy induced outward shift in demand for high-efficiency washers (Figure 3.2). Here, changes in price and quantity resulting from a shift in demand occur along the supply curve so the slope of demand is not identifiable. As a result, the most we could measure is how marginal cost changes with quantity sold. The problem is we cannot disentangle the shift in demand from movement along the demand curve, both of which are needed to measure the change in consumer surplus.

Nevertheless, we can see that the smaller the price increase, the more likely there will be a net gain in consumer surplus in the high-efficiency market, as shown in Panel A of Figure 3.2. In this case, the blue trapezoid is likely to be greater than the red trapezoid. If prices actually decrease due to economies of scale (Figure 3.2; Panel B), there are unambiguous consumer gains in the high-efficiency market. We cannot identify how much they gain because we cannot discern how much of the quantity and price changes are due to movement along the demand curve verses a shift in demand. Still we can obtain a lower bound estimate for the gain using the area of the lower blue trapezoid traced out by points $(0, P_0^{he})$, $(0, P_1^{he})$, (Q_1^{he}, P_1^{he}) and (Q_0^{he}, P_0^{he}). This calculation assumes the shift in demand makes up an arbitrarily small share of observed changes

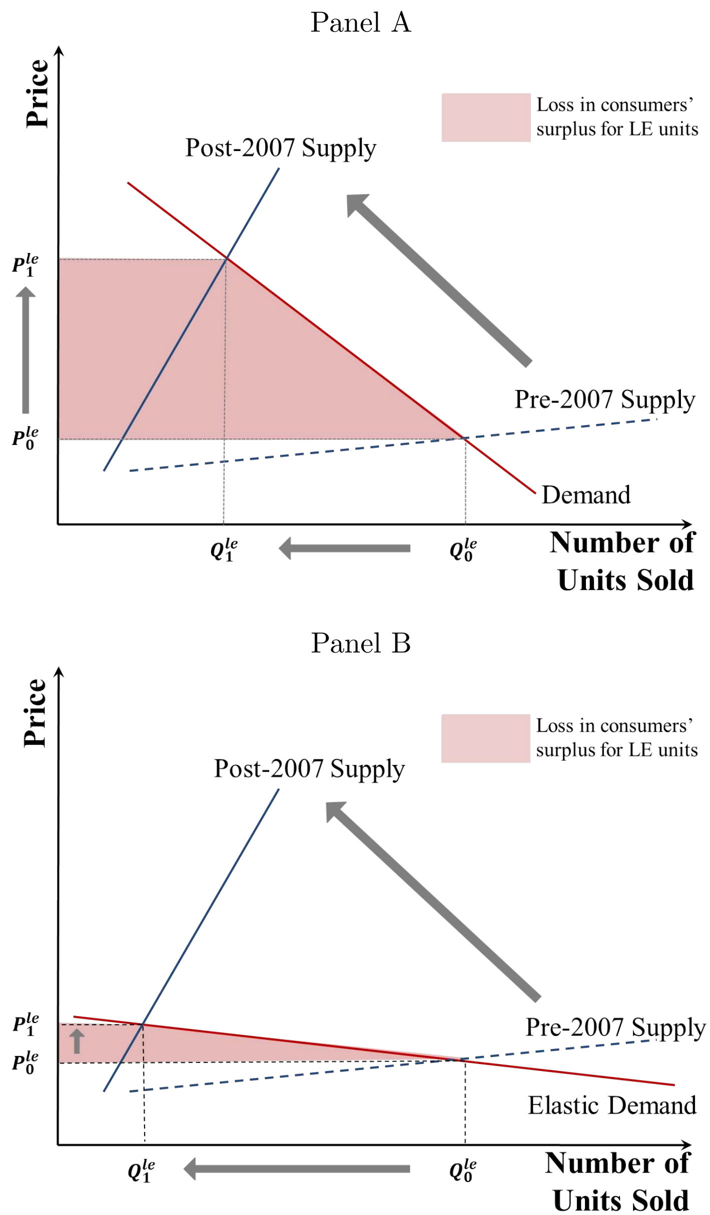

Figure 3.1: The standard change reduces supply of low-efficiency units.

The top panel shows a hypothetical market with supply and demand for low-efficiency washer units. The policy change causes supply curve to shift inward, causing a loss in consumer surplus equal to the shaded area. In the bottom panel, demand for low-efficiency units is flatter and the consumer surplus loss is greater. But we cannot identify the supply curve from policy's effects on price and quantity.

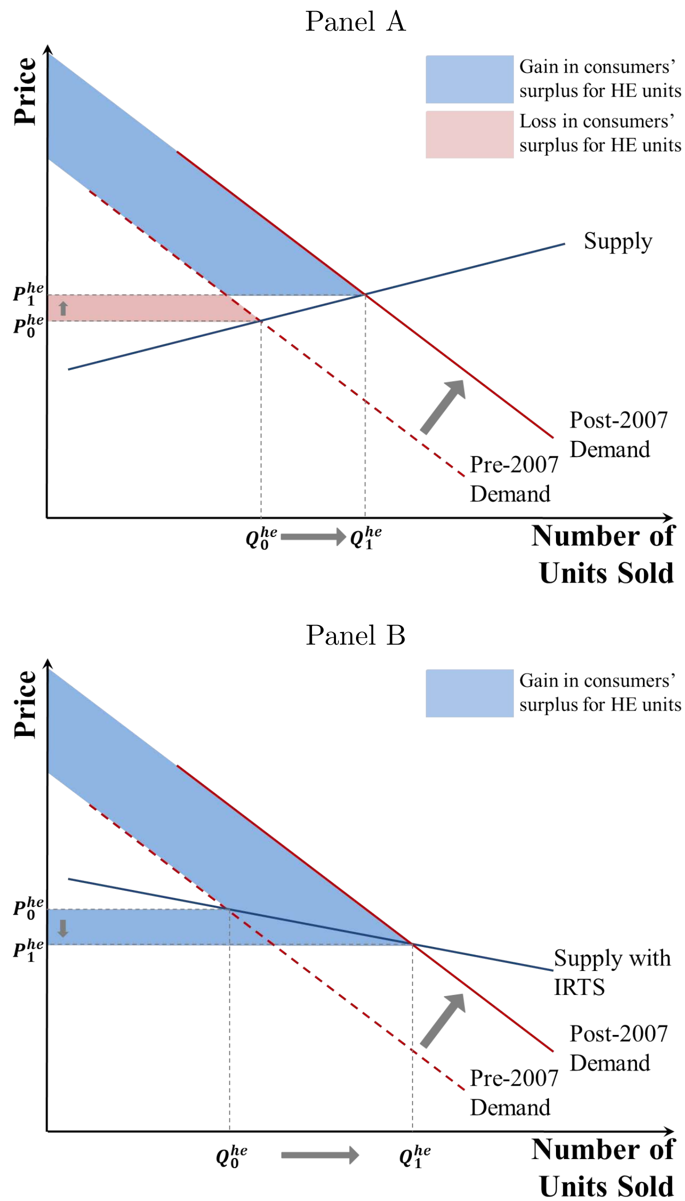

Figure 3.2: The standard change shifts out demand for high-efficiency units.

The top panel shows a hypothetical market with supply and demand for high-efficiency units. The policy change causes prices of low-efficiency units to increase, which causes demand for high-efficiency units to shift out, which increases quantity and price and consumer and producer surpluses. In the bottom panel, supply is downward sloping due to increasing returns to scale, and price falls rather than increasing. We cannot identify the size of surplus benefits because the slope of the demand curve cannot be identified from the policy's effects on price and quantity.

in price and quantity while movement along the demand curve makes up an arbitrarily large share of observed changes. This equals the lower bound of the total consumer surplus change, because it excludes the positive welfare gain from the upper blue trapezoid area in Panel B of Figure 3.2, which is not identifiable since we cannot identify the slope of the demand curve for higher efficiency washers.

We can generalize the lower bound estimate of total change in consumer surplus to account for every washer model sold in the market, as following:

$$\text{Minimum Change in Consumers' Surplus} = \sum_i \frac{1}{2}\left(p_0^i - p_1^i\right)\left(q_0^i + q_1^i\right) \tag{3.1}$$

where p_0^i and p_1^i are prices for washer model i before and after the policy change, and q_0^i and q_1^i are their respective quantities sold.

The equation gives the sum of trapezoid areas outlined by prices and quantities from the pre- and post-standard periods for each washer model. For models with an increase in price after the standard, the corresponding trapezoid area will be assigned with a negative value, implying a loss in consumer surplus, as illustrated by the red-shaded area in Figure 1. For models with a decrease in price after the standard, the associated trapezoid area will be assigned a positive value, which is akin to the lower blue-shaded trapezoid area in Panel B of Figure 2. Given model-specific estimates of changes in prices and quantities, we can therefore construct a lower-bound estimate of the net impact of total consumer surplus due to efficiency standards.

Note that inference about the lower bound of consumer surplus change does not depend on industry structure. Although we have drawn the figure as if the washer market is a perfectly competitive, possibly one with increasing returns to scale, the welfare analysis only concerns consumers' surplus and the structure of demand. It need not be the case that observed price equals marginal cost before or after the policy change. Shrinking producer margins from increased competition may also explain falling prices and increasing quantities.

Another way to view the lower-bound approximation is to draw on ideas akin to Le Chatelier's principle (Samuelson 1947). In banning a portion of the washer market, a worst-case scenario with respect to consumers is to assume extremely limited substitution between washer types. In the extreme case, demand across all submarkets would be fixed: prices and availability of different washers would not shift demands of any other washers. We could then evaluate the change in consumer surplus assuming all changes in prices and quantities were along demand curves and none of the changes came from shifts in demand. If, however, consumers' demands are responsive to prices of substitute washers, then consequent shifts in demand must improve consumer welfare relative to the case where we assume no substitutability. This is precisely the calculation used to calculate the minimum change in consumers' surplus.

