Three Essays on the Objective Function in Economics

by

Daniel A. Goldstein

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Public Policy and Economics)
in The University of Michigan
2007

Doctoral Committee:

Professor Kathryn Mary Dominguez, Co-Chair
Professor Miles S. Kimball, Co-Chair
Professor Robert B. Barsky
Professor Yu Xie
Assistant Professor Christopher L. House

© Daniel A. Goldstein

All rights reserved
2007

To my parents,
for their countless gifts of opportunity, encouragement, and love.

Acknowledgements

This work would not have been possible without the support of many people.

I particularly want to thank Miles Kimball and Yu Xie, who were instrumental in helping me through a period when I was really stuck. Without their support I do not know when or if I would have completed this work.

I am also grateful to the other members of my committee, Kathryn Dominguez, Robert Barsky, and Chris House, for their extremely valuable comments on earlier drafts of these (and other) essays. Each individual's contributions have really helped to shape these papers.

Countless thanks go to Olivier Coibion, my co-author and driving force for Chapter 2. I am so glad you thought of me when this project was in its infancy.

Other thanks go to Emily Greenman, Yuriy Gorodnichenko, Peter Morrow, King Banian, Serena Ng, Mary Corcoran, my Friday morning third year paper group, and attendees of various seminars and workshops at the University of Michigan. Each one provided invaluable insights and support during this process. I apologize for the inevitable omissions from this list.

Lastly, I could not have even begun my work here at Michigan but for the financial support of the Gerald R. Ford School of Public Policy and the Horace H. Rackham School of Graduate Studies.

Table of Contents

List of Figures

List of Tables

List of Appendices

Introduction

One of the pillars of economic thought is the concept of constrained optimization. The economist says: if I know what it is that you want (your objective), and I know what choices are available to you (your constraints), then I can predict how you will behave. How? Simple. Determine which of those possible choices gives you the “most” of what you want, and that is what you will choose. After all, why would an individual intentionally choose any other outcome?

I personally find this logic appealing. Perhaps that is why I have chosen to study economics. Constrained optimization – in which individuals (or organizations) choose from available options so as to maximize some objective – is an extremely tidy way of viewing the world, and potentially a powerful one. Difficult problems of choice ultimately boil down to the solution of a math problem – albeit often a difficult problem. However, the method is fragile in (at least) two fundamental ways.

The first is the assumption that agents do in fact make choices so as to maximize their objective. This is what I will call the assumption of rationality. Unfortunately, while this usage of the term rational is a common one, there is another, related usage which critics of this approach tend to call upon. Rationality in those contexts is used to describe *how* agents make choices: it implies that individuals are little more than cold-blooded computers that go through a comprehensive deliberative process before making any decisions. This caricature bears so little resemblance to our common experiences that it may lead some to quickly discard the optimization framework altogether. Surely emotional factors play a role in our everyday lives, and do we really carefully process all available information before making even the smallest decisions?

Now, when *I* speak of rationality, I do not make any claims as to *how* individuals make decisions. Rather, I am merely referring to the principle that individuals do not in general intentionally choose second-best (or worse) alternatives. To do so is akin to making the statement: "I know I prefer to do X, but I choose Y." I am reluctant to accept that this is the standard model of how humans behave. If it is, then the vast majority of economic theory is an irrelevant intellectual exercise, even as an approximation. As for the question of the procedures that go on in an individual's brain when determining a course of action, I do not believe an economist's skills are well suited to such a question. Neuroscience and psychology would seem more appropriate, and in fact there has been progress recently with respect to how the brain actually processes information and makes decisions (see Damasio 1994 and 2003 for a non-technical overview). At any rate, to avoid the possibility of confusion, from this point forward I will use the terms such as "optimal" or "neoclassical" to refer to my narrower definition of rationality.

Actually, an agent would probably never make the above statement "I know I prefer X, but choose Y." Instead, we, as outsiders, might see an agent choose Y; while we believe X is actually optimal. This leads to the second potentially fragile assumption underlying the method of constrained optimization. What is the individual's objective? In practice, economists will typically assume that an agent's objective function takes on some convenient mathematical form, spin it through the standard toolbox, and study the implications. This makes the problems we encounter tractable, and tractability is of no small importance.

However, tractability does not necessarily make the procedure accurate. Take, for instance, the caricature of the homo-economicus. This type of agent is only concerned with his own material well-being. That is, in his utility function you'd only find items like consumption, or money. If we observe this homo-economicus giving money to charity, then he is behaving in a non-optimal manner, since we have pinned down his objective function as not having any preference for such a non-material object. Fine. However, if we observe another agent – for whom we have not pinned down such a clean set of preferences – give money to charity, we should not apply the homo-economicus's preferences to her and conclude she too is being non-optimal. Her actual preferences could require charitable giving as a means of maximizing her overall well-being.

While the three essays which follow are strictly independent of each other, the common thread uniting them all is that each essay suggests how modifying an agent's objective functions away from a more conventional form can bring theoretically "optimal" behaviors more in line with potentially puzzling observations. First, Chapter 1 suggests that terrorists – even suicide terrorists – cannot be necessarily dismissed as completely irrational. It appears that economic conditions – as represented by the price of crude oil – seem to be related to the frequency of attacks. I argue that deteriorating economic health in the region may tip some individuals who are receptive to the notion of being a terrorist into actually being a terrorist. This reminds us that policymakers might well benefit from economic-style models that explicitly note that some individuals – that is, prospective terrorists – have a preference for making destructive, symbolic attacks against their enemies. We should not be quick to dismiss them as being simply "crazy."

Chapter 2, on monetary policy, follows a (vaguely) similar line of reasoning. Here, we present a surprising finding that interest rates chosen by the central banks of the US and the EU appear to respond to the variance of unemployment across regions. While we propose several theoretical explanations for this phenomena, the one which most accurately describes the Federal Reserve's policy (but not the ECB's) is the one in which we modify the Federal Reserve's objective function to one that includes a preference for less variance in overall welfare across states.

It is in chapter 3 where the above discussion on the importance of correctly describing preferences is most relevant. Here, I describe how it is technically straightforward to include emotions explicitly in an agent's utility function. I go on to explore some of the implications of augmenting the objective functions of individuals in this way, and find that this technique seems to be an appropriate in describing many emotionally-linked behaviors. That the framework can be useful in a variety of contexts suggests that it might be an appropriate baseline model for describing behavior. While the technique may not be universally applicable, it seems reasonable to try and expand the standard framework of constrained optimization to describe more than just the behavior of cold-blooded automatons. Indeed, acknowledging that humans might have preferences over hot-blooded passions has the potential to open up new worlds for economic-style analyses of behavior.

Chapter 1

Can Oil Prices Help Predict Terrorism?

Abstract

While acts of terrorism are usually attributed to motivations of political, religious, or other ideological nature, this paper provides evidence suggesting that certain types of terrorist attacks are correlated with changes in a readily available economic variable. Particularly, over the last thirty years significant terrorist attacks directed against the United States or Western European targets by Middle Eastern organizations occur more often after the price of crude oil falls than after it rises. A series of logistic regressions estimate how the probability of attack varies based purely on changes in oil prices over various time horizons. While the probabilities of attack implied by these regressions are not extraordinarily large even after significant changes in oil prices, the relative change in probabilities is considerable. The paper further informally discusses two hypotheses for why there might be a causal relationship between oil prices and terrorist events of the sort suggested by the empirical results. The most plausible of these is an expected utility model, in which those that are predisposed to committing terrorist acts – but have yet to do so – may be triggered to act when local economic conditions deteriorate.

1 Introduction

The motivations for terrorist attacks are usually attributed to political, religious, and/or ideological conflicts. There can also be little doubt that clashes in these dimensions are fundamental in spurring terrorist activity. After all, terrorism is but one form of rebellion – albeit a horrific one – which requires at least a minimal level of discontent with the status quo. Unfortunately, conflicts along these dimensions are often unchanging for long periods of time, making it difficult to monitor their changes so as to try and gauge when the likelihood of a terrorist strike is high. This is regrettable, as an improved ability to predict terrorist events based on changing conditions can only benefit efforts to detect and prevent such attacks.

Moreover, while the literature is fairly thin, economic explanations of terrorist activity have not generally been successful. In their analysis of cross-sectional data from Lebanon and the occupied territories in Israel, Krueger and Maleckova (2002) argue that the participation in and support for terrorism is not correlated with poverty. Indeed, they find that Palestinian suicide bombers tend to be drawn disproportionately from households not in poverty. This finding is echoed by Atran (2003). While Blomberg, Hess and Weerapana's (2002) results suggest that "terrorism from within" is correlated with the business cycle, their finding is limited to high income nations, and lacks the international dimension that dominates the current policy debate.

This paper diverges from these literatures. The evidence presented here suggests that in one instance readily available economic data can aid in predicting the timing of a certain type of terrorist event. Specifically, the key empirical finding below is that, over the last thirty years significant terrorist attacks aimed against the United States or Western European nations[1] by organizations based in the Middle East are considerably more frequent after a decline in crude oil prices than after an increase. A series of logistic regressions verifies this observation, and ultimately allows one to assign a probability as to whether or not a terrorist attack will occur next month based solely on recent movements in oil prices. While this paper does not take the position that changes in oil prices are a primary cause of anti-Western terrorism, the evidence and discussion presented here suggests why it is sensible for oil price movements of this nature to in fact be a contributing factor to the periodic surges in Middle-Eastern terrorist activity against the West.

The bulk of what follows is devoted to describing the very simple approach used to uncover this somewhat surprising finding, along with the results themselves and a brief demonstration of the straightforward conversion of these results to probabilities of an attack. Section 2 covers some of the nuts and bolts of the data and its construction. This is followed in Section 3 by a description of the general flavor of data, with Section 4 presenting the key quantitative results. While the paper's main contribution is as a piece of descriptive statistics, the discussion in Section 5 considers some of the implications of these findings and informally suggests two

[1] Which will be referred to jointly as "the West" or "Western Nations." I include here attacks against Western Europe, the United States, Canada, and Australia.

causal hypotheses for why this apparent inverse relationship between oil prices and terrorism might exist.

2 Event Selection and Timing

As we are interested in how oil prices move before (and, to a lesser extent, after) terrorist events, it is obviously important how we choose those events. Terrorism is defined by the United States government as premeditated, politically motivated violence against civilians or non-mobilized military personnel by either sub-national groups or clandestine agents, usually intended to influence an audience. In the analysis that follows, the following additional criteria is imposed when selecting events: A) the principle target of the attack was civilians or military personnel of Western Nations, B) the attack was instigated by groups with Middle Eastern origins, and C) the attack was of significant size.[2] A dummy variable (E) was coded 1 for each month where a qualifying event occurred, and zero otherwise. Table 1.1 lists 43 terrorist events selected since 1974 that meet these criteria, based on several different data sources. There are 353 possible months between January 1974 and May 2003, so an event occurs in about 12.2% of the months in this time period.

The measure of oil prices used in this paper is the monthly average of Persian Gulf crude oil prices paid by American refineries. This data is provided by the U.S.

[2] By significant size, I impose the requirement that at least 5 individuals are held hostage, kidnapped, injured or killed in a given event.