3.3 Assessing the Value of Washer Attributes

Some concern has been raised regarding the impact of policy changes on the quality and characteristics of washing machines. Although some attributes of washers; such as energy efficiency, capacity, soil sensors and spin speed; arguably have improved as a result of standards, other features may have become less desirable, like the longer duration of the washing cycle or possibly poorer cleaning performance. Quality, after all, depends on the tastes and preferences of individual buyers. Changes perceived as improvements to some consumers may be seen as detriments to others. The impact of standards on consumer surplus should embody the economic value of these quality changes, which may affect the demand of high-efficiency washers. For example, if new washer designs made available due to the 2007 policy change were particularly undesirable to the majority of consumers, then the relative scarcity of banned washers that were still available after 2007 should have led to a significant increase in old washer prices and this would have resulted in greater consumer loss, as shown in Panel A of Figure 1. On the other hand, if the new washers were close substitutes to the banned washers in terms of washer attributes and functionalities, demand for low-efficiency washers would become relatively elastic, as shown in Panel B of Figure 1.

One might be able to obtain direct engineering measurements of changes in washer quality. Engineering measurements typically apply standardized tests to evaluate washer performance. Engineers, for example, could test whether washers manufactured before the ban tended to last longer without malfunction, test the length of wash, and so on. Consumer surveys could also be used to evaluate consumer preferences towards different washers. Alternatively, values of attributes might be estimated using a hedonic approach that relates attributes of washers to their prices, but this approach must assume preferences are homogeneous. However, we were not able to acquire detailed engineering data or comprehensive consumer reports associated with washer models; therefore, these approaches are not included in this study.

Instead, we use regression models with fixed-effects to estimate washer-specific changes in quantities and prices. Revealed preferences give direct lower-bound measures of consumer surplus without imposing strong assumptions about preferences, transaction costs, or discount rates. This approach does not require underlying assumptions about homogeneous preferences, like the hedonic regression approach. Similarly, it does not require econometric assumptions that are as strong as cross-sectional hedonic regression model approach. However, this approach does assume that consumers generally make rational decisions when purchasing clothes washers so that policy-induced changes in prices and quantities reflect individual values. In the following sections, we detail the data source used in this clothes washer study and the application of the fixed-effect model.

A possible concern with our approach is that many other factors changed concurrently with the 2007 policy change that may have affected washer prices and units sold, including shifts in demand from non-policy factors and technological changes driven by non-policy factors. These factors include changes in energy prices, the collapse of the housing market, the recession that began at the end of 2007, and the subsequent financial crisis and prolonged downturn in the aggregate economy. We attempt to control for these factors as well as possible using standard panel regression techniques and time trends, but they may nevertheless confound the estimated effects of the 2007 policy change on prices and quantities sold. At worst we believe our estimates of changes in consumer surplus are biased too low.

3.4 Data and its Limitations

We obtained point-of-sale (POS) data on total sales of washers purchased in the United States from NPD Group, with measures for units sold and average price broken down by washer brand and model. NPD collects monthly POS data from major household appliance retailers in the U.S. market, including units sold and sales revenue of each model. The data include a limited set of attributes for each model, including washer size, color and type. For a large majority of models, this information was merged with a data set with a Federal Trade Commission (FTC) assessment of energy efficiency measured as kilowatt-hours per year (kWh/y) for standard usage. The FTC measure of energy efficiency cannot be converted to the U.S. Department of Energy's (DOE's) MEF measure, because the FTC measurement does not fully account for the energy used for removing the remaining moisture in the clothes, which reduces subsequent drying time.

If precise MEF data were available and matched to each washer, one approach might be to consider separately four categories of washers: (1) those that met the minimum standard before and after the 2007 change (2) those that met the minimum standard before 2007 but not after; (3) those that met the Energy Star threshold before and after 2007; and (4) units that met the Energy Star threshold before 2007 but not after the policy change. Because we do not have MEF data, we proxy for washers that may have been closer and further from the 2007 minimum and Energy Star standards using the FTC measure of efficiency (kWh/y). We do not have this measure for all units, but we do have it for a significant majority of models in our data. Specifically, we divide units into Energy Star and non-Energy Star, and we further divide each of these subsamples by the median FTC measure of efficiency (kWh/y). This gives us four categories of washers including those that consume less than and more than the median energy use among the Energy Star qualified units and those that consume more than and less than the median energy used among the non-Energy Star qualified units.

We also consider a limited sample of washers that includes only models that were sold both before and after the 2007 policy change. The limited sample facilitates an "apples to apples" comparison of washer prices before and after the 2007 policy changes. Summary statistics for these data, broken down by each sub-group and category, are reported in are reported in table 3.1 and table 3.2.

Table 3.1: Summary Statistics of the Full Sample

	Energy Star				non Energy Star			
	High Efficiency		Low Efficiency		High Efficiency		Low Efficiency	
	pre-2007	post-2007	pre-2007	post-2007	pre-2007	post-2007	pre-2007	post-2007
Units Sold/month	3,314	93,352	34,306	62,851	8,885	61,383	57,937	14,163
Average Prices (2012 $)	948.11	731.16	932.48	725.15	500.13	442.81	398.94	420.30
Std. Deviation of Price	321.64	242.89	324.23	238.22	291.26	151.96	136.35	101.14
Front Loading (%)	100.0	85.6	64.2	34.6	1.2	12.6	0.3	0.0
Top Loading (%)	0.0	14.4	35.7	65.4	98.8	87.4	99.7	100.0
Average FTC (kwh/year)	161.29	140.22	295.41	237.40	390.56	339.35	761.30	474.67
Std. Deviation of FTC	12.62	20.49	137.89	59.38	40.89	96.20	200.38	26.53
Average # of Cycles	7.83	10.32	11.79	13.20	9.57	10.14	8.60	11.49
Std. Deviation of # of Cycles	2.59	3.50	4.74	4.03	3.87	4.29	4.43	5.30
Average RPMs	1136.95	1156.34	965.74	1042.40	999.62	856.74	631.23	757.91
Std. Deviation of RPMs	180.38	179.14	175.12	190.04	266.09	267.24	39.97	110.37
Side by Side (%)	23.7	20.9	54.6	68.1	69.2	81.9	96.3	97.3
Portable Side by Side (%)	0.0	0.0	0.7	0.0	6.3	0.5	0.6	0.0
Stackable (%)	71.2	78.5	40.6	31.5	1.3	10.5	1.2	0.0
Prestacked (%)	0.0	0.2	3.7	0.2	23.0	7.0	1.9	2.7
Portable Stackable (%)	2.2	0.2	0.1	0.0	0.0	0.1	0.0	0.0
Washer/Dryer Combo (%)	2.2	0.1	0.3	0.0	0.2	0.0	0.0	0.0
Portable Washer/Dryer Combo (%)	0.8	0.0	0.1	0.2	0.0	0.0	0.0	0.0
Not Specified (%)	0.0	0.1	0.0	0.0	0.0	0.0	0.0	0.0
Load Sensor (%)	87.2	57.5	33.9	35.8	0.2	4.4	0.1	0.0
Soil Sensor (%)	0.0	7.8	16.7	14.6	0.5	4.9	4.8	14.5
Soil and Load Sensor (%)	8.7	13.4	4.2	28.9	0.9	0.0	0.0	0.0
No Sensor (%)	0.0	19.3	26.4	18.3	94.8	88.3	81.1	85.0
Sensor Type Unspecified (%)	0.0	0.0	9.1	0.0	0.0	0.0	5.1	0.0

Continued on next page

Table 3.1 – continued from previous page

	Energy Star				non Energy Star			
	High Efficiency		Low Efficiency		High Efficiency		Low Efficiency	
	pre-2007	post-2007	pre-2007	post-2007	pre-2007	post-2007	pre-2007	post-2007
Not Specified (%)	4.1	1.9	9.7	2.3	3.6	2.4	9.0	0.5
Average Capacity (Cu. Ft.)	3.45	4.03	3.56	4.02	3.37	3.40	3.13	3.31
Std. Deviation of Capacity	0.47	0.37	0.42	0.38	0.59	0.40	0.36	0.25
Electronic Control (%)	82.3	90.1	70.2	76.2	10.4	12.2	2.5	12.0
Mechanic Control (%)	16.4	9.9	29.5	23.7	89.4	87.8	96.0	88.0
Not Specified (%)	1.3	0.0	0.3	0.0	0.2	0.1	1.4	0.0

Notes: Full sample includes all data with and FTC energy efficiency rating. High and low efficiency values are selected according to the median value of the FTC energy rating taken across model numbers in this particular sample during the period from 2001 to 2011. Note that the data include a varying number of retailers over time and the sample of retailers may not be representative of the population. Monthly data from 2001 - 2011 are used in this analysis. All averages and standard deviations are weighted by total units sold during that period. Characteristics such as # of Cycles, RPMs, and Capacity are recorded as ranges; middle point of the range is used when calculating the average. For example: 1099.5 RPMs is used for those washers with RPMs in the range 1000-1199. For the ranges with one boundary, the boundary is used in the calculation. For example: 2.5 cubic feet is used for those with less than 2.5 cubic feet and 4.5 is used for those with 4.5+ cubic feet. Data with "Not Specified" characteristics are treated as missing data and not included in the calculation of the characteristic average.