Energy Information Administration.[3] In an attempt to keep these prices in constant dollars, they are deflated by the United States consumer price index (CPI) excluding energy.[4]

While the level of oil prices is important in this study, the key analyses that follow focus on how prices were changing around terrorist events. One potentially important consideration will be how we consider the timing of oil price changes around the terrorist events. *A priori*, we do not want to exclude the possibility that oil prices adjust immediately in response to a terrorist event. Given that the terrorist groups in this analysis are of Middle Eastern origin, and the Middle East provides a large fraction of the world's supply of oil, oil traders may view terrorist activities – or lack thereof – as indicative of changing stability in oil supplies. Thus, a terrorist event may impact petroleum prices immediately. If prices are indeed immediately influenced by terrorist events, then the price of oil during the month an event takes place would not give us a good measure of what prices would have been in the absence of the event. This would be particularly true if the event occurred early in a month.

---

[3] While Persian Gulf Crude prices are used predominantly throughout this paper, the same analysis was performed with a composite OPEC price, a U.S. domestic crude oil price, and a composite U.S. refiner acquisition cost. The results are similar for all price series, which is unsurprising given the high correlation of prices across types. However, there do appear to be small timing differences across each series, which is likely a result of the data collection process.
[4] As Middle East Oil transactions are typically denominated in American dollars, and the currencies of Middle Eastern nations are pegged to the dollar, deflating by a U.S. price index is appropriate. Using the version of the CPI that excludes energy prices is also reasonable, as oil prices might otherwise independently contribute to movements of the CPI, which we want to avoid.

In an attempt to allay this concern, when measuring how prices change around a given month I use the price of oil in the preceding month as a reference point. Thus, for clarity two different terms are used throughout this paper. Each observation in the data represents one month, which I'll call a "month of interest." In the end, it is these "months of interest" that we care about, and there may, or may not be, a qualifying terrorist event within that month. The month preceding each "month of interest" I call a "reference month." All changes in oil prices around a month of interest (time t) are calculated relative to the price in the preceding (or "reference") month (time t-1). I thus define the variable $P_t^k = \ln(price_{t-k}) - \ln(price_{t-1})$ as the change in the price of oil between time t and horizon k. If prices rise over this time period, the value of P_t^k will be negative. In short, P_t^k is a measure of how much higher, in percentage terms, prices were k months ago.

There is one caveat, however, for this "month of interest" and "reference month" relationship. Note that a number of the events took place in the last week of a given month (see Table 1.1). When an event occurs at the end of a month, even if oil prices do indeed respond immediately to this event, the *average* prices for the month in which the event occurred would not be expected to change significantly. For those cases where events occur at the end of the month, the price of oil during the actual event month would seem to be a better measure of the prices that precede an event than the prices during the previous month. Mechanically, I attribute the "month of interest" for these end-of-month events to the *following* month. These attributions are included in Table 1.1.

3 General Trends

Having selected a set of events that meet the above criteria, next on the agenda is to examine how oil prices generally tend to move around these terrorist events. As an overview of the data, Figure 1.1 plots the real Persian Gulf Crude Oil price level over the period January 1974 – May 2003, with vertical bars representing the dates of the terrorist attacks in Table 1.1. One impression taken from this chart is that terrorist attacks were particularly frequent during two time spans – January 1983 through January 1987, and January 1997 through January 1999 – during which crude prices were falling significantly. Another overt feature is that relatively few events occurred during the price run-up of the second oil crisis.

This impression is further strengthened by scrutinizing how prices were changing on average around the dates of the terrorist attacks. (See Figure 1.2.) The variables graphed here are the average and median value of the price change variable P_t^k, which was defined above, at various horizons of k months. The average and median values are calculated both for months where a terrorist event occurs and months where one does not. The downward trend in oil prices prior to months when terrorist events occur differs strikingly from the price trend before months where events do not occur.

It turns out that the differences – both for the averages and the medians – are statistically significant. Table 1.2 provides an event-by-event breakdown of oil price changes in the months preceding the events. (This is the data underlying Figure 1.2

for the 12 months preceding the events.) T-tests of the means and Pearson Chi-Squared tests of the medians reveal the statistical significance of the differences between event/non-event months. Furthermore, it is apparent from the last two rows of this table that oil prices fall prior to an event considerably more often than they rise. For example, six months before an event oil prices were more than 10% higher than in the reference month in 15 out of 43 cases and were 10% or lower for only 4. By the same criteria, at nine and twelve month horizons the ratio is 21 against 6 – a ratio greater than three to one.[5]

4 Main Results

4.1 Core Regressions

From this preliminary "eyeballing" of the data, there is a strong suggestion that oil price movements tend to be systematically different prior to months where a qualifying terrorist event occurs than prior to months where one does not. The next objective is to gauge whether these changes in oil prices alone can in fact help improve a prediction as to whether a terrorist event is going to occur or not.

---

[5] Three additional Middle East security events, which are similar in many ways to the terrorist events included in this analysis but do not qualify as terrorist events by the criteria employed, were also preceded by sharp declines in crude oil prices. These are: Iraq's invasion of Kuwait in August 1990, the assassination attempt on George HW Bush in April 1993, and the ambush of American troops in Mogadishu Somalia in October 1993 (this is the "Black Hawk Down" incident, and was, some believe, supported by Al Qaida operatives). Furthermore, not included are numerous kidnappings of individual Americans that occurred in Lebanon during the period between 1982 and 1987, which coincided with the long oil price recession of the mid-1980s.

To this end, we want to estimate the following equation that specifies the potential relationship between a change in the oil prices and the occurrence of a terrorist event: $E_t = \beta_0 + \beta_1 P_t^k + \varepsilon_t$.

As described above, E_t is a dummy variable coded 1 for a month of interest t in which a qualifying terrorist event occurs, and zero otherwise. P_t^k is, as defined above, the change in the real oil price of Persian Gulf crude between an observation at time t and horizon k. With a dichotomous variable as the dependent variable, it is appropriate to use logistic regression to estimate this relationship. Here, the logit is solved via the method of maximum likelihood, Huber-White robust standard errors are calculated, and a Z-statistic is estimated based on these parameters. This Z-statistic can be used to test the probability that the null hypothesis – the β_1 coefficient is equal to zero – is true given the data.[6]

Table 1.3a provides a summary of the results for various versions of this logit. The logit coefficients listed in this table correspond to the β_1 values from the regressions over various price change horizons (k). Here, the β_1 values have been exponentiated to provide the odds ratio interpretation of the logit coefficient. If this coefficient has a value greater than one, it implies that the higher prices are at time k relative to the reference month – or equivalently, the more prices fell over the k months preceding a month of interest – the greater the odds of an event occurring.

---

[6] As one would expect, employing a probit specification yields virtually identical results, but the logit permits the convenient odds ratio interpretation.

More precisely, the values of these coefficients represent the change in the odds of attack from a 1% increase in the price difference over a given horizon.

Table 1.3b manipulates these results further, estimating the marginal effect of changes in prices on the probability of an attack. Each coefficient here represents the slope of the probability response curve implied by the logit, evaluated at the mean value of the price difference.[7] For example, consider the results for Persian Gulf Crude prices at a nine month horizon. The value of the marginal effect is 0.196. This means that, compared to the case where the price had actually changed by its mean amount (which is a little greater than zero), if prices were actually 1% higher eight months ago this would increase the probability of an attack by 0.19 percentage points. By extension, a 5% difference increases the probability of attack by roughly one percentage point.

There are two key results evident in Tables 1.3a and 1.3b. First, the odds ratios are greater than one in every specification – and by extension, the marginal effects are greater than zero in every specification – and are in general statistically significant at a level of at least 1% for time horizons between five and twelve months. As discussed above, values for the odds ratio greater than one indicate that the likelihood of attack rises if prices have fallen.

We can see this result visually by using the logit models to calculate the probability of attack over a range of hypothetical price changes, and graphing the results.[8] Figures 3 and 4 graph some examples of the probability response curves implied by their respective logit results. Here, the y-axis measures the probability of an event next month, given the change in log prices on the x-axis. The upward sloping curves demonstrate the implication that the more prices fall over a given time period, the higher the probability of an attack next month. The slopes of these curves also offer a visual representation of the marginal effect of changes in oil prices on the event probabilities. Due to the non-linearity inherent in a logistic regression, the slope of the curve becomes steeper as the difference in prices increases.

The second observation from these tables is that in general over shorter time horizons the logit coefficients tend to be larger. This suggests that given an equal sized percentage drop in prices over two different time horizons, the probability of an attack increases more for the shorter time horizon. This is evident in Figure 1.3, in which the slope of the implied probability response curve is steeper for the shorter horizon measure.

Of course, the variance in oil price changes is smaller over shorter horizons. A 15% change in prices over one month is more than two standard deviations of this series, whereas that same change over one year is roughly one-half standard deviation. It is thus of interest to compare how sensitive the probability of attack is to

---

[7] This is roughly equal to zero for short horizons, and slightly positive for longer horizons (reflecting the fact that on average oil prices decreased slightly over the time period.)

[8] The implied probability q of the event given oil price change P^{ik} is the solution to: $\ln(q/(1-q)) = \beta_0 + \beta_1 P^k$.

changes in prices of comparable relative magnitude over different time horizons To this end, a second set of regressions was run after standardizing each price series. These results are summarized in Tables 4a and 4b. Note the Z-values are identical to those in Tables 1.3a and 1.3b, as one would expect after dividing each series by a constant term. However, the interpretation of these coefficients is somewhat different, in that each coefficient now represents the change in the odds ratio from a one standard deviation change in price, rather than a one percent change in price. In the case of Persian Gulf crude prices, the probability of attack next month is most sensitive to a one standard deviation change in prices over an eight month interval (that is, when k=9). This standard deviation is a change in prices of approximately 25% over this span.

Visually, it is clear from Figure 1.4 that the probability of terrorism is more sensitive to the standardized change over an eight month interval than the two month interval, which is exactly the opposite conclusion that one would make from inspection of Figure 1.3. Still, the differences in the slopes of the curves are not extraordinarily large, and are certainly not different from each other in a statistically significant sense. This suggests that the rate of change in prices matters more than the magnitude of that change.

With most time-series analyses, the fundamental concern is that there may be serial correlation. In the context of this paper, we would have serial correlation if the recent history of terrorist activity has predictive power. It appears, for instance, that events tend to occur in clusters. (That these clusters occur during oil price recessions is a big part of what is driving the results.) These clusters may lead the errors from the regression to be correlated with one another. Not accounting for the history of events to eliminate the correlation would lead to bias in the coefficient on oil price changes.

The traditional way to eliminate concerns of serial correlation is to include lags of the dependent variable on the right hand side of the regression equation. To this end, Tables 1.3a-1.4b also offer the results from a dynamic logit version, which includes 12 lags of the event dummy variable on the right-hand-side, in addition to our price change variable P_t^k. In none of these formulations to the 12 lags turn out to be even jointly significant – that is, the recent history of terrorist events has no predictive power. While inclusion of the lagged event dummies absorbs some of the sensitivity of terrorist events to oil prices, it only does this to a small degree. The baseline results remain robust to this modification, suggesting that serial correlation – or, the clumping of events – is not a source of bias. (See De Jong (2004) for discussion of the properties of these types of dynamic logits.)

An alternative means of countering concerns about clustering is to remove from the data the observations from a suspicious cluster and see if the regression results are unchanged. As mentioned, the surge in terrorist activity in the early/mid-eighties – which coincides with deterioration in OPEC's ability to maintain high prices – is a particularly striking feature of the data. To determine if this cluster of events is indeed what is driving the overall pattern, I excise observations between the peak of oil prices in February 1981 and the trough in June 1986 from my sample.