Table 3.2: Summary Statistics of the Limited Sample

	Energy Star				non Energy Star			
	High Efficiency		Low Efficiency		High Efficiency		Low Efficiency	
	pre-2007	post-2007	pre-2007	post-2007	pre-2007	post-2007	pre-2007	post-2007
Units Sold/month	12,464	14,974	24,413	8,523	18,246	11,779	23,854	1,260
Average Prices (2012 $)	989.65	805.63	847.33	822.92	406.07	380.61	376.23	535.39
Std. Deviation of Price	218.45	268.27	327.73	173.28	213.33	99.52	113.10	286.01
Front Loading (%)	87.2	90.5	49.4	25.2	0.1	0.0	0.0	0.0
Top Loading (%)	12.8	9.5	50.6	74.8	99.9	100.0	100.0	100.0
Average FTC (kwh/year)	189.47	169.87	300.94	321.58	409.57	392.28	547.15	519.92
Std. Deviation of FTC	18.77	23.29	78.53	72.16	35.38	54.77	35.49	51.07
Average # of Cycles	8.52	8.63	13.57	11.75	7.85	7.71	10.65	9.15
Std. Deviation of # of Cycles	3.09	2.73	4.51	3.30	3.21	4.18	5.12	4.83
Average RPMs	1091.64	1128.71	928.68	1025.57	1006.40	664.93	626.54	625.96
Std. Deviation of RPMs	111.21	139.16	195.20	132.08	296.24	151.78	9.89	8.41
Side by Side (%)	23.2	38.1	53.4	76.3	73.0	96.6	98.4	70.6
Portable Side by Side (%)	0.0	0.0	0.0	0.0	5.1	1.5	0.2	0.0
Stackable (%)	75.6	59.8	42.4	22.4	0.0	0.0	0.0	0.0
Prestacked (%)	0.4	0.1	3.8	1.1	21.8	1.9	1.3	29.4
Portable Stackable (%)	0.5	1.4	0.0	0.0	0.0	0.0	0.0	0.0
Washer/Dryer Combo (%)	0.0	0.0	0.4	0.3	0.0	0.0	0.0	0.0
Portable Washer/Dryer Combo (%)	0.3	0.7	0.0	0.0	0.0	0.0	0.0	0.0
Load Sensor (%)	62.6	63.3	36.1	66.7	0.2	0.1	0.2	0.0
Soil Sensor (%)	10.3	4.1	27.6	7.6	0.5	16.7	13.6	6.9
Soil and Load Sensor (%)	9.0	21.9	3.4	17.3	0.0	0.0	0.0	0.0
No Sensor (%)	17.6	10.6	23.9	6.3	98.0	82.5	78.3	90.0
Sensor Type Unspecified (%)	0.3	0.0	0.2	0.0	0.0	0.0	0.0	0.0
Not Specified (%)	0.2	0.0	8.8	2.1	1.4	0.6	7.9	3.0

Continued on next page

Table 3.2 – continued from previous page

	Energy Star				non Energy Star			
	High Efficiency		Low Efficiency		High Efficiency		Low Efficiency	
	pre-2007	post-2007	pre-2007	post-2007	pre-2007	post-2007	pre-2007	post-2007
Average Capacity (Cu. Ft.)	3.60	3.55	3.53	4.01	3.23	3.20	3.25	3.08
Std. Deviation of Capacity	0.27	0.37	0.31	0.44	0.62	0.31	0.18	0.41
Electronic Control (%)	93.4	94.0	59.4	86.7	0.7	0.4	0.3	0.0
Mechanic Control (%)	6.6	6.0	40.6	13.3	99.3	99.6	99.7	100.0

Notes: Limited sample is a subset of the full sample that includes all washer models observed in the data at least once before and after the 2007 policy change data. High and low efficiency values are selected according to the median value of the FTC energy rating taken across model numbers in this particular (limited) sample during the period from 2004 to 2011. Note that the data include a varying number of retailers over time and the sample of retailers may not be representative of the population. Monthly data from 2001 - 2011 are used in this analysis. All averages and standard deviations are weighted by total units sold during that period. Characteristics such as # of Cycles, RPMs, and Capacity are recorded as ranges; middle point of the range is used when calculating the average. For example: 1099.5 RPMs is used for those washers with RPMs in the range 1000-1199. For the ranges with one boundary, the boundary is used in the calculation. For example: 2.5 cubic feet is used for those with less than 2.5 cubic feet and 4.5 is used for those with 4.5+ cubic feet. Data with "Not Specified" characteristics are treated as missing data and not included in the calculation of the characteristic average.

3.5 Price and Quantity Changes Over Time

The data described above are plotted in Figure 3.3 through Figure 3.8. The figures show prices, units sold, and shares of units sold for each of the four categories of washer, for both the full sample and the limited sample that only includes units sold both before and after the 2007 policy change. All prices are adjusted for inflation using the Consumer Price Index and reported in 2012 dollars. The average prices for different categories across time (individual washers weighted by quantity sold) are plotted in the figures with the shaded band showing +/- 1 standard deviation; and the weighted average price taken across all units sold in the sample are also shown for comparison. The data appear to show a slight discontinuity in quantities and prices around the time of the policy change. Prices decline for high-efficiency units and prices increase slightly for low-efficiency units. Note that all prices are in current dollars. The upward trend in the market share of more efficient washers, with a strong shift toward Energy Star units, started before the 2007 policy change, and the trend continued well after the amended standards. The longer-run trends are obviously difficult to disentangle from other factors influencing technological advance and demand. Despite the marked shift in purchasing behavior toward more-efficient units, real average prices show little if any trend upward, even in the full sample. As the more-efficient units appear to be falling in price faster than less-efficient units, the concurrent trends in prices and quantities provide some evidence of scale economies.

3.6 Regression Models

Fixed-effect models are used to explore the relationship between dependent variables and independent variables that are both time-variant and have within-individual variation. Here we attempt to discern model-specific effects of the policy on price and quantity of units sold using a fixed-effect, regression-discontinuity framework. This regression includes fixed effects for each washer model and smooth, or spline, functions of time to account for other factors changing around the time of the policy. Modeling fixed effects removes all model-specific features of washers so that price effects cannot be attributed to changes in quality or characteristics over time. With this type of analysis, we attempt to explain how individual model prices and quantities change over time, and whether these changes were unusually sharp around the 2007 policy change.

The effect of the policy itself is estimated with a dummy variable to capture effects around the time of the policy change under the maintained assumption that the change itself is exogenous conditional on a flexible time-trend control. We also included an interaction term between the 2007 policy dummy variable and the energy efficiency level, measured in kWh per year, because this allows us to control for unobserved time-varying effects related to energy efficiency

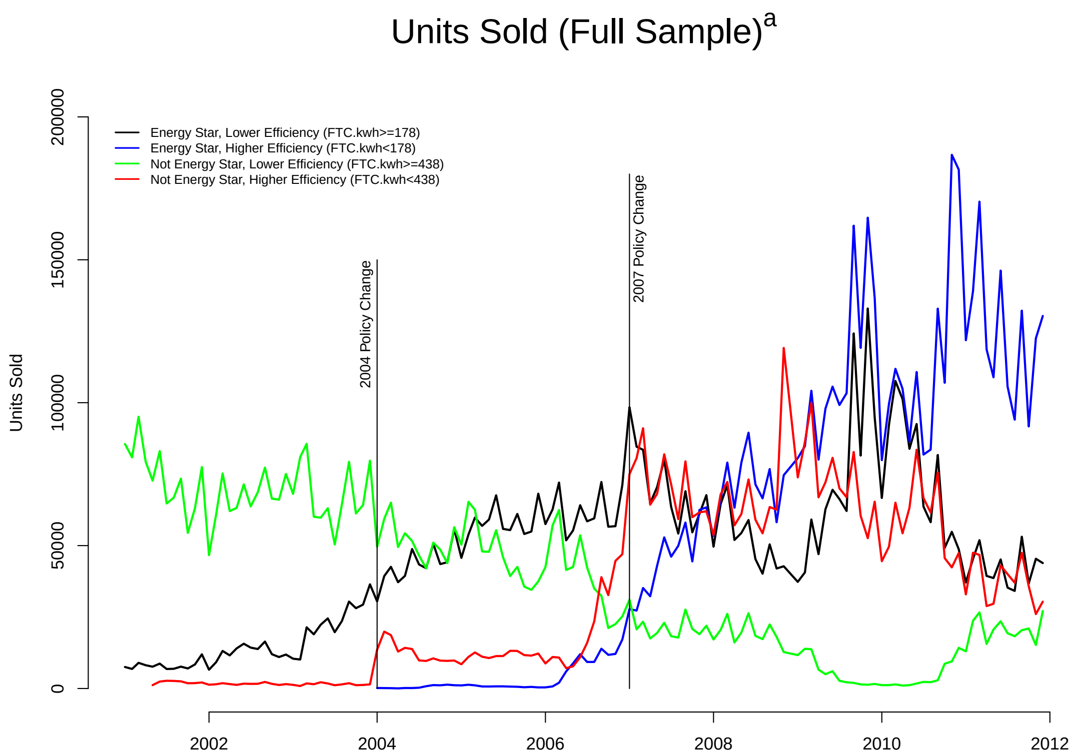

Figure 3.3: Washer Units Sold — Full Sample

[a]Full sample includes all data with and FTC energy efficiency rating. High and low efficiency values are selected according to the median value of the FTC energy rating taken across model numbers in this particular sample during the period from 2001 to 2011. Note that the data include a varying number of retailers over time and the sample of retailers may not be representative of the population.