Doing this cuts out 14 events, or about 1/3 of the total sample of terrorist attacks. Despite this, neither the direction nor the statistical significance of the results are substantially changed when rerunning the regressions over this time period (see Tables 1.3a and 1.3b). This is further evidence that the apparent relationship between oil prices and attacks is independent of the recent history of attacks. As removing data from one's sample is rarely a desirable step, I will be using the full sample for the remainder of this piece.

Another possible specification of interest is that there may be a nonlinear relationship between the change in oil prices and terrorism. Do large changes in prices have an increasingly large impact on the likelihood of an attack? Table 1.6 offers an augmented version of the baseline univariate model that includes a term for the k-month percentage change squared. The value for the squared coefficients is less than one at almost every horizon, suggesting that as the change in price gets larger and larger, the increase in likelihood gets smaller. However, the squared change in price is only statistically significant for one horizon, suggesting that the linear specification is capturing most of the relationship.

To get a sense for how the baseline empirical model functions as a predictive device, Figure 1.5 plots the probability of attack implied by one of these regressions (Persian Gulf Crude change over a nine month horizon) with vertical bars denoting actual event dates. The model seems to do a reasonably good job. Most of the periods where the predicted probability of attack is high are marked by one or more attacks, whereas relatively few of the low probability periods are marked by attacks –

and when attacks do occur during low-probability periods they are generally isolated. The only noteworthy "false positive" would be the period in mid-1991 when oil prices fell precipitously in the wake of the first Iraq war, yet there were no terrorist attacks. Significant "false negatives" would have been two events in late 1979 – including the Iran Hostage crisis – the UTA airliner bombing in mid-1989, and the Bali nightclub bombing in late 2002, each of which occurred after periods when oil prices had risen significantly.

4.2 Implications: Example Calculation of the Probability of Attack

The relative changes in probability implied by these results are striking. This is most readily apparent from the fact that the probability response curves implied by these regressions (Figures 1.3 and 1.4) are fairly steep. A hypothetical numerical example can help emphasize this. Suppose we decide to look back in time at Persian Gulf oil prices eight months ago in order to try and gauge the probability of a terrorist event occurring next month (i.e. looking back nine months before month of interest). This particular specification is chosen because it appears (via the results in Table 1.4b) that the likelihood of an event occurring is most sensitive to changes over the eight month horizon.

As a baseline, suppose that prices are unchanged between today and eight months ago ($P_t^9=0$). The logit results indicate that if prices are unchanged then the probability of an event next month is about 11.2%. (The 95% confidence interval for this result is between 8.1% and 15.1 %.) Now suppose, however, that eight months ago prices were 20% higher than they are today. 20% is a slightly less than one standard deviation of the actual series of price changes over a nine month period, so it represents a sizeable, but not uncommon change in prices. Based on the results of this logit, the implied probability of a terrorist event next month is 15.6% (95% C.I.: 11.4% - 21.0%). Similarly, if prices were actually 20% lower nine months ago than they are today, this implies a probability of attack of 7.9% (95% C.I.: 4.9%-12.4%). By extension, we can conclude that the probability of an attack is almost twice as great if prices were 20% higher nine months ago than today (15.6%) compared to if prices were 20% lower nine months ago (7.9%).

Granted, even after this relatively large hypothetical drop in prices, a 15.6% probability of attack does not resolve terribly much of the uncertainty as to whether or not a terrorist event of this type is going to occur. Still, given that terrorism is largely considered a political phenomenon, it is noteworthy that a basic economic variable such as oil prices can apparently significantly aid prediction at all.

5 Discussion

5.1 The Question of Causality

If this statistically significant correlation between oil prices and terrorism is not merely a result of chance, then there must be some type of causal mechanism at work. It should be stressed, however, that despite the statistical significance of oil price changes in the above regressions, this does not by any means prove that there is a direct causal linkage between the two. The above analysis is univariate, which implicitly assumes that oil price changes are independent of any other "terrorism-causing" variables. However, the changes in oil prices observed may not be random in this respect, and thus the possibility remains that there is some other variable that is both correlated with oil prices and also prompts higher rates of terrorism.[9] That is, we may have omitted variable bias. If this is the case, however, it may not diminish the value of these results in aiding terrorism prediction efforts. After all, oil prices may be capturing the impact of such an omitted variable, and thus serve as a useful proxy. An omitted variable problem would, of course, affect any policy implications of these results.

Dealing with the omitted variable concern is non-trivial. This is exacerbated by the fact that, to my knowledge, not even in the political science literature has attempted to approach a time-series of widely-varying terrorist events to look for common political or social causes that might serve as an omitted variable. Analysis of terrorist activities generally take the form of scrutinizing a particular event or cluster of terrorist events (for instance, the violence in Lebanon during the 1980s) and looking into the recent past for political events that might have triggered that violence

---

[9] Suppose, for example, that some pattern of behavior by Western Nations – perhaps political or economic pressure or military activity – contributes to a drop in oil prices, and also directly angers terrorist organizations and leads them to respond with attacks. The above empirical results would be consistent with this scenario, but would the relation between oil prices and terrorism would not necessarily be causal. Ehrlich and Liu (2002) do in fact argue that U.S. efforts to maintain oil supply stability (and by extension keep prices low) have in fact inflamed some Muslims.

(in this case, Israel's invasion of Lebanon and the increased presence of U.S. forces in Beirut in the aftermath). Such an approach is entirely reasonable.

However, for the purpose of ascribing common causes to a variety of disparate events, this type of analysis has at least one major shortcoming: the political events that are deemed important triggers of terrorism are chosen only *after* the terrorist attacks take place. That is, history remembers what "caused" a terrorist attack only if the attack occurs in the first place. A political event that might *a priori* seem likely to spark violence but does not is quickly forgotten, whereas a similar event that is followed by a wave of terrorism will certainly be remembered and prominently recorded in history. Therefore, constructing a time-series of important political events that may play a role in terrorism is likely to be extremely biased toward a collection of events that are, indeed, followed by terrorism. This means that the β coefficient on a "political event" variable would be biased as well. (Assuming such a variable could be satisfactorily constructed in the first place. Political history is a continuum that does not lend itself to quantitative pigeon-holing.) This is essentially a measurement problem, and may explain why there are no time-series analyses of this sort in the literature.

What are the implications of this discussion on the present analysis of oil and terrorism? First is that if Middle Eastern political factors are indeed an omitted variable, it would be extremely difficult to establish this, and the measurement error could ultimately obscure any true relationship between oil and terrorism. Furthermore, remember that there are two requirements for there to be omitted variable bias: One is for the omitted variable to be correlated with the dependent variable, which is what I have discussed to this point, and the other is for the omitted variable to correlated with the independent variable in question – in this case changes in oil prices. While political factors are almost certainly correlated with terrorism (even if it may be difficult to quantitatively establish this), it is far from clear that political events are linked with oil prices.

Hamilton (see for example, 2003) is one prominent proponent of the argument that political events in the Middle East can trigger movements in oil prices. The idea here is that major political events in the Middle East can either a.) interrupt production and lower oil supplies directly or b.) increase uncertainties in oil supplies, which boost precautionary demand. Either of these outcomes will yield higher oil prices in the aftermath of these political events. However, Barsky and Killian (2004) compile a list of many political events resembling the ones suggested by Hamilton that are not in fact followed by increases in oil prices. While such evidence does not by any means close the door on possible systematic correlation between political events in the Middle East and oil prices, it does highlight the difficulty alluded to above in choosing relevant political events that might be important, and the bias this could potentially bring into the current analysis.

5.2 Potential Causal Linkages

For the remainder of the paper, let's put aside this omitted variable concern. Doing so, we are left with the question of what causal chain might exist between changes in oil prices and Middle Eastern terrorism against the West. It would seem that we can quickly rule out reverse causality: that is, the causal direction must be from changing oil prices to terrorism, not terrorism to oil prices. Due to the timing conventions employed here, the latter would imply that oil prices fall in anticipation of terrorism. Given the potentially destabilizing effect of a Middle East terrorist attack on oil supplies, *a priori* one would expect that if the market "knows" a terrorist act is imminent, the price would be expected to rise. Oil prices falling in anticipation of attacks does not seem reasonable.

Going in the other direction, I see at least two potential explanations for how changes in oil prices might contribute to the frequency of Middle Eastern terrorism against Western targets. The first of these explanations, which seems more likely at this time, calls on the expected utility theories common throughout research into the economics of criminal behavior. In the other, falling oil prices directly stimulate certain agents to engage in or otherwise promote terrorism in an effort to interrupt the fall in prices. In what follows I will informally describe these hypotheses, and discuss the merits and demerits of each. While neither hypothesis is fully consistent with all the facts, they seem to be intuitively reasonable explanations for the empirical relationship we observe.

5.2.1 Expected Utility Hypothesis

Since Becker's landmark article "Crime and Punishment" (1968), economists studying crime have usually sought to describe criminal behavior as being the outcome of some rational decision process. While Becker's focus was in how negative incentives (i.e. punishments) might influence an individual's choice as to whether to commit a crime, the basic approach has flavored criminal research more generally to this day. According to this philosophy, a prospective criminal, when making the calculation of whether to proceed with his crime, considers the utility he expects to receive from the crime, and compares it to the utility he expects to receive from not participating.

Consider an individual who, for whatever reason, is predisposed to commit an act of terrorism, but has not yet done so. By predisposed, I mean that his utility increases (or he believes it will increase) as a result of becoming a terrorist. As he goes forward and is deciding whether to actually perform a terrorist act, his deliberations might be interpreted in the above light.[10] If his (subjective) expected utility from terrorism exceeds his expected utility from not engaging in terrorism (and instead engaging in legitimate economic (or non-economic) activities), we would predict that he would participate in the terrorist act. We might down write this decision rule in a straightforward way as:

$$Commit\,Terrorist\,Act\,if: E_0\left[\sum_{t=0}^{\infty}\beta^t[U_t\,|Terror\,attack\,at\,t=0]\right] > E_0\left[\sum_{t=0}^{\infty}\beta^t[U_t\,|No\,Terror\,attack\,at\,t=0]\right].$$

[10] Sandler et al [1983] and Grossman [1991] build theoretical models terrorist behavior in a manner consistent with this philosophy, though the details of their approaches differ from the one suggested below.

At time 0, the prospective terrorist considers his lifetime stream of expected utility (discounted at some rate β between 0 and 1) in two states of the world: the world where he commits the terrorist act at time zero and the one where he does not. Now suppose we have a society of heterogeneous agents that follow this decision rule. While most individuals would never consider committing a terrorist act, for those that might be particularly inclined to terrorism (and have a large left hand side to this decision rule) – but have not yet felt compelled to – the right-hand side of the expression may only be slightly larger than the left. For these individuals, new information or events that either raise the value of the left hand side or lowers the value of the right hand side of their decision may trigger them to commit a terrorist act. On average, we would thus expect a higher incidence of terrorism if either the left or right sides of the expression changed appropriately for that society.

Note that it is possible that if there are no long-term repercussions from engaging in the terrorist act, (for instance, say the perpetrator does not expect to be captured or injured), then it follows that $[U_t|\text{Terror at } t=0] = [U_t|\text{No Terror at } t=0]$ for all $t>0$. That is, he can engage in legal activities for the rest of his life in the same way as he would if he hadn't committed the terrorist act. The opposite extreme would be the case of a potential a suicide-terrorist, who can no longer engage in activities, economic or otherwise, after an attack.[11, 12]

[11] This does not necessarily mean that $E[U_t \mid \text{Terror at } t=0]= 0$ for $t>0$ in the case of the suicide bomber, as he may expect positive utility even in periods following his death.

In connecting this expected utility argument to changes in Middle Eastern oil prices, first note that the economic performance of many Middle Eastern nations is closely related to the price of oil. This is most readily apparent from inspection of Figure 1.6, which plots the average annual (nominal) Persian Gulf crude price and the nominal per-capita GDP of a prominent oil-producing nation, Saudi Arabia. The two series track each other well, with a correlation coefficient of about 0.8. There is not an obvious strong trend to either series, so the correlation is not likely spurious.

To the extent that an individual's economic prospects are impacted by the ebbs and flows of the broader economy, one would expect a strengthening or weakening economy to raise or lower, respectively, the average person's expected utility from participating in legitimate economic activities. If, say, falling oil prices typically contributes to deteriorating economic conditions, as is apparently the case in oil producing nations, then an individual's expectations of his own future economic well-being will often be tied to the price of oil.

---

[12] While some readers may recoil at the idea of ascribing any form of rationality to terrorists, from an economist's perspective it seems to be a sensible approach. Certainly, the behavior of terrorists, especially suicide attackers, would seem to mark the epitome of irrationality. However, I would argue that what is more likely is that the decision to be a terrorist is rational, but only given a utility function that a non-terrorist would think of as being bizarre, or evil. That is, the apparent irrationality of a suicide attacker actually results from the inability of a non-terrorist to comprehend what utility could possibly be gained by such actions. To the prospective terrorist, however, the gains from terrorism may be perfectly clear. It is in understanding an individual's underlying motivations – or how a person's utility function is formed – where economics usually breaks down, and this is no exception. Other disciplines, such as psychology, sociology, or political science, can aid in understanding how the terrorist inclinations are developed. Given that these inclinations are there, however, the economist can then make inferences as to behavior, which is what I suggest below.

Even those not in the work force may feel the impact of falling oil prices, as unusually generous state-sponsored welfare programs – which oil-rich nations are famous for – will be under increasing pressures. In Saudi Arabia, for instance, there is a significant linkage between the price of oil and government expenditures. Figure 1.7 plots this relationship since 1974. The correlation over the whole time period is .30, which is non-negligible. Since 1990, the correlation between non-defense spending and oil prices is even stronger, at .60.

While this reasoning might hold well within oil-producing nations, it is less clear how a person in a Middle Easterner country that does not produce oil might be impacted. The channel in this case could be changes in the level of charitable donations received by those in non-oil producing nations. Many Muslims are annually required to pay *Khums*, a charitable donation to the needy equal to 1/5 of their income. This includes income derived from mines or minerals. Charitable flows out of the oil producing nations would thus be profoundly affected by changes in oil prices. Those who receive those donations, primarily the poor, would in turn suffer during downturns.

This proposed impact on charitable giving from lower oil prices is not merely conjecture. Rooney (2006) views high oil revenues as being the primary source of funds to charitable organizations (both legitimate charities, and terrorist fronts) in the Middle East. He argues that elevated oil prices leads to a large pool of charitable funds. A report by the Institute for the Analysis of Global Security (2004) reiterates the linkage between oil prices and charitable giving in the Middle East.

If a drop in oil prices impacts the area's economies immediately through the channels suggested above, then a typical person's economic outcomes will be negatively impacted immediately. Even if the transmission mechanism between changes in oil prices to economic outcomes is a slow one, expectations of future economic wellbeing will be hurt by these falling oil prices. In either case, the right hand side of the expected utility decision rule will fall, and for some individuals on the margin who are contemplating committing a terrorist act by using the above decision rule, this change may lead them to act.

One might also argue for this expected utility theory from the other direction. That is, falling oil prices could also *raise* the left hand side of the above decision rule. It is well known that in many Middle Eastern nations there is significant underlying enmity towards the United States and other Western European nations. In such an environment, as the price of oil falls and economic prospects deteriorate, this would increase the general level of frustration in society, and some of these anti-Western sentiments may more readily surface. As this happens, the society as a whole may begin to view strikes against Western targets more favorably. For an individual considering whether or not to commit a terrorist act, then even if he is not personally impacted by weakening economic conditions, we would expect that a social climate that views the terrorist activities more favorably would *increase* the expected utility from participating in such an act. Such an explanation would seem to be consistent with findings of Atran (2003) who shows that suicide bombers are typically motivated by social and institutional pressures, not necessarily personal economic pressures.

In sum, whether it is by raising the utility from engaging in a terrorist act or lowering the expected utility from participating in legitimate economic activities, by this expected utility argument we would expect a decrease in oil prices to increase the number of individuals willing to become terrorists. In such an environment terrorist organizations would be better able to recruit members, develop plans of attack, and, as frustration with the economy rises, may even garner (tacit or explicit) support from the local government or business community. The perception that the West, particularly the United States, actively uses its economic and military power in an effort to control Middle East oil supplies may serve inflame anti-Western sentiment in the face of falling crude prices. (See, for example, Ehrlich and Jianguo (2002).)

Turned around, the argument seems even more compelling: when oil prices are rising and many in the Middle East are prospering as a result, there will be less urgency or desire within these nations to rock the boat and strike Western targets at that time (risking military and/or economic reprisal), and less tolerance for those who desire to do so. By the same token, when the local economies are doing well, the opportunity cost of committing a terrorist act – and forgoing legitimate economic activities – is larger.

As mentioned, Atran's (2003) central thesis – as well as Krueger and Maleckova's (2003) – is that terrorists are frequently not those who are the worst off economically. First, I'll note that, with the exception of those directly impacted by changes in flows of charitable giving, the expected utility argument does not talk about an agent's overall level of economic well-being but changes in those levels. That is, in principal, poverty is *not* a necessary condition for an individual to be "tipped" into committing a terrorist act. This distinction implies that the results here do not necessarily contradict the findings that poverty is not a good predictor of whether an agent is likely to commit terrorist acts.

In addition, note that the expected utility argument requires there to be a strong, pre-existing anti-Western sentiment. Clearly, if the utility derived from attacking a Western target is zero, then no matter how poorly the economy is performing or how low expected utility from legitimate economic activities may be, we would not expect any anti-Western terrorism to occur. Indeed, many developing regions of the world experience economic downturns that do not result in increasing numbers of terrorist activities against the West. It would seem that there is something special in the Middle East – where decidedly negative attitudes towards the West are coupled with fringes of violent Islamic Fundamentalism. Under these conditions, an economic downturn may allow these sentiments to flourish, with higher levels of terrorism following.

Recent empirical work by Blomberg and Hess (2002) and Neumeyer (2003) lends some support to this expected utility argument, as both cross-country studies find that deteriorating economic conditions leads to increased instability. While Neumeyer is focused on criminal behavior more generally, Blomberg and Hess's findings are directly relevant to this discussion, as they find evidence that argue that internal and external political conflict is considerably more common during recession than expansion. They find that this linkage is particularly strong for developing and non-democratic nations, where the probability of conflict more than doubles. A paper by Blomberg, Hess and Weerapana (2002) also directly tests whether the incidence of "terrorism from within" increases following recessions, and find that it does. However, the findings of the latter are that the connection between the business cycle and terrorism is only significant for high income democracies, which is not the pattern observed here.

Krueger and Maleckova (2002) provide evidence counter to this proposed relationship between economic expectations and terrorism. They find that six months prior to the September 2000 intifada in Israel economic expectations in the occupied territories were relatively high and unemployment relatively low. That the intifada occurred after these favorable observations suggests that economics had little to do with the rising hostilities. Unfortunately, the economic data they provide is not available closer to the actual initiation of violence, and given the volatility of the unemployment data shown, a rapid deterioration in economic conditions is not out of the question, especially given the bursting of Israel's technology bubble and subsequent recession that began in mid-2000, only months before the start of the intifada. Even Atran (2003) who in general is skeptical of expected utility explanations of terrorist behavior, concedes that there was at least one episode of increased violence in Israel that followed significant deteriorating economic conditions amongst the Palestinians. While Krueger and Maleckova do provide some time-series evidence on a *positive* relationship between GDP and terrorism levels in Israel during an earlier time period, they note that their finding is not robust to small changes in specification.

Another significant drawback to this expected utility hypothesis is that the planning and preparation for a major terrorist attack may take place over a long time-period. The attacks on September 11, 2001, for instance, are believed to have been planned and prepared for over a period of years. The individuals carrying out this attack doubtless decided to participate in it long before they knew what oil prices would be around September 11. Therefore, while an expected-utility decision rule might have been in place for these terrorists, the price of oil near September 11 would not likely have been a factor in it. For smaller-scale attacks – such as kidnappings, car bombings, gun attacks against restaurants or tourist groups – that would seem to require less preparation, the argument still seems reasonable.

Indeed, running the regression analysis of Section 4 on a sub-sample of "low-preparation"[13] events, the coefficients are even larger than on the full sample of events, suggesting even greater sensitivity for these types of attacks to changes in oil prices. (See Table 1.5.) The same analysis performed on "high-preparation" events

[13] There is inevitably some arbitrariness in the selection of "low-" vs. "high-" preparation events, but an effort was made to break events into categories based on the author's perception of the amount of planning that would be required. For example, parking a car full of explosives next to an embassy is probably not going to require much planning, whereas an armed assault on a European airport probably required considerably training and preparation.

yields smaller odds-ratios, which, while still greater than one (that is, indicating the same pattern as for the overall) are no longer statistically significant, due to a combination of the smaller coefficients and the fewer number of events.

The expected utility argument could be further bolstered by scrutiny of higher-frequency data on movements in output and incomes in Middle Eastern nations. While the annual averages of per capita GDP cited in Figure 1.6 indicate a tight connection between oil prices and economic performance in the Middle East, using quarterly or monthly data could help determine whether the timing of changes in these variables mirrors that of oil prices, and might themselves be used in a similar way to help predict terrorist activity against the United States. Data on actual living standards (say, a measure of median income), rather than the blanket average of per capita income, may also be illuminating, since the fluctuations in oil prices may have a differential impact on the various strata of the economic hierarchy. Unfortunately, this type of detailed time series data is scarce in developing regions.

Furthermore, if an environment of deteriorating economic prospects does indeed prompt higher levels of terrorism, we might be able to find similar trends in other areas. That is, would we find that, for example, terrorist activities originating in Chechnya seem to be correlated with the economic travails of the former Soviet Union? As mentioned above, Blomberg, Hess and Weerapana's (2002) cross-country time-series results hint that we may, though their economic data (annual) does not lend itself to the higher-frequency, focused analysis performed here. While, as mentioned, Krueger and Maleckova's (2002) cross-section results in Israel and

Lebanon suggest that poverty has little to do with terrorist activity, much of their analysis simply lacks the time-series dimension that would allow one to conclude that changes in the economic situation are irrelevant. That is, the question is still open as to whether terrorist strikes against Israel are more frequent when material conditions in the West Bank and Gaza Strip worsen.