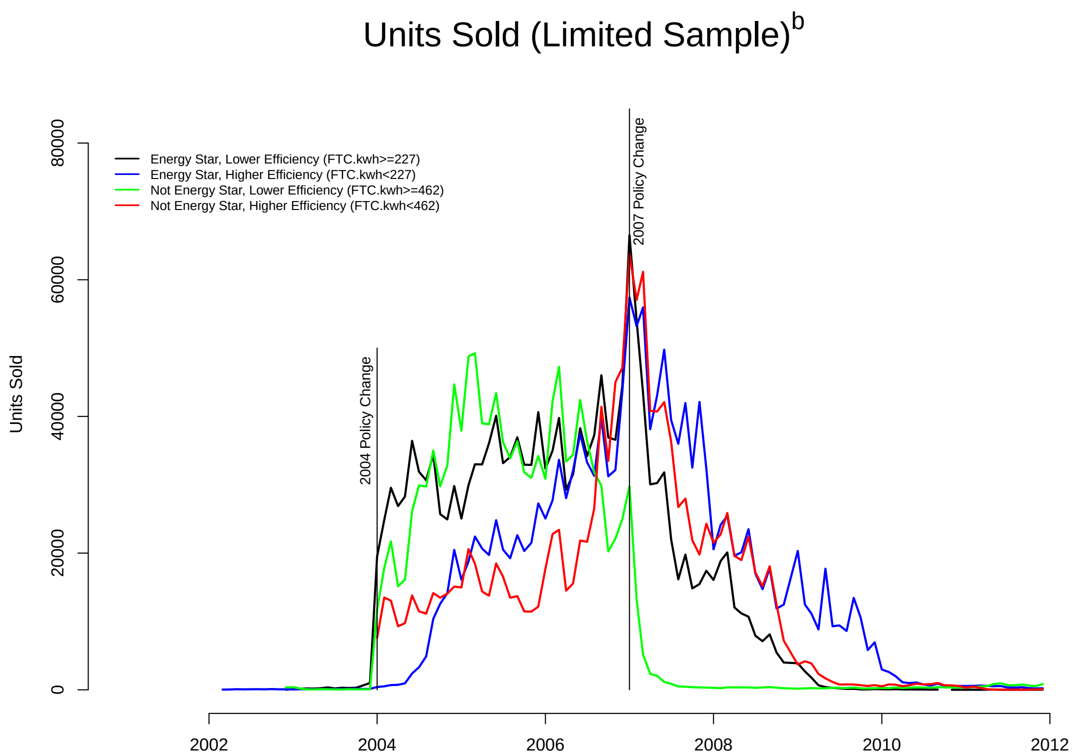

Figure 3.4: Washer Units Sold — Limited Sample (models existed before and after 2007)

[b] Limited sample is a subset of the full sample that includes all washer models observed in the data at least once before and once after the 2007 policy change data. High and low efficiency values are selected according to the median value of the FTC energy rating taken across model numbers in this particular (limited) sample during the period from 2004 to 2011. Note that the data include a varying number of retailers over time and the sample of retailers may not be representative of the population.

Washer Prices (Full Sample)[a]

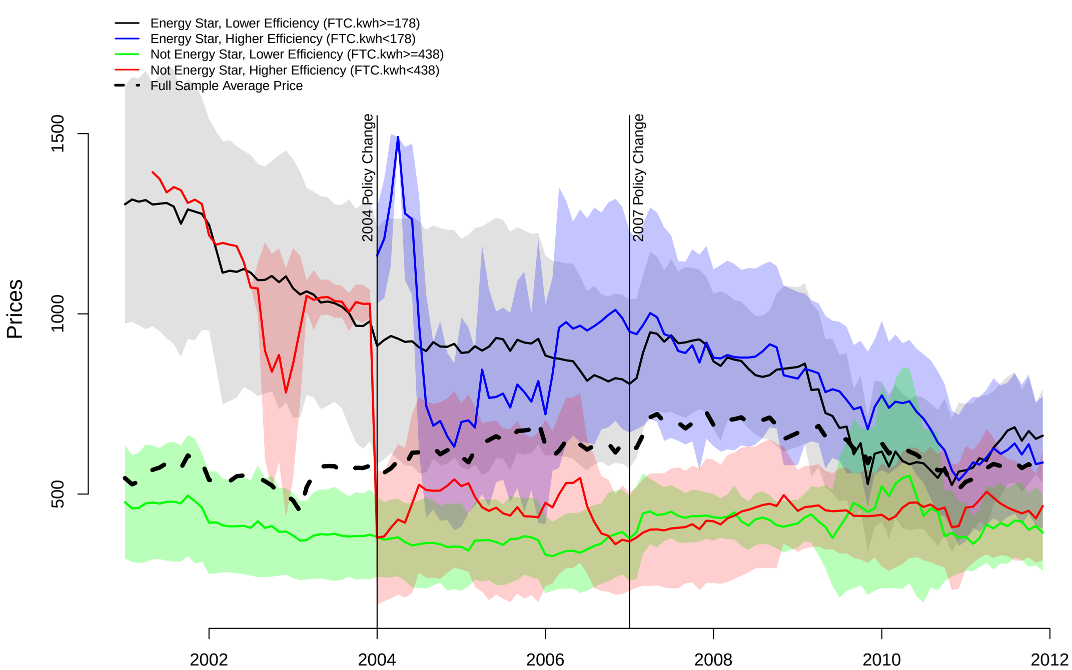

Figure 3.5: Average Washer Prices — Full Sample

[a]Full sample includes all data with and FTC energy efficiency rating. High and low efficiency values are selected according to the mean value of the FTC energy rating taken across model numbers in this particular sample during the period from 2001 to 2011. Note that the data include a varying number of retailers over time and the sample of retailers may not be representative of the population. The lines for each category of washer indicate the value for the mean washer sold in that category, and the band around the line marks +/- 1 standard deviation. The thick dashed line is the average price across all models, with models weighted by number of units sold.

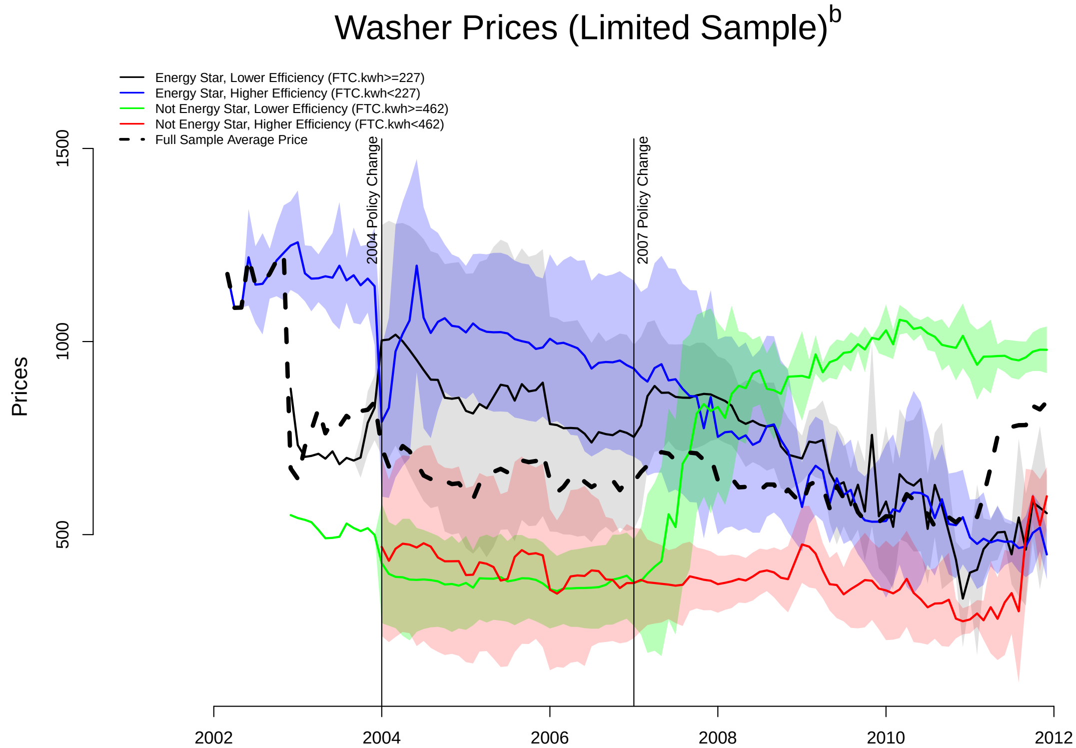

Figure 3.6: Average Washer Prices — Limited Sample (models existed before and after 2007)

[b]Limited sample is a subset of the full sample that includes all washer models observed in the data at least once before and after the 2007 policy change data. High and low efficiency values are selected according to the mean value of the FTC energy rating taken across model numbers in this particular (limited) sample during the period from 2004 to 2011. Note that the data include a varying number of retailers over time and the sample of retailers may not be representative of the population. The lines for each category of washer indicate the value for the mean washer sold in that category, and the band around the line marks +/- 1 standard deviation. The thick dashed line is the average price across all models, with models weighted by number of units sold.

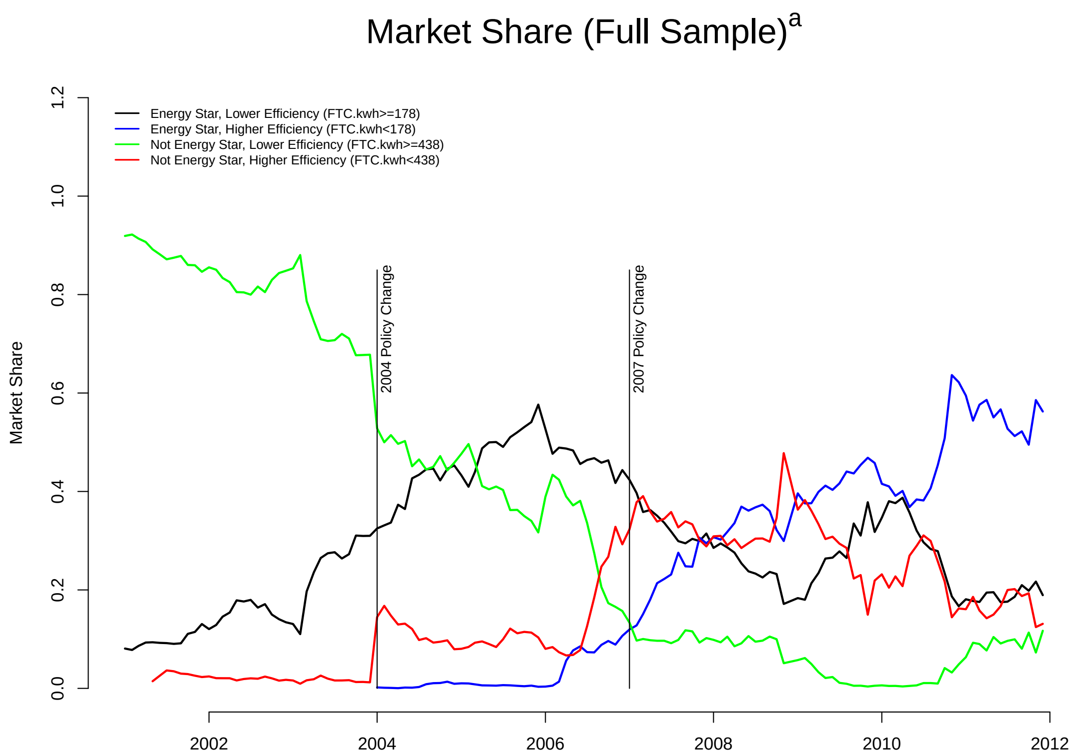

Figure 3.7: Market Share — Full Sample

[a]Full sample includes all data with and FTC energy efficiency rating. High and low efficiency values are selected according to the median value of the FTC energy rating taken across model numbers in this particular sample during the period from 2001 to 2011. Note that the data include a varying number of retailers over time and the sample of retailers may not be representative of the population.

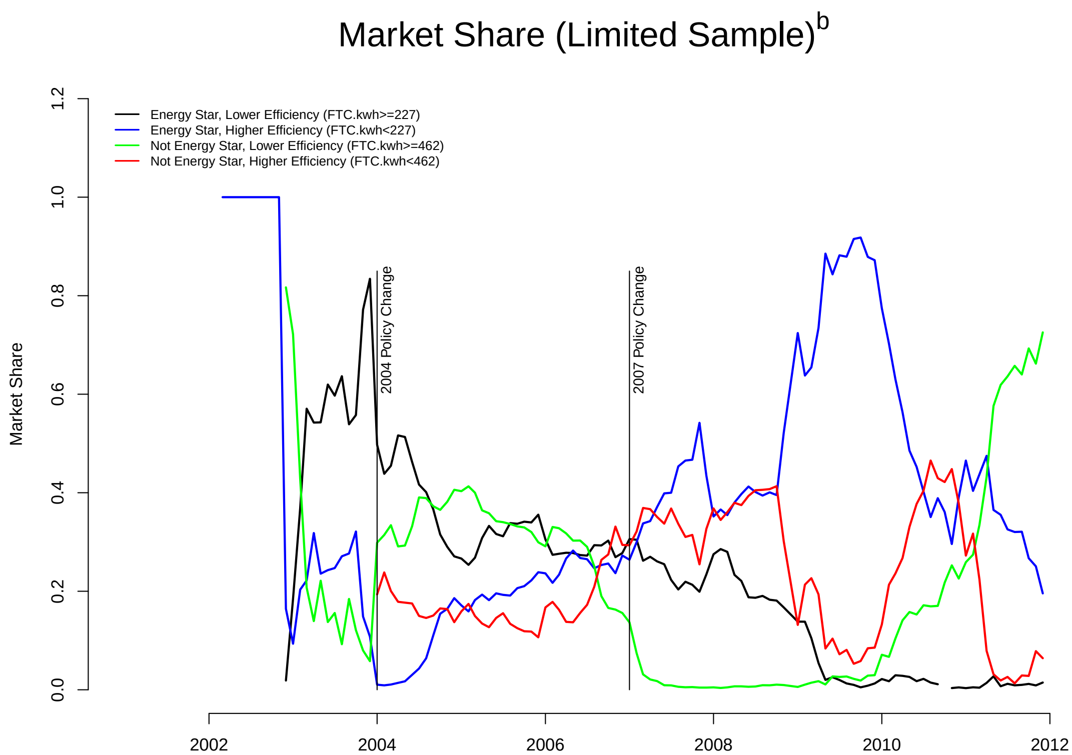

Figure 3.8: Market Share — Limited Sample (models existed before and after 2007)

[b] Limited sample is a subset of the full sample that includes all washer models observed in the data at least once before and after the 2007 policy change data. High and low efficiency values are selected according to the median value of the FTC energy rating taken across model numbers in this particular (limited) sample during the period from 2004 to 2011. Note that the data include a varying number of retailers over time and the sample of retailers may not be representative of the population.

and gain insights into the variations in price effects by efficiency level. It is important to note that the regression framework - akin to a regression discontinuity - is not especially well suited to this particular policy change due to the factors described in the opening section: (a) the change was anticipated in advance and likely influenced behavior long before and long after the event of the change, and (b) many other factors were likely influencing the appliance industry around the time of the policy change. As we detail in the discussion below, these factors suggest that our estimated minimum change in consumers' surplus may be biased too low. A spline function with four knots was added to the model to impose continuity restrictions at the joint points and force the regression to have a smoother transition when an exogenous shock takes place.

The baseline regression models are specified as follows:

$$q_{it} = \alpha_i + f(t) + \beta_q d2007_t + \gamma_q(d2007 \times kWh/y_i) + \epsilon_{it} \tag{3.2}$$

$$p_{it} = \alpha_i + f(t) + \beta_p d2007_t + \gamma_p(d2007_t \times kWh/y_i) + \epsilon_{it} \tag{3.3}$$

where $d2007_t$ is a dummy variable indicating the policy change in 2007, $f(t)$ is a spline function of time, and kWh_i is energy efficiency measured in kilowatt hours per year for each washer model.

We also estimated models with more robust controls for time, achieved by interacting smooth functions of time with each of the four categories of washers described in section 3.4, denoted as $f_c(t)$. The models with trend interactions are:

$$q_{it} = \alpha_i + f_c(t) + \beta_q d2007_t + \gamma_q(d2007_t \times kWh/y_i) + \epsilon_{it} \tag{3.4}$$

$$p_{it} = \alpha_i + f_c(t) + \beta_p d2007_t + \gamma_p(d2007_t \times kWh/y_i) + \epsilon_{it} \tag{3.5}$$

All regressions are estimated using both the full sample of washers and the limited sample of washers with transactions observed both before and after the 2007 policy change. To avoid mixing the effects of an early policy change in 2004, we use the data from 2004 to 2011 in this session; and study the effect of 2004 policy change using data from 2001 to 2006 (results shown in section 3.9.1 in the Appendix). The smooth functions of time are estimated using a natural cubic spline, each with four knots and three degrees of freedom.

Results from these regressions are reported in table 3.3. The results suggest the policy change caused prices for most washers to fall and quantities to increase. The positive sign of the coefficient for the interaction term in the price model implies that the price of high-efficiency washers would decrease faster compared to less-efficient washers, and the effect reverses for low-efficiency washers. The sign reverses for the most inefficient units. Analogously, the negative sign of the coefficient for the interaction term in the quantity model implies that the quantity sold of high-efficiency washers increased faster compared to low-efficiency washers, and the effect reverses for low-efficiency washers. The most inefficient units with positive price effects and negative quantity effects were likely banned as a result of the policy change, as these units vanish from the data set, or nearly so, in later months.

Table 3.3: Summary of Regressions

	Price Effect of Policy			Quantity Effect of Policy		
	β	γ	Adj. R^2	β	γ	Adj. R^2
			Estimate/(t-statistic)			
Basic — Full Sample	-264.97	0.57	0.80	399.17	-1.19	0.37
	(-22.29)	(18.24)		(3.58)	(-4.05)	
Robust — Full Sample	-101.70	0.13	0.82	62.46	-0.21	0.38
	(-7.58)	(3.57)		(0.48)	(-0.62)	
Basic — Limited Sample	-224.10	0.55	0.90	931.92	-1.41	0.46
	(-20.01)	(21.90)		(7.51)	(-5.06)	
Robust — Limited Sample	-52.09	0.06	0.92	929.91	-1.52	0.48
	(-4.08)	(1.74)		(5.96)	(-3.86)	

Notes: All regression models include fixed effects for each model of washing machine. *Basic* models include a natural spline of time (on a monthly scale) with four knots. *Robust* models include a separate natural spline of time for each of the four classes of washer as classified in the data section and plotted in the above figures. The β coefficient is for the dummy variable indicating the post-2007 time period. The γ coefficient is the effect on the interaction of the year-2007 dummy variable and the unit's FTC energy-efficiency rating, measured in kWh of energy use per year.

3.7 Minimum Change in Consumers' Surplus

In section 3.2 above we described an estimate for the minimum loss in consumers' surplus as:

$$\text{Minimum Change in Consumers' Surplus} = \sum_i \frac{1}{2}\left(p_0^i - p_1^i\right)\left(q_0^i + q_1^i\right) \tag{3.1}$$

To derive this value from the estimated regressions, we define baseline quantities and prices and then infer the policy-induced difference in quantities and prices from the data and regression coefficients. Although the change in prices and quantities is clearly defined by the regression coefficients, there are several ways to select baseline prices and quantities, given the fact that some washer models entered and exited the database at least in part due to the policy. We therefore consider three baseline prices and quantities using: (a) average value shortly before and after the 2007 policy change; (b) average ex-ante prices and quantities; and (c) average ex-post prices and quantities. We added the estimated difference derived from the regression to find ex-post prices and quantities and subtracted the estimated difference derived from the regression to find ex-ante prices and quantities.

Specifically, we define:

$$\Delta p^i = \hat{\beta}_p + \hat{\gamma}_p kWh/y_i \tag{3.6}$$

$$\Delta q^i = \hat{\beta}_q + \hat{\gamma}_q kWh/y_i \tag{3.7}$$

and consider three alternative ways to define ex-ante and ex-post prices (p)and quantities (q).

(a) 2006-2007 baseline: $x_0^i = \overline{x}^i_{2006-07} - \frac{1}{2}\Delta x^i$ $\quad x_1^i = \overline{x}^i_{2006-07} + \frac{1}{2}\Delta x^i$

(b) 2004-2006 baseline: $x_0^i = \overline{x}^i_{2004-06}$ $\quad x_1^i = \overline{x}^i_{2004-06} + \Delta x^i$

(c) 2007-2011 baseline: $x_0^i = \overline{x}^i_{2007-11} - \Delta x^i$ $\quad x_1^i = \overline{x}^i_{2007-11}$

where $x \in \{p, q\}$ and $\overline{x}^i_{200x-0y}$ indicates the average monthly (quantity q or price p) in years $200x$ to $200y$. Because each washer entered and exited the market at different periods of time, we modify the price/quantity estimation slightly to accommodate washers that were sold only before or only after the 2007 policy change. For example, if washer models were introduced to the market after 2007, the ex-ante price (p_0^i) estimate would be zero under scenario (b), and the ex-post price could be expressed as max $\left\{0, p_0^i + \Delta p^i\right\}$. Thus, all derived prices and quantities are truncated at zero, which assume free disposal and existence of a choke price above which no sales occur.