5.2.2 Business Practices Hypothesis

Oil prices take a more active role in the second hypothesis for why oil prices and frequency of terrorism may be linked. This is a simple, argument. When oil prices are falling, this is harmful to oil producing interests, who see their profit margins deteriorate. Oil traders are sensitive to issues of political stability in the Middle East, as so much of the world's oil supply comes from there. Thus, some oil producers may reason that promoting the appearance of instability in the Middle East may help support higher oil prices. Sponsoring terrorism against Western targets may be one way to accomplish this. In other words, a terrorist attack can be thought of as a business response to deteriorating industrial conditions.

On the one hand, this argument would seem to be immune to one key concern of the above expected utility argument – that long preparation times before major terrorist events makes it unreasonable to surmise that the participants are being influenced by oil fluctuations around an event date. While a long planning period may still be required for major attacks like 9/11, the decision to proceed on a certain date could be made by the sponsors of the attack based on recent news. That is, after the training and planning for a major attack has taken place, the leadership can elect to proceed with the attack at relatively short notice in response to changing oil markets.

On the other hand, if terrorism is indeed actively sponsored by oil-interests in an attempt to raise prices during downturns, it does not appear to have been an especially successful strategy. Inspection of Figure 1.2 suggests that the downward trend in average oil prices prior to events does not reverse itself in the short-run following these terrorist events. This finding is also reflected when scrutinizing price changes on an event by event basis (not shown). In general, whatever direction oil prices were moving prior to an event is how they tend to move in the months following an event. Another mark against this hypothesis is that by instigating attacks, particularly major ones, the sponsors risk military or economic reprisal that may lower profitability.

It should be noted, however, that while the downward price trend is not abruptly reverse on average following these events, it does appear (in Figure 1.2) that on average the rate of decrease may slow following the attacks. Also, since historically a large fraction of the funding for Middle East terrorist organizations comes out of some of the biggest oil producing states (e.g. Saudi Arabia, Iran, Iraq, Libya) a hypothesis of the sort that links terrorism to the oil industry should not be completely discounted at this point. Scrutiny of futures market data may yet indicate a spike following events that suggests that oil prices move upward in the months following these terrorist events, which might allow oil interests "in-the-know" to earn a profit. Ideally, detailed data on the local oil industry's costs and revenues could help determine whether or not the profitability of oil producers in the region tends to increase or stabilize in the months following terrorist events. This kind of data is not likely to be forthcoming, however.

6 Conclusion

There appears to be a statistical connection between changes in real crude oil prices and terrorist strikes by Middle Eastern organizations against the United States and Western Europe. Particularly, attacks are more common after prices drop than after they rise. This connection allows us to use a logistic specification to deduce the probability of attack by Middle Eastern agents against the West based solely on the recent history of oil price movements. The empirical model's predictions as to whether or not a terrorist event will occur seems to do a good job qualitatively of describing the actual history of terrorist events of this type. While the relative differences in the probability of an attack implied by the empirical results are large for reasonable changes in oil prices, the probabilities themselves remain small, and as such these results can only supplement efforts to anticipate terrorist action.

While a univariate relationship alone cannot hope to establish causality, the findings here would be consistent with at least two reasonable causal hypotheses, particularly the standard philosophy underlying theories of the economics of criminal behavior. At the very least, the results here call for further investigation of both the phenomena described here, and whether a pattern of comovements in relevant economic variables and terrorist activities is demonstrable elsewhere. It may turn out that where cross-sectional analysis has thus far failed to demonstrate a linkage between economic factors and terrorism, a time-series approach may yet succeed.

Table 1.1
Listing and Description of Qualifying[1] Terrorist Events
January 1974 - May 2003

Event	Date	Attributed Date	Description
British Airways Hijacking 1	3/3/1974	Mar-74	British Airways flight to London hijacked by Arab National Yourth Organisation for the Liberation of Palestine. 102 Passengers aboard
British Airways Hijacking 2	11/23/1974	Nov-74	British DC-10 airliner hijacked at Dubai, UAE, by Palestinian Rejectionist front terrorists and eventually flown to Tunisia where a German passenger was killed.
Attack on Orly Airport, France	1/19/1975	Jan-75	Arab terrorists attack Orly airport, Paris, France, seizing ten hostages in a terminal bathroom. Eventually the French provided the terrorists with a plane to fly them to safety in Baghdad, Iraq.
Air France hijacking 1	6/27/1976	Jul-76	An Air France airliner is hijacked by a joint German Baader-Meinhof/Popular Front for the Liberation of Palestine terrorist group and its crew are forced to fly to Entebbe airport in Uganda. Some two hundred and fifty eight passengers and crew are held hostage.
Lufthansa hijacking	10/13/1977	Oct-77	Four Palestinian terrorists hijack a German Lufthansa Boeing 737 and order it to fly around a number of Middle East destinations for four days. All the ninety hostages are eventually rescued and three terrorists killed.
Alitalia hijaacking	9/7/1979	Sep-79	3 Shi'ite Muslims hijack flight protesting disappearance of leader Musa-al-Sadr
Iran Embassy Attack/Hostages	11/4/1979	Nov-79	Iranian radicals seize the US Embassy in Tehran, taking sixty-six American diplomats hostage.
French Lebanon Embassy Bombing	5/24/1982	May-82	Car bombing outside French Embassy, 1 employee and 12 others killed.
Attack on Paris Synagogue	8/9/1982	Oct-82	Abu-Nidal attacks on Paris Synagogue and Jewish Restaurant. 6 killed, 27 injured.
U.S. Lebanon Embassy Bombing	4/18/83	Apr-83	Sixty three people, including the CIA's Middle East Director, are killed and 120 injured in a 400 lb suicide truck bomb attack on the US Embassy in Beirut, Lebanon. Islamic Jihad claims responsibility.
Beirut Military Compound Bombing	10/23/1983	Oct-83	Simultaneous suicide truck bombs on American and French compounds in Beirut, Lebanon. A 12,000 lb bomb destroys a US Marine Corps base killing two hundred and forty one Americans. Islamic Jihad claims responsibility.
Attack on US Embassy in Kuwait	12/12/1983	Dec-83	US Embassy in Kuwait was targeted by Iranian backed Iraqi Shia terrorist who attempted to destroy the building with a truck bomb. The attack was foiled by guards and the device exploded in the Embassy fore-court killing five people.
Restaurant Bombing in Torrejon, Spain	4/12/1984	Apr-84	Eighteen US servicemen killed and eighty three people injured in bomb attack on restaurant near USAF base in Torrejon, Spain. Responsibility claimed by Hezbollah as revenge for March bombing in Beirut attributed to the US government.
Air France hijacking 2	7/31/1984	Aug-84	Islamic Jihad hijacks Air France jetliner out of West Germany and diverts to Iran.
Bombing of US Embassy in East Beirut	9/20/1984	Sep-84	Suicide bomb attack on US Embassy in East Beirut kills twenty three people and injures twenty one others. Attributed to the Iranian backed Hezbollah group.
TWA Flight 847	6/1/1985	Jun-85	TWA Boeing 727 was hijacked enroute to Rome, Italy, from Athens, Greece, by two Lebanese Hezbollah terrorists. Two thirds of passengers American.
Achille Lauro Hijacking	10/7/1985	Oct-85	Four Palestinian Liberation Front terrorists seize the Italian cruise liner, Achille Lauro, during a cruise in the eastern Mediterranean, taking more than seven hundred people hostage.
Bomb attacks in Paris	12/1985	Dec-85	Series of bomb attacks in Paris attributed to Hezbollah.
Attacks on Italian, Austrian Airports	12/27/1985	Jan-86	Suicide grenade and gun attacks against passenger terminals at Rome and Vienna airports by the Abu Nidal terrorist group results in sixteen people being killed and more than 100 civilians injured.
Berlin Discotheque Bombing	4/5/1986	Apr-86	Two US soldiers are killed and seventy nine American servicemen are wounded in Libyan bomb attack on a night club in West Berlin, Germany.
Pan Am Flight 73 Hijacking	9/5/1986	Sep-86	Hijacking of Pan Am fllight 73 out of Pakistan. Attirubted to Abu Nidal. 3 Americans killed, amongst others.
Pan Am/Lockerbie Bombing	12/21/1988	Dec-88	Pan Am Boeing 747 blown up over Lockerbie, Scotland. Agents from Libya deemed responsible.
UTA Airliner Bombing	9/19/1989	Sep-89	One hundred and seventy people killed when French UTA airliner explodes in mid-air over Niger. The French government issued warrants for the arrest of four Libyans.

Table 1.1 (Continued)
Listing and Description of Qualifying[1] Terrorist Events
January 1974 - May 2003

Event	Date	Attributed Date	Description
WTC Bombing	2/26/1993	Mar-93	World Trade Centre in New York, USA, badly damaged by a massive bomb planted by Islamic terrorists. The car bomb was planted in an underground garage and left six people dead and more than one thousand people injured.
Belgian Relief Center Bombing in Iraq	12/13/1993	Dec-93	Belgian Relief Center Bombed in Iraq. 5 Killed.
Tour Bus Attack in Cairo	12/27/1993	Jan-94	Attack on tour bus in Cairo. 8 Austrians Wounded.
Saudi Arabia Military Compound Bombing	11/14/1995	Nov-95	Seven foreigners, including a number of US servicemen, are killed in bomb attack on US-run National Guard training centre at Riyadh, Saudi Arabia.
Kidnapping of French Tourists	1/26/1996	Feb-96	Al-Aslam tribesmen kidnap 17 French tourists. (WHERE?)
Attack on Greek Tourists	4/18/1996	Apr-96	Greek tourists attacked in Cairo (CASUALTIES?)
Saudi Arabia Khobar Towers Bombing	6/25/1996	Jul-96	Islamic radical terrorists opposed to the western military presence in the Gulf region, explode a truck bomb next to a USAF housing area at Dhahran, Saudi Arabia, killing 19 American servicemen and 385 injuring more.
Tourist Kidnapping in Yemen 1	3/4/1997	Mar-97	Yemeni tribesmen kidnap 6 German tourists
Tourist Kidnapping in Yemen 2	8/13/1997	Aug-97	Yemeni tribesmen kidnap 10 Italian tourists
Western tourists attacked in Egypt 1	9/18/1997	Sep-97	Egyptian gunmen attack tourist bus in front of Egyptian National Antiquities Museum in Cairo. 9 German tourists killed.
Western tourists attacked in Egypt 2	11/17/1997	Nov-97	58 western tourists killed and dozens injured in gun attack on party visiting historic monuments in southern Egypt. Attack blamed on Islamic guerillas after six are killed in shoot out with police.
Tourist Kidnapping in Yemen 3	6/18/1998	Jun-98	Yemen tribesmen kidnap 9 Italian Tourists
Africa Embassy Bombings	8/7/1998	Aug-98	US Embassies in Nairobi, Kenya, and Dar-es-Salem, Tanzania, heavily damaged by massive bomb attacks. In the Nairobi attack 247 people were killed, including 12 Americans, and 4,000 injured. Ten people were killed and 74 injured in Tanzania incident. US in
Western tourists attacked in Yemen	12/28/1998	Jan-99	Yemini militants kidnap a group of western tourists, including 12 Britons, 2 Americans, and 2 Australians on the main road to Aden. Four victims were killed during a rescue attempt the next day.
USS Cole Attack	10/12/2000	Oct-00	USS Cole bombed in Persian gulf by Al Qaida operatives. 17 soldiers killed, 37 wounded.
WTC and Pentagon Strikes	9/11/2001	Sep-01	Al Qaida operatives commandeer 4 airplanes and crash 3 of them into World Trade Center and Pentagon. Both towers of WTC collapse, killing thousands.
Suicide bombing in Saudi Arabia	10/6/2001	Oct-01	Suicide Bomber in shopping area in Saudi arabia kills 1 American, injuring 2 American, 2 Britons and 2 Filipinos
Bali Bombing	10/14/2002	Oct-02	Series of bombings in nightclub district of Bali, Indonesia, frequented predominantly by Australians and Western Europeans. 15 Australians killed, amongst others.
Riyadh Bombing	5/5/2003	May-03	Suicide bombers detonate explosives in American housing compounds in Riyadh, Saudi Arabia. 8 Americans killed, 34 injured.