In Table 3.4 we report lower-bound changes in consumer surplus, as measured in dollars per month, under the three baseline scenarios, two model specifications, and two sample sizes as described in the previous sections. The values differ in accordance with the data sample used (limited or full sample), the regression model from which changes in quantities and prices were derived (basic or robust regression), and the baseline prices and quantities used.

3.8 Discussion

We can estimate some, but not all, welfare effects of the 2007 changes in energy efficiency and Energy Star standards using panel regression techniques and a regression discontinuity design that exploits changes in prices and quantity that occur around the time of the policy change. Our analysis suggests that the standard changes reduced prices and increased quantities of more-efficient units, and reduced quantities and slightly increased prices of less-efficient units, some of which were banned from manufacture in 2007 but still sold after 2007.

Because the policy change would have increased demand for the more-efficient units, the

Table 3.4: Estimates of Minimum Change in Consumer Surplus

Baseline Price and Quantity	Full Sample Basic Model	Full Sample Robust Model	Limited Sample Basic Model	Limited Sample Robust Model
	Dollars Per Month/ (Percent of Total Sales)			
$\bar{p}$ & $\bar{q}$ equal average	16,210,650	12,243,790	5,675,561	4,708,101
of 2006 & 2007	(8.78)	(6.63)	(5.69)	(4.72)
p_0 & q_0 equal average	15,908,970	8,683,736	6,868,341	6,966,556
of 2004-2006	(4.98)	(2.72)	(6.02)	(6.10)
p_1 & q_1 equal average	35,680,760	25,002,440	1,820,717	1,621,604
of 2007-2011	(14.59)	(10.22)	(4.64)	(4.13)

Notes: The table reports lower bound changes in consumer surplus as described in the text, measured in dollars per month. The values differ according the sample of data used (limited or full), the regression model from which changes in quantities and prices were derived (basic or robust), and the baseline prices and quantities used. The numbers in parenthesis give the minimum change in surplus as a percent of average sales, where the denominator is $\frac{1}{2}(p_0q_0 + p_1q_1)$ and p_0, q_0, p_1 and q_1 are as defined in the text.

observed fall in prices would seem possible under these scenarios: (a) there are significant economies of scale in production, which caused average production costs to decline with higher sales; (b) washer markets are imperfectly competitive, and the policy change caused not just an outward shift in demand for higher efficiency units, but a change in the elasticity of demand for higher efficiency units. In the latter case, it is easy to imagine that the new demand for higher efficiency units would be more elastic to price than the previous demand would have been. Specifically, the new much larger pool of high-efficiency shoppers may have preferences that are less brand specific or feature specific than buyers of high-efficiency washers before the policy change. A more elastic demand could incite monopolistically competitive producers to lower prices. Of course, a combination of these two explanations may also account for the changes we observed in the more efficient washer market.

Other factors likely changing around the time of the policy change might bias these inferences. The most logical alternative explanation for the decline in prices is the rapid slowdown in the housing market and the aggregate economy, which also took place around this time. The resulting downward shift in demand could have encouraged firms to generally lower prices. This explanation, however, cannot reconcile the strong shift in the share of sales toward more-efficient washers. In general, we see larger growth in sales and larger declines in prices for higher efficiency units as compared to less-efficient units. This pattern cannot be easily reconciled by a general inward shift in demand due to a slowing housing market or broader economy.

Another possibility is that a technological advance in the production of relatively more-efficient machines occurred simultaneously with the policy. If this were the case, then the policy may not have *caused* the decline in price, but the advance would nevertheless have made the welfare costs of the policy change much smaller. Yet another possible explanation for the increasing demand is that rising energy prices in 2007 and 2008 precipitated higher demand for high-efficiency washers. But even if energy prices were more responsible than the policy change for the shift in demand away from low-efficiency units and toward higher efficiency units, the observed decline in washer prices would only seem possible if there were significant economies of scale, changes in strategic pricing, or both by monopolistically competitive firms. Thus, the effect of the policy would be much the same as the change in energy prices.

We have attempted to control for all these possible confounding factors described previously by using a fixed-effect model that incorporates a flexible time-trend function. The market would usually take some time to adjust and react to the anticipated policy change; hence, the time trends may be capturing effects of the policy, not just potential confounding factors. Thus, even our lower-bound estimated welfare impacts may be conservative. Indeed, when robust washer-type-specific trends are included, estimated minimum consumer surplus gains decline sharply, especially in the full sample. Yet even with these more conservative estimates we find consumer gains to be equal to at least 3 percent of total sales.

In general, these results are difficult to reconcile with the idea that the policy change caused significant welfare losses and are consistent with plausible scenarios in which the policy change improved consumer welfare. However, due to the basic challenges of identification described previously, it is not possible to use these data to estimate the impact of the policy change on firm profits.

APPENDIX

3.9 Appendix

3.9.1 Analysis of 2004 Policy Change

Table (3.5) summarizes the features of washer models that exist both before and after 2004 policy change during the period from 2001 to 2011. The four categories are defined the same way as in the analysis of 2007, but using median values from 2001 to 2006 to avoid the effects of 2007 policy change.

Table 3.5: Summary Statistics of the Limited Sample

	Energy Star				non Energy Star			
	High Efficiency		Low Efficiency		High Efficiency		Low Efficiency	
	pre-2004	post-2004	pre-2004	post-2004	pre-2004	post-2004	pre-2004	post-2004
Units Sold/month	619	1,140	2,946	1,204	3,011	1,380	16,878	374
Average Prices (2012 $)	828.54	719.86	1106.34	890.60	448.90	386.02	426.28	375.58
Std. Deviation of Price	176.43	143.95	515.19	300.21	173.34	101.44	116.31	102.78
Front Loading (%)	85.4	13.5	42.4	1.3	10.8	0.3	0.0	0.0
Top Loading (%)	14.3	86.5	56.8	98.7	89.2	99.7	100.0	100.0
Average FTC (kwh/year)	254.81	209.85	608.46	394.64	808.58	572.84	966.70	964.69
Std. Deviation of FTC	21.11	22.81	303.70	81.43	173.81	65.91	8.19	9.13
Average # of Cycles	11.56	5.81	12.82	9.74	5.12	8.01	8.91	8.75
Std. Deviation of # of Cycles	3.01	2.61	4.93	1.65	3.60	2.86	3.27	1.69
Average RPMs	932.20	1106.86	1010.83	648.79	665.58	624.50	655.13	844.52
Std. Deviation of RPMs	139.54	39.61	152.79	72.16	140.99	0.00	77.16	33.09
Side by Side (%)	26.3	87.9	68.8	5.5	95.8	97.6	96.5	92.6
Portable Side by Side (%)	72.9	0.3	0.0	0.0	0.0	0.0	0.0	0.0
Stackable (%)	0.8	0.4	0.0	0.3	2.0	0.1	0.0	0.0
Prestacked (%)	0.0	0.0	31.2	94.2	1.0	2.3	3.5	7.4
Portable Stackable (%)	0.0	9.5	0.0	0.0	0.0	0.0	0.0	0.0
Washer/Dryer Combo (%)	0.0	1.8	0.0	0.0	1.2	0.0	0.0	0.0
Soil and Load Sensor (%)	0.0	9.5	0.0	0.0	0.0	0.0	0.0	0.0
No Sensor (%)	73.1	86.5	31.2	87.7	48.8	1.3	89.9	0.0
Sensor Type Unspecified (%)	21.5	0.0	55.0	0.0	46.0	0.0	1.2	0.0
Not Specified (%)	5.3	3.9	13.8	12.3	5.2	98.7	8.9	100.0
Average Capacity (Cu. Ft.)	2.71	3.56	4.29	3.21	3.17	3.21	3.36	3.82
Std. Deviation of Capacity	0.60	0.39	1.17	0.15	0.78	0.33	0.63	0.65

Continued on next page

Table 3.5 – continued from previous page

	Energy Star				non Energy Star			
	High Efficiency		Low Efficiency		High Efficiency		Low Efficiency	
	pre-2004	post-2004	pre-2004	post-2004	pre-2004	post-2004	pre-2004	post-2004
Electronic Control (%)	26.1	97.5	83.5	1.8	45.4	78.7	4.6	8.0
Mechanic Control (%)	73.6	0.7	16.2	98.2	53.4	21.3	95.4	92.0
Not Specified (%)	0.3	1.8	0.3	0.0	1.2	0.0	0.0	0.0

Notes: Limited sample is a subset of the full sample that includes all washer models observed in the data at least once before and after the 2004 policy change data. High and low efficiency values are selected according to the median value of the FTC energy rating taken across model numbers in this particular (limited) sample during the period from 2001 to 2006. Note that the data include a varying number of retailers over time and the sample of retailers may not be representative of the population. Monthly data from 2001 - 2011 are used in this analysis. All averages and standard deviations are weighted by total units sold during that period. Characteristics such as # of Cycles, RPMs, and Capacity are recorded as ranges; middle point of the range is used when calculating the average. For example: 1099.5 RPMs is used for those washers with RPMs in the range 1000-1199. For the ranges with one boundary, the boundary is used in the calculation. For example: 2.5 cubic feet is used for those with less than 2.5 cubic feet and 4.5 is used for those with 4.5+ cubic feet. Data with "Not Specified" characteristics are treated as missing data and not included in the calculation of the characteristic average.

We also plotted prices, unit sold and market share for the limited 2004 sample in Figure 3.9 to Figure 3.11.