Note:

[1]The terrorist events listed above meet the U.S. government's definition of a terrorist event and, for the purposes of this paper, satisfy 3 additional criteria. These are A) the apparent principle target of the attack was civilians or non-mobilized military personnel of Western Nations, B) the attack was instigated by groups with Middle Eastern origins, and C) the attack was of significant size. By significant size, I impose the requirement that at least 5 individuals are held hostage, kidnapped, injured or killed in a given event.

Source:

There are various sources of data for the events list above, including the Center for Defense and International Security Studies Database, the Almanac of Modern Terrorism, the State Department's Patterns of Global Terrorism, and Mickolas (various years, Terrorism in the 1980s series).

Table 1.2
Percent Differences in Oil Prices Prior to Terrorist Events
Real Price of Landed Persian Gulf Crude Oil
January 1974 - May 2003

Event	Attributed Date	Percent Difference in Price k-Months Prior to Event Relative to Event Date[1]											
		12	11	10	9	8	7	6	5	4	3	2	1
British Airways Hijacking 1	Mar-74											-23.02	0
British Airways Hijacking 2	Nov-74			-3.961	19.06	13.606	8.6246	8.5449	7.3017	4.4163	2.3758	-1.833	0
Attack on Orly Airport, France	Jan-75	-4.622	18.399	12.945	7.9638	7.8842	6.6409	3.7556	1.715	-2.493	-0.661	0.4531	0
Air France hijacking 1	Jul-76	-0.906	0.8907	-2.617	2.5203	4.0891	2.0476	1.7729	1.2164	1.1181	0.7908	-0.037	0
Lufthansa hijacking	Oct-77	-2.628	-2.825	-3.007	-0.535	0.5413	1.1396	0.5562	0.5914	-0.928	-0.398	-0.952	0
Alitalia hijaacking	Sep-79	-35.11	-35.99	-36.07	-36.37	-32.53	-32.25	-26.6	-21.43	-19.88	-3.759	-1.093	0
Iran Embassy Attack/Hostages	Nov-79	-41.48	-41.78	-37.95	-37.67	-32.01	-26.84	-25.29	-9.17	-6.504	-5.411	1.78	0
French Lebanon Embassy Bombing	May-82	4.4401	4.8722	1.2571	0.5034	-0.135	3.1551	3.4294	2.4646	0.8431	-0.176	1.7664	0
Attack on Paris Synagogue	Oct-82	2.6262	1.988	5.2779	5.5522	4.5874	2.9658	1.9467	3.8892	2.1228	-1.765	2.0177	0
U.S. Lebanon Embassy Bombing	Apr-83	19.279	15.391	19.174	17.156	14.041	14.986	16.15	16.522	17.945	9.0588	4.5648	0
Beirut Military Compound Bombing	Oct-83	21.038	21.409	22.832	13.946	9.4521	4.8874	2.5628	2.5311	2.803	2.7042	0.2993	0
Attack on US Embassy in Kuwait	Dec-83	23.249	14.363	9.869	5.3042	2.9796	2.948	3.2198	3.121	0.7161	0.4169	0.8736	0
Restaurant Bombing in Torrejon, Spain	Apr-84	3.9547	3.9231	4.1949	4.0961	1.6912	1.392	1.8486	0.9751	-2.157	-0.17	-2.631	0
Air France hijacking 2	Aug-84	5.5363	5.237	5.6937	4.8202	1.6877	3.6756	1.2138	3.8451	2.9823	2.2493	2.0355	0
Bombing of US Embassy in East Beirut	Sep-84	2.8062	3.2629	2.3894	-0.743	1.2448	-1.217	1.4143	0.5515	-0.182	-0.395	-2.431	0
TWA Flight 847	Jun-85	14.312	12.276	14.707	12.679	10.393	9.3961	7.3109	8.8311	9.8556	7.2239	5.1735	0
Achille Lauro Hijacking	Oct-85	9.1519	8.1546	6.0693	7.5896	8.614	5.9823	3.9319	-1.242	-6.028	-4.831	-2.896	0
Bomb attacks in Paris	Dec-85	12.678	14.198	15.223	12.591	10.541	5.3671	0.5806	1.7781	3.7127	6.6086	5.5715	0
Attacks on Italian, Austrian Airports	Jan-86	30.522	31.546	28.914	26.864	21.69	16.904	18.101	20.036	22.932	21.895	16.323	0
Berlin Discotheque Bombing	Apr-86	75.147	69.974	65.187	66.385	68.319	71.215	70.178	64.607	48.283	19.261	2.3982	0
Pan Am Flight 73 Hijacking	Sep-86	88.617	87.58	82.009	65.685	36.663	19.8	17.402	14.489	-2.202	-3.19	1.0466	0
Pan Am/Lockerbie Bombing	Dec-88	25.639	16.26	16.732	16.74	19.089	20.456	15.573	10.326	6.4703	-2.965	-3.959	0
UTA Airliner Bombing	Sep-89	-27.59	-28.58	-24.62	-10.22	-1.841	2.8677	6.8533	15.058	7.767	4.097	2.1049	0
Arco Chemical Factory Explosion	Jul-90	11.822	9.7166	12.408	17.053	18.504	24.939	23.786	17.606	7.1129	-6.032	-10.55	0
WTC Bombing	Mar-93	-4.07	7.0179	10.436	14.458	11.872	9.4657	8.1856	5.3833	-0.644	-3.274	-4.458	0
Belgian Relief Center Bombing in Iraq	Dec-93	26.364	25.18	29.638	30.995	29.916	24.34	20.386	13.899	14.565	9.0553	8.3697	0
Tour Bus Attack in Cairo	Jan-94	37.393	41.851	43.208	42.129	36.553	32.599	26.112	26.779	21.268	20.583	12.213	0
Saudi Arabia Military Compound Bombing	Nov-95	2.0316	-0.138	6.6895	8.8002	10.12	15.781	11.485	3.8476	-0.003	1.5619	2.5448	0
Kidnapping of French Tourists	Feb-96	-1.95	-0.63	5.0316	0.7348	-6.902	-10.75	-9.188	-8.205	-10.75	-5.267	-6.885	0
Attack on Greek Tourists	Apr-96	-5.076	-9.373	-17.01	-20.86	-19.3	-18.31	-20.86	-15.37	-16.99	-10.11	-5.173	0
Saudi Arabia Khobar Towers Bombing	Jul-96	-14.66	-13.09	-12.11	-14.65	-9.17	-10.79	-3.903	1.0314	6.2044	9.1048	3.7082	0
Tourist Kidnapping in Yemen 1	Mar-97	10.393	13.293	7.8968	4.1886	8.0617	12.567	18.313	22.423	21.547	18.74	13.212	0
Tourist Kidnapping in Yemen 2	Aug-97	22.06	27.806	31.916	31.039	28.232	22.705	9.4924	8.16	3.8123	1.8334	0.3578	0
Western tourists attacked in Egypt 1	Sep-97	20.867	24.977	24.101	21.294	15.767	2.5542	1.2218	-3.126	-5.105	-6.58	-6.938	0
Western tourists attacked in Egypt 2	Nov-97	28.474	25.667	20.139	6.9268	5.5944	1.2467	-0.732	-2.208	-2.566	4.3726	3.1353	0
Tourist Kidnapping in Yemen 3	Jun-98	46.094	45.736	52.674	51.437	48.302	38.448	33.79	22.411	18.744	13.765	7.1773	0
Africa Embassy Bombings	Aug-98	52.2	50.962	47.827	37.974	33.316	21.936	18.269	13.29	6.7027	-0.475	-2.546	0
Western tourists attacked in Yemen	Jan-99	42.504	38.837	33.857	27.27	20.093	18.021	20.567	22.861	27.445	22.244	11.225	0
USS Cole Attack	Oct-00	-27.93	-23.65	-19.52	-9.952	-12.14	-14.71	-15.55	-10.21	-9.411	-6.833	-5.302	0
WTC and Pentagon Strikes	Sep-01	32.447	26.099	20.892	12.979	8.0081	5.3653	3.0222	1.6949	2.6045	1.2929	1.1554	0
Suicide bombing in Saudi Arabia	Oct-01	36.516	31.309	23.396	18.425	15.782	13.439	12.112	13.022	11.71	11.572	10.417	0
Bali Bombing	Oct-02	-39.67	-44.94	-46.92	-34.93	-20.77	-9.042	-5.879	-5.118	-4.528	0.3334	3.1292	0
Riyadh Bombing	May-03	5.529	6.1193	10.98	13.776	10.647	7.1298	10.77	18.317	26.216	27.711	21.296	0
Average for Event Months		12.4	12.3	11.7	11.0	9.6	7.9	7.0	7.0	5.0	3.7	1.5	0.0
Average for Other Months		0.3	0.1	0.0	-0.1	-0.1	0.0	-0.1	-0.3	-0.1	-0.1	0.0	0.0
P-value for difference (t-test)		0.01	0.01	0.01	0.00	0.00	0.01	0.01	0.00	0.02	0.01	0.21	
Median for Event Months		10.39	12.28	10.71	10.70	9.03	5.67	3.84	3.85	2.70	1.04	1.16	0.00
Median for Other Months		2.02	1.18	1.41	2.21	2.16	1.77	0.94	0.87	0.91	0.39	0.31	0.00
P-value for difference (Pearson chi-squared test)		0.06	0.01	0.00	0.00	0.01	0.00	0.01	0.02	0.31	0.74	0.24	0.00
Number of events where price is 10% higher in month t-k than on reference date (month t-1)		21	21	22	21	20	15	15	15	10	8	6	
Number of events where price is 10% lower in month t-k than on reference date (month t-1)		6	6	7	6	5	6	4	3	3	1	2	

[1] For an event occurring in month t, each cell in the table is calculated as $\ln(price_{t-k}) - \ln(Price_{t-1})$.

Source: Oil Price Data from Energy Information Administration, Terrorism Dates assembled from CDISS, Patterns of Global Terrorism yearbook, Mickolas.

Table 1.3a
Summary of Logit Results (Odds Ratios)
Based on Changes in Persian Gulf Crude Prices
Dependent Variable: Event Dummy

	Independent Variable: Percent Difference of Real Oil Price in Given Month Relative to Reference Month										
	Number of Months Before Month of Interest (k)										
Version	12	11	10	9	8	7	6	5	4	3	2
Univariate											
Coefficient (Odds Ratio)	1.016	1.017	1.018	1.020	1.019	1.018	1.019	1.025	1.025	1.034	1.040
Z-statistic	(2.67)**	(2.76)**	(2.77)**	(2.88)**	(2.88)**	(2.65)**	(2.64)**	(3.08)**	(2.42)*	(2.48)*	(1.28)
Including Event Lags (12 months)											
Coefficient (Odds Ratio)	1.015	1.017	1.018	1.018	1.017	1.016	1.017	1.022	1.023	1.032	1.061
Z-Statistic	(2.39)*	(2.53)*	(2.62)**	(2.57)*	(2.50)*	(2.16)*	(2.14)*	(2.56)*	(2.02)*	(2.12)*	(1.99)*
Dropping Interval 2/1/81-6/1/86											
Coefficient (Odds Ratio)	1.01478	1.01726	1.01849	1.02103	1.02163	1.02026	1.02149	1.02659	1.02413	1.03749	1.0382
Z-Statistic	(2.09)*	(2.23)*	(2.25)*	(2.45)*	(2.54)*	(2.34)*	(2.30)*	(2.69)**	(1.98)*	(2.21)*	(0.93)

Notes: * significant at 5%; ** significant at 1%
Z-statistics based on robust standard errors.

How to Read this Table: Each coefficient refers to the odds ratio associated with the specified independent variable in a univariate logistic regression. This table thus represents 11*3=33 regressions. For example, if we look at column 7 for the Persian Gulf Crude series, this is the odds ratio in the regression of our Event dummy variable on the percentage difference between the Persian Gulf price of oil in the reference month (which is one month before the month of interest) and the price of oil six months earlier. A one percent increase in the difference between the oil price over this time increases the odds off attack by .019.

Table 1.3b
Marginal Effect of Oil Price Changes on Terrorist Events (Implied by Logit Results)
Based on Changes in Persian Gulf Crude Prices
Dependent Variable: Event Dummy

	Independent Variable: Percent Difference of Real Oil Price in Given Month Relative to Reference Month										
	Number of Months Before Month of Interest (k)										
Version	12	11	10	9	8	7	6	5	4	3	2
	(%)										
Univariate											
Coefficient (Marginal Effect)	0.155	0.171	0.185	0.196	0.192	0.181	0.194	0.246	0.252	0.340	0.412
Z-statistic	(2.77)**	(2.86)**	(2.88)**	(2.94)**	(2.89)**	(2.65)**	(2.63)**	(3.04)**	(2.43)*	(2.48)*	(1.31)
Including Event Lags (12 months)											
Coefficient (Marginal Effect)	0.143	0.159	0.174	0.173	0.168	0.154	0.164	0.216	0.221	0.313	0.577
Z-Statistic	(2.42)*	(2.57)*	(2.66)**	(2.58)*	(2.47)*	(2.13)*	(2.12)*	(2.52)*	(2.01)*	(2.10)*	(2.03)*
Dropping Interval 2/1/81-6/1/86											
Coefficient (Odds Ratio)	0.126	0.145	0.16	0.179	0.185	0.176	0.186	0.227	0.211	0.32	0.346
Z-Statistic	(2.20)*	(2.38)*	(2.41)*	(2.60)**	(2.62)**	(2.38)*	(2.33)*	(2.66)**	(2.00)*	(2.25)*	(0.96)

Notes: * significant at 5%; ** significant at 1%.
Z-statistics based on robust standard errors.
Marginal Effects are calculated at the mean value of the price deviation, which is roughly zero for each variable.

How to Read this Table: Each logit coefficient refers to the Marginal Effect associated with the specified independent variable in a univariate logistic regression on the event dummy variable. This table thus represents 11*2=22 regressions. For example, if we look at column 7 for the Persian Gulf Crude series, this is marginal effect of a one percent increase in the six month price difference on the probability of attack next month. Here, this one percent difference increases the probability of attack next month by about .181 percentage points

Table 1.4a
Summary of Logit Results (Odds Ratios)
Normalized Price Differences
Dependent Variable: Event Dummy

	Independent Variable: Percent Difference of Real Oil Price in Given Month Relative to Reference Month										
	Number of Months Before Month of Interest (k)										
Version	12	11	10	9	8	7	6	5	4	3	2
Univariate											
Coefficient (Odds Ratio)	1.657	1.695	1.693	1.703	1.646	1.558	1.558	1.632	1.522	1.477	1.250
Z-statistic	(2.91)**	(3.00)**	(3.01)**	(3.17)**	(3.22)**	(3.06)**	(3.06)**	(3.43)**	(2.72)**	(2.53)*	(1.08)
Including Event Lags (12 months)											
Coefficient (Odds Ratio)	1.50919	1.54458	1.57236	1.52773	1.47175	1.38904	1.38049	1.46896	1.38149	1.40565	1.4412
Z-Statistic	(2.39)*	(2.53)*	(2.62)**	(2.57)*	(2.50)*	(2.16)*	(2.14)*	(2.56)*	(2.02)*	(2.12)*	(1.99)*

Notes: * significant at 5%; ** significant at 1%
Z-statistics based on robust standard errors.

How to Read this Table: Each coefficient refers to the odds ratio associated with the specified independent variable in a univariate logistic regression. This table thus represents 11*2=22 regressions. For example, if we look at column 7 for the Persian Gulf Crude series, this is the odds ratio for the regression of our Event dummy variable on the (standardized) percent difference between the Persian Gulf oil price in the reference month and the price of oil six months earlier. If prices fall by one standard deviation over six months the odds off attack increase by .021.

Table 1.4b
Marginal Effect of Oil Price Changes on Terrorist Events (Implied by Logit Results)
Normalized Price Differences
Dependent Variable: Event Dummy

	Independent Variable: Percent Difference of Real Oil Price in Given Month Relative to Reference Month										
	Number of Months Before Month of Interest (k)										
Price Change Variable (i)	12	11	10	9	8	7	6	5	4	3	2
	(%)										
Univariate											
Coefficient (Marginal Effect)	4.355	4.492	4.544	4.542	4.299	3.898	3.835	4.070	3.326	2.887	1.187
Z-statistic	(3.05)**	(3.15)**	(3.11)**	(3.19)**	(3.16)**	(2.99)**	(2.96)**	(3.21)**	(2.49)*	(2.21)*	(0.64)
Including Event Lags (12 months)											
Coefficient (Marginal Effect)	4.35	4.531	4.648	4.643	4.314	3.818	3.763	4.266	3.633	3.629	2.553
Z-Statistic	(2.77)**	(2.86)**	(2.88)**	(2.94)**	(2.89)**	(2.65)**	(2.63)**	(3.04)**	(2.43)*	(2.48)*	(1.31)

Notes: * significant at 5%; ** significant at 1%.
Z-statistics based on robust standard errors.
Marginal Effects are calculated at the mean value of the price deviation, which is slightly greater than zero for each variable.

How to Read this Table: Each logit coefficient refers to the Marginal Effect associated with the specified independent variable in a univariate logistic regression on the event dummy variable. This table thus represents 11*2=22 regressions. For example, if we look at column 7 for the Persian Gulf Crude series, this is marginal effect of a one standard deviation increase in the six month price difference on the probability of attack next month. Here, if prices fall by one standard deviation the probability of attack next month rises by about 3.9 percentage points.

Table 1.5
Summary of Logit Results For Selected Subsets of Terrorist Events
Persian Gulf Crude
Dependent Variable: Event Dummy

	Independent Variable: Percent Difference of Real Oil Price in Given Month Relative to Reference Month — Number of Months Before Month of Interest (k)										
Regression Version	12	11	10	9	8	7	6	5	4	3	2
All Events (43 events)											
Coefficient (Odds Ratio)	1.016	1.017	1.018	1.020	1.019	1.018	1.019	1.025	1.025	1.034	1.040
Z-statistic	(2.67)**	(2.76)**	(2.77)**	(2.88)**	(2.88)**	(2.65)**	(2.64)**	(3.08)**	(2.42)*	(2.48)*	(1.28)
"Small Scale" Events (22 events)											
Coefficient (Odds Ratio)	1.019	1.021	1.023	1.022	1.023	1.022	1.024	1.031	1.039	1.058	1.102
Z-Statistic	(2.96)**	(2.94)**	(2.91)**	(2.71)**	(2.79)**	(2.54)*	(2.60)**	(2.92)**	(2.79)**	(3.14)**	(2.84)**
"Large Scale" Events (21 events)											
Coefficient (Odds Ratio)	1.00595	1.00729	1.00751	1.0106	1.00974	1.00968	1.01009	1.0136	1.00349	0.99783	0.96395
Z-Statistic	(0.68)	(0.78)	(0.82)	(1.24)	(1.28)	(1.29)	(1.27)	(1.74)	(0.37)	(0.17)	(1.02)

Notes: * significant at 5%; ** significant at 1%
Z-statistics in parenthesis based on robust standard errors.

The sub-sample groupings of events were categorized as follows:

"Small Scale" Sub-category that includes events that would seem to have relatively short preparation times, such as suicide car bombings, or gun attacks on un-fortified targets in the Middle East (tourist groups or museums).

"Large Scale" Sub-category that includes events that would seem to have relatively long preparation times, such as airline hijackings, massive bombings, or gun-attacks against secured facilities (airports, embassies outside the middle east).

Table 1.6
Alternative Specification: Including Squared Change in Oil Price (Odds Ratios)
Dependent Variable: Event Dummy

	Independent Variable: Percent Difference of Real Oil Price in Given Month Relative to Reference Month — Number of Months Before Month of Interest (k)										
Variable	12	11	10	9	8	7	6	5	4	3	2
Oil Price Change	1.01781	1.01945	1.02097	1.02601	1.03028	1.02906	1.03076	1.04378	1.03716	1.05779	1.03157
Z-statistic	(2.24)*	(2.09)*	(2.06)*	(2.18)*	(2.35)*	(2.27)*	(2.26)*	(2.77)**	(2.31)*	(2.75)**	(1.16)
Oil Price Change Squared	0.99993	0.99994	0.99993	0.99983	0.99967	0.99962	0.99957	0.99931	0.99944	0.99865	1.0014
Z-Statistic	(0.52)	(0.41)	(0.43)	(0.95)	(1.51)	(1.40)	(1.25)	(1.39)	(1.31)	(2.26)*	(0.91)

Notes: * significant at 5%; ** significant at 1%
Z-statistics based on robust standard errors.

How to Read this Table: Each column reflects a single multivariate logit, including as independent variables the percentage change in oil price over horizon k, and the squared change in price over that same interval.