The baseline regression models are specified as follows:

$$q_{it} = \alpha_i + f(t) + \beta_q d2004_t + \gamma_q (d2004 \times kWh/y_i) + \epsilon_{it} \tag{3.8}$$

$$p_{it} = \alpha_i + f(t) + \beta_p d2004_t + \gamma_p (d2004_t \times kWh/y_i) + \epsilon_{it} \tag{3.9}$$

And the robust models are:

$$q_{it} = \alpha_i + f_c(t) + \beta_q d2007_t + \gamma_q (d2007_t \times kWh/y_i) + \epsilon_{it} \tag{3.10}$$

$$p_{it} = \alpha_i + f_c(t) + \beta_p d2007_t + \gamma_p (d2007_t \times kWh/y_i) + \epsilon_{it} \tag{3.11}$$

where $d2004_t$ is a dummy variable indicating the policy change in 2004, $f(t)$ is a spline function of time, and kWh_i is energy efficiency measured in kilowatt hours per year for each washer model.

Table 3.6 shows the results of the regression model, and the policy change in 2004 had very similar but smaller effects in the market as the one in 2007 causing prices to fall and quantity to increase.

Units Sold (Limited Sample)[c]

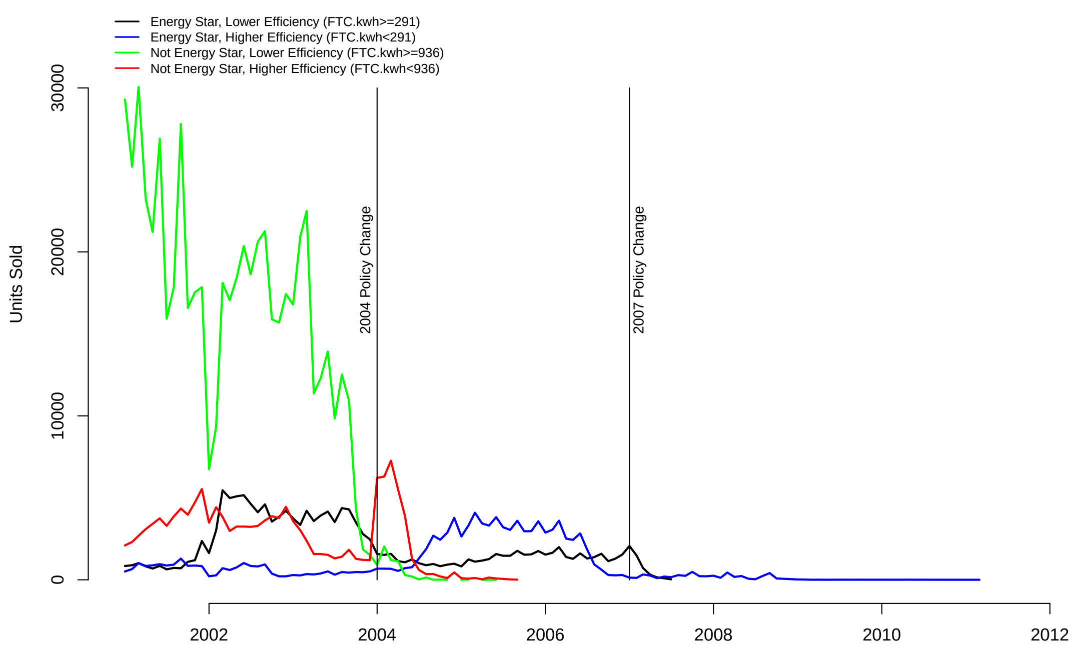

Figure 3.9: Washer Units Sold — Limited Sample (models existed before and after 2004)

[c]Limited sample is a subset of the full sample that includes all washer models observed in the data at least once before and once after the 2004 policy change data. High and low efficiency values are selected according to the median value of the FTC energy rating taken across model numbers in this particular (limited) sample during the period from 2001 to 2006. Note that the data include a varying number of retailers over time and the sample of retailers may not be representative of the population.

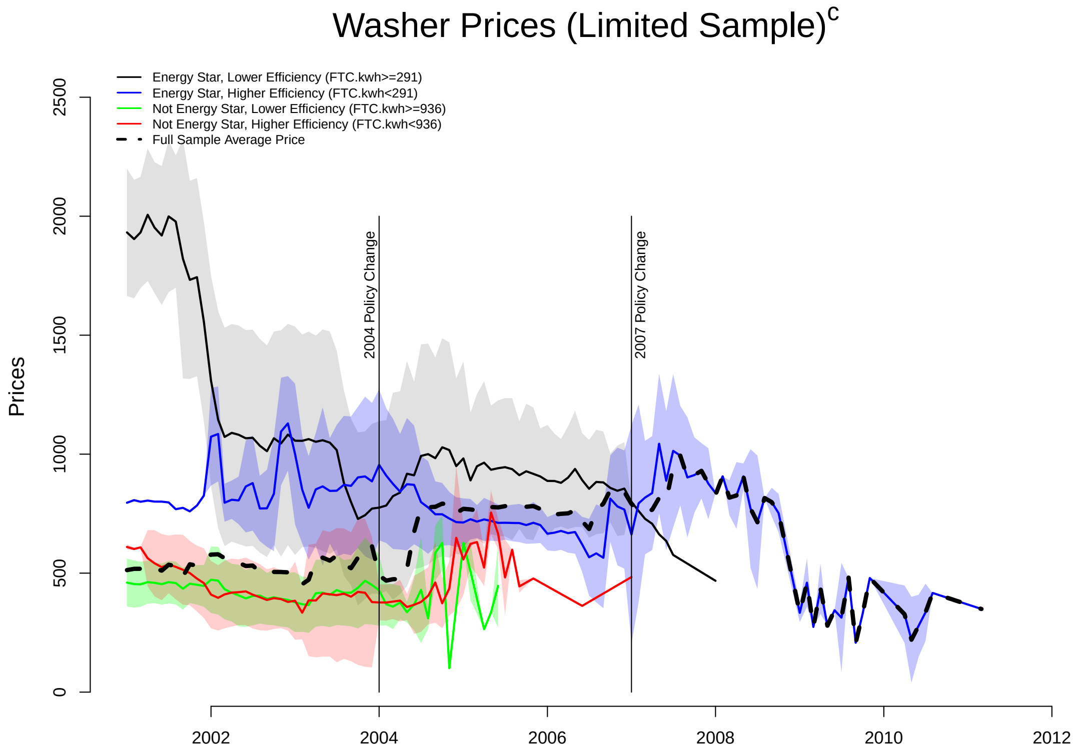

Figure 3.10: Average Washer Prices — Limited Sample (models existed before and after 2004)

[c]Limited sample is a subset of the full sample that includes all washer models observed in the data at least once before and after the 2004 policy change data. High and low efficiency values are selected according to the mean value of the FTC energy rating taken across model numbers in this particular (limited) sample during the period from 2001 to 2006. Note that the data include a varying number of retailers over time and the sample of retailers may not be representative of the population. The lines for each category of washer indicate the value for the mean washer sold in that category, and the band around the line marks +/- 1 standard deviation. The thick dashed line is the average price across all models, with models weighted by number of units sold.

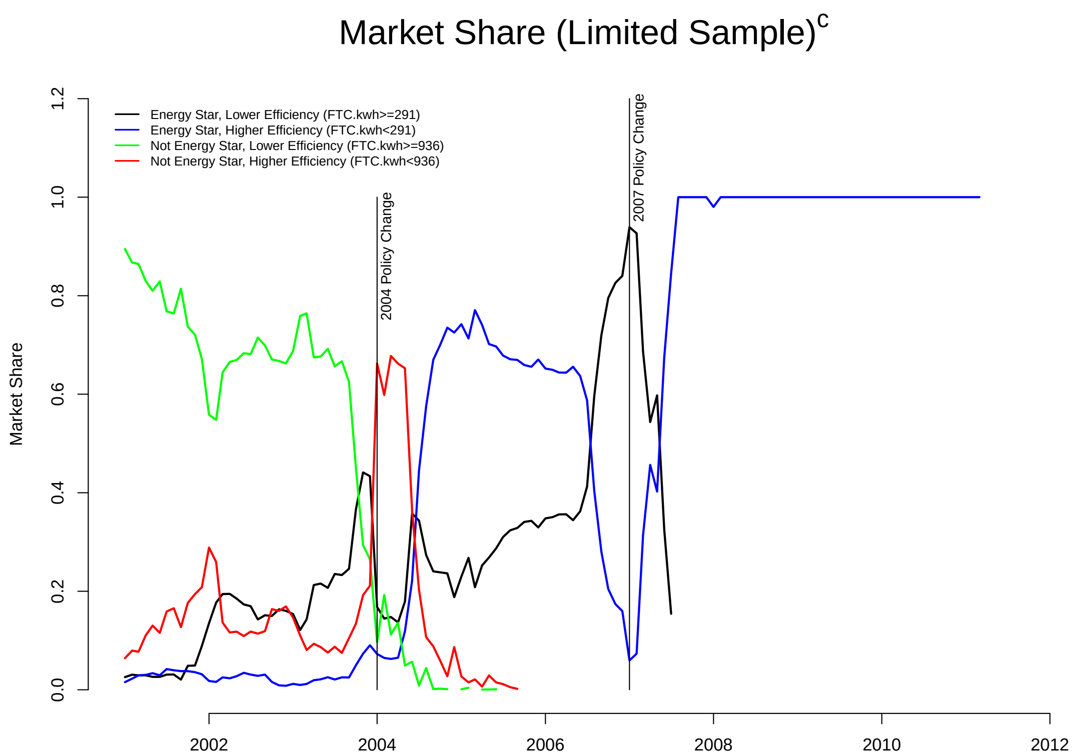

Figure 3.11: Market Share — Limited Sample (models existed before and after 2004)

[c]Limited sample is a subset of the full sample that includes all washer models observed in the data at least once before and after the 2004 policy change data. High and low efficiency values are selected according to the median value of the FTC energy rating taken across model numbers in this particular (limited) sample during the period from 2001 to 2006. Note that the data include a varying number of retailers over time and the sample of retailers may not be representative of the population.

Table 3.6: Summary of Regressions (2004 Policy Change)

	Price Effect of Policy			Quantity Effect of Policy		
	β	γ	Adj. R^2	β	γ	Adj. R^2
	Estimate/(t-statistic)					
Basic — Full Sample	-99.62	0.03	0.97	929.29	-1.54	0.63
	(-10.02)	(1.93)		(6.80)	(-7.60)	
Robust — Full Sample	14.95	-0.11	0.97	611.06	-1.32	0.63
	(1.26)	(-6.31)		(3.53)	(-5.08)	
Basic — Limited Sample	-106.51	0.04	0.96	585.47	-1.35	0.47
	(-6.59)	(2.15)		(3.12)	(-5.73)	
Robust — Limited Sample	-6.10	-0.10	0.96	100.82	-0.48	0.50
	(-0.29)	(-3.52)		(0.39)	(-1.40)	

Notes: All regression models include fixed effects for each model of washing machine. *Basic* models include a natural spline of time (on a monthly scale) with four knots. *Robust* models include a separate natural spline of time for each of the four classes of washer as classified in the data section and plotted in the above figures. The β coefficient is for the dummy variable indicating the post-2004 time period. The γ coefficient is the effect on the interaction of the year-2004 dummy variable and the unit's FTC energy-efficiency rating, measured in kWh of energy use per year.

Table 3.7: Estimates of Minimum Change in Consumer Surplus (2004 Policy Change)

Baseline	Full Sample		Limited Sample	
Price and Quantity	Basic Model	Robust Model	Basic Model	Robust Model
	Dollars Per Month/ (Percent of Total Sales)			
$\bar{p}$ & $\bar{q}$ equal average	7,037,929	3,357,163	820,007	976,705
of 2003 & 2004	(7.47)	(3.56)	(7.59)	(9.04)
p_0 & q_0 equal average	-1,024,409	-3,926,984	2,542,662	3,405,117
of 2001-2003	(-0.47)	(-1.82)	(12.32)	(16.50)
p_1 & q_1 equal average	16,628,850	5,946,159	1,422,502	1,223,729
of 2004-2006	(8.32)	(2.97)	(31.06)	(26.72)

Notes: The table reports lower bound changes in consumer surplus as described in the text, measured in dollars per month. The values differ according the sample of data used (limited or full), the regression model from which changes in quantities and prices were derived (basic or robust), and the baseline prices and quantities used. The numbers in parenthesis give the minimum change in surplus as a percent of average sales, where the denominator is $\frac{1}{2}(p_0q_0 + p_1q_1)$ and p_0, q_0, p_1 and q_1 are as defined in the text.

REFERENCES

Allcott, H. 2011. "Consumers' Perceptions and Misperceptions of Energy Costs." *The American Economic Review*, 101(3): 98–104.

Anderson, Soren T, Ryan Kellogg, and James M Sallee. 2011. "What Do Consumers Believe About Future Gasoline Prices?" National Bureau of Economic Research.

Barsky, Robert, and Lutz Kilian. 2004. "Oil and the Macroeconomy Since the 1970s." National Bureau of Economic Research Working Paper 10855.

Black, Fisher, M C Jensen, and Myron Scholes. 1972. "The Capital Asset Pricing Model: Some Empirical Tests." *Studies in the theory of capital markets*, 81(3): 79121.

Blume, Marshall E., and Irwin Friend. 1973. "A New Look at the Capital Asset Pricing Model." *The Journal of Finance*, 28(1): pp. 19–33.

Bodie, Zvi, and Victor I. Rosansky. 1980. "Risk and Return in Commodity Futures." *Financial Analysts Journal*, 36(3): pp. 27–31+33–39.

Breeden, Douglas T. 1979. "An intertemporal asset pricing model with stochastic consumption and investment opportunities." *Journal of financial Economics*, 7(3): 265–296.

Chen, Nai-Fu. 1991. "Financial Investment Opportunities and the Macroeconomy." *The Journal of Finance*, 46(2): 529–554.

Chen, Nai-Fu, Richard Roll, and Stephen A. Ross. 1986. "Economic Forces and the Stock Market." *The Journal of Business*, 59(3): pp. 383–403.

Cochrane, John H. 2001. *Asset pricing.* {Princeton University Press}.

Cook, J. R., and L. A. Stefanski. 1994. "Simulation-Extrapolation Estimation in Parametric Measurement Error Models." *Journal of the American Statistical Association*, 89(428): 1314–1328.

Cox, John C, Jonathan E Ingersoll Jr, and Stephen A Ross. 1985. "An intertemporal general equilibrium model of asset prices." *Econometrica: Journal of the Econometric Society*, 363–384.

Darby, Michael R. 1982. "The Price of Oil and World Inflation and Recession." *The American Economic Review*, 72(4): pp. 738–751.

Dasgupta, Partha, and Eric Maskin. 2005. "Uncertainty and Hyperbolic Discounting." *The American Economic Review*, 95(4): pp. 1290–1299.

Dusak, Katherine. 1973. "Futures Trading and Investor Returns: An Investigation of Commodity Market Risk Premiums." *Journal of Political Economy*, 81(6): pp. 1387–1406.

Fama, Eugene F., and James D. MacBeth. 1973. "Risk, Return, and Equilibrium: Empirical Tests." *Journal of Political Economy*, 81(3): pp. 607–636.

Fama, Eugene F., and Kenneth R. French. 1992. "The Cross-Section of Expected Stock Returns." *The Journal of Finance*, 47(2): pp. 427–465.

Fama, Eugene F., and Kenneth R. French. 1996. "The CAPM is Wanted, Dead or Alive." *The Journal of Finance*, 51(5): pp. 1947–1958.

Hamilton, James D. 1983. "Oil and the Macroeconomy since World War II." *Journal of Political Economy*, 91(2): pp. 228–248.

Hamilton, James D. 1996. "This is what happened to the oil price-macroeconomy relationship." *Journal of Monetary Economics*, 38(2): 215 – 220.

Hamilton, James D. 2008. "Understanding Crude Oil Prices." National Bureau of Economic Research Working Paper 14492.

Hamilton, James D. 2009. "Causes and Consequences of the Oil Shock of 2007-08*."

Hamilton, J.D., and J.C. Wu. 2011. "Risk premia in crude oil futures prices." Working paper. University of California at San Diego.

Hansen, Lars Peter, and Scott F Richard. 1987. "The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models." *Econometrica: Journal of the Econometric Society*, 587–613.

Harvey, Andrew, Esther Ruiz, and Enrique Sentana. 1992. "Unobserved component time series models with Arch disturbances." *Journal of Econometrics*, 52(1?): 129 – 157.

Hotelling, Harold. 1931. "The Economics of Exhaustible Resources." *Journal of Political Economy*, 39(2): pp. 137–175.

Hughes, Jonathan E., Christopher R. Knittel, and Daniel Sperling. 2008. "Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand." *The Energy Journal*, 29(1): 113–134.

Jagannathan, Ravi, and Zhenyu Wang. 1996. "The Conditional CAPM and the Cross-Section of Expected Returns." *The Journal of Finance*, 51(1): pp. 3–53.

Jensen, Michael C. 1968. "The Performance of Mutual Funds in the Period 1945-1964." *The Journal of Finance*, 23(2): pp. 389–416.

Keim, Donald B., and Robert F. Stambaugh. 1986. "Predicting returns in the stock and bond markets." *Journal of Financial Economics*, 17(2): 357 – 390.

Kilian, Lutz. 2008. "The Economic Effects of Energy Price Shocks." *Journal of Economic Literature*, 46(4): 871–909.

Kilian, Lutz. 2009. "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market." *American Economic Review*, 99(3): 1053–69.

Kim, Chang-Jin, and Charles R. Nelson. 1999. *State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications.* Vol. 1 of *MIT Press Books*, The MIT Press.

Leland, H.E. 1979. “Quacks, lemons, and licensing: A theory of minimum quality standards.” *The Journal of Political Economy*, 1328–1346.

Lettau, Martin, and Sydney Ludvigson. 2001. “Resurrecting the (C)CAPM: A Cross-Sectional Test When Risk Premia Are Time-Varying.” *Journal of Political Economy*, 109(6): pp. 1238–1287.

Lewellen, Jonathan, and Stefan Nagel. 2006. “The conditional CAPM does not explain asset-pricing anomalies.” *Journal of Financial Economics*, 82(2): 289 – 314.

Merton, Robert C. 1973. “An intertemporal capital asset pricing model.” *Econometrica: Journal of the Econometric Society*, 867–887.

Newell, Richard G., and William A. Pizer. 2003. “Discounting the distant future: how much do uncertain rates increase valuations?” *Journal of Environmental Economics and Management*, 46(1): 52 – 71.

Nordhaus, William D. 2007. “A Review of the ”Stern Review on the Economics of Climate Change”.” *Journal of Economic Literature*, 45(3): pp. 686–702.

Ronnen, U. 1991. “Minimum quality standards, fixed costs, and competition.” *The RAND Journal of economics*, 490–504.

Shiller, Robert J. 2005. *Irrational exuberance.* Princeton University Press.

Stern, Nicholas. 2006. *Stern review on the economics of climate change.* Stern Review on the Economics of Climate Change, UK Treasury.

Tol, Richard S.J. 2006. “The Stern Review of the economics of climate change: a comment.” *Energy & Environment*, 17(6): 977–981.

Weitzman, Martin L. 1998. “Why the Far-Distant Future Should Be Discounted at Its Lowest Possible Rate.” *Journal of Environmental Economics and Management*, 36(3): 201 – 208.

Weitzman, Martin L. 2001. “Gamma Discounting.” *The American Economic Review*, 91(1): pp. 260–271.

Weitzman, Martin L. 2007. “A Review of ”The Stern Review on the Economics of Climate Change”.” *Journal of Economic Literature*, 45(3): pp. 703–724.