Figure 1.1
Plot of Real Persian Gulf Crude Price and Terrorist Event Dates
January 1974 - May 2003

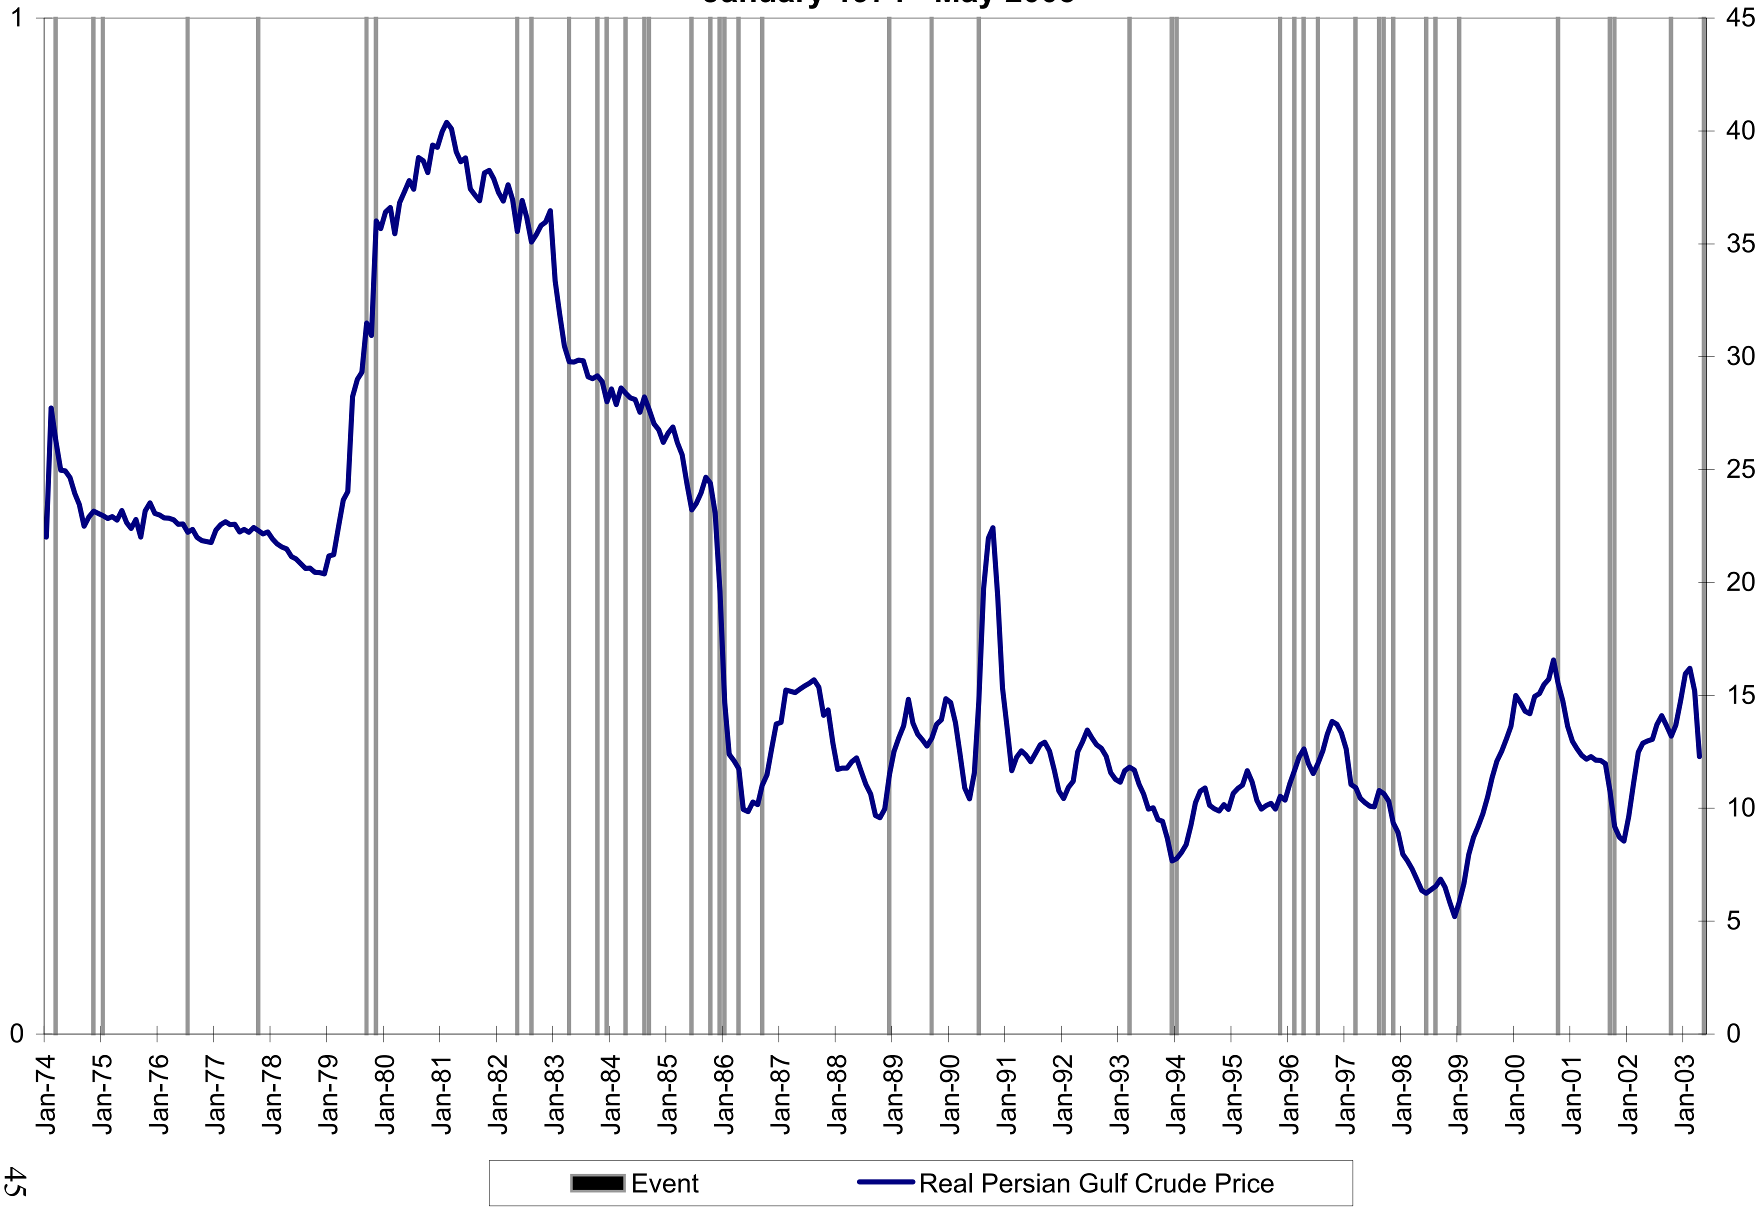

Figure 1.2
Average Percent Difference in Real Persian Gulf Oil Price
Relative to Reference Month
Event Months vs. Non-Event Months
With 95% Confidence Intervals

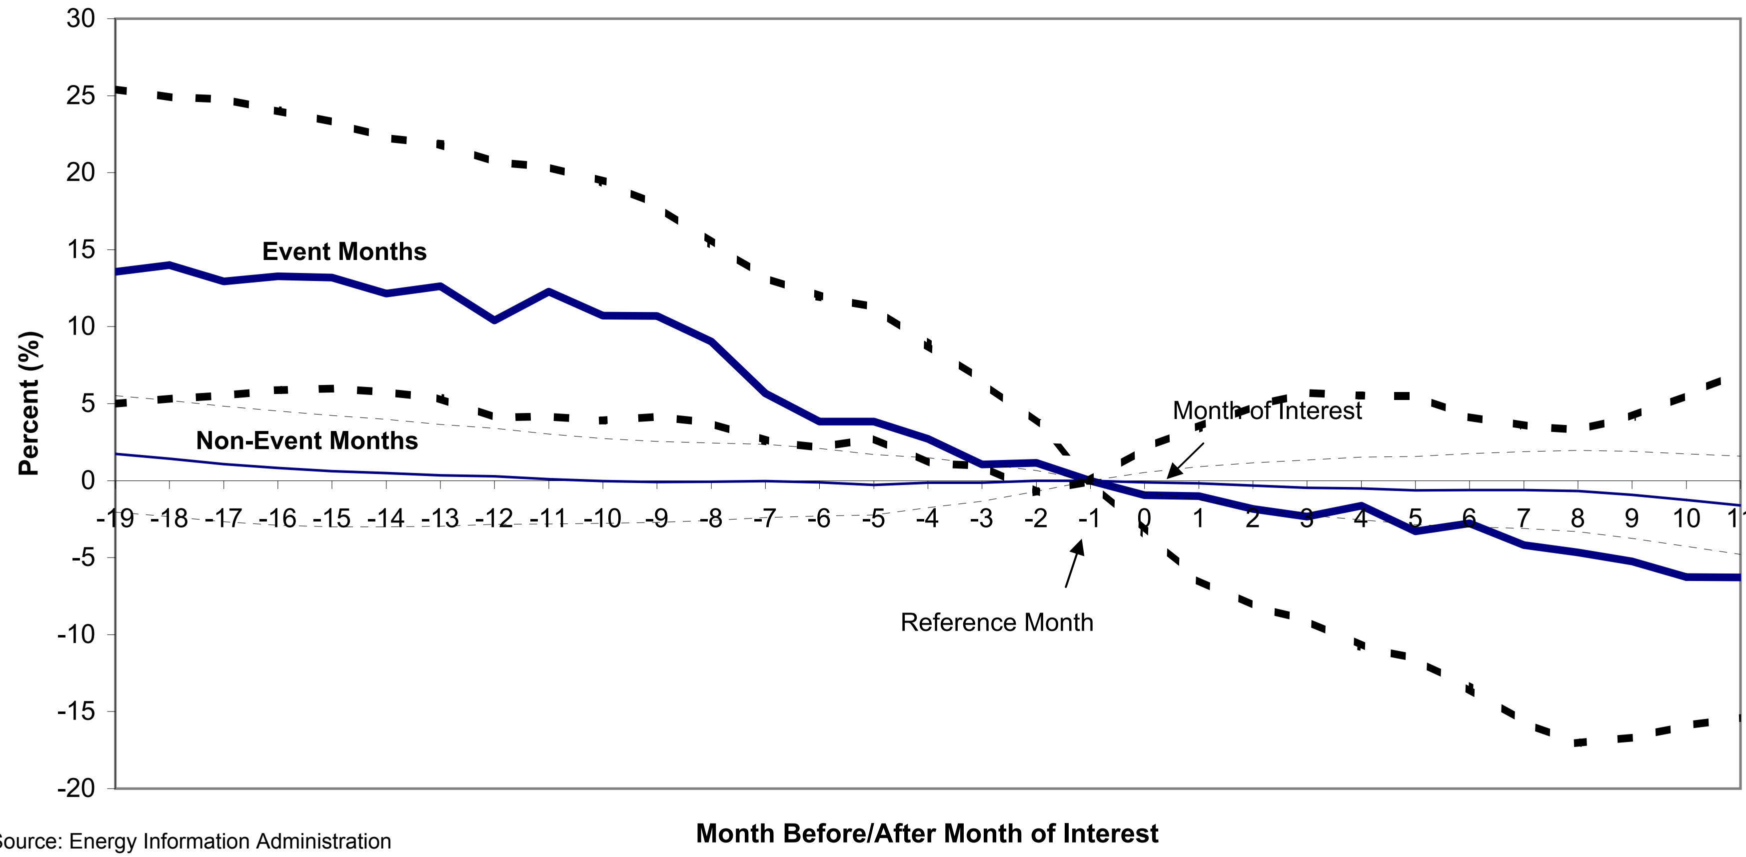

Source: Energy Information Administration

Figure 1.3
Predicted Probability of Terrorist Attack "Next Month"
Using Persian Gulf Crude Prices 3 and 9 Month Ago
(Based on Logit Results)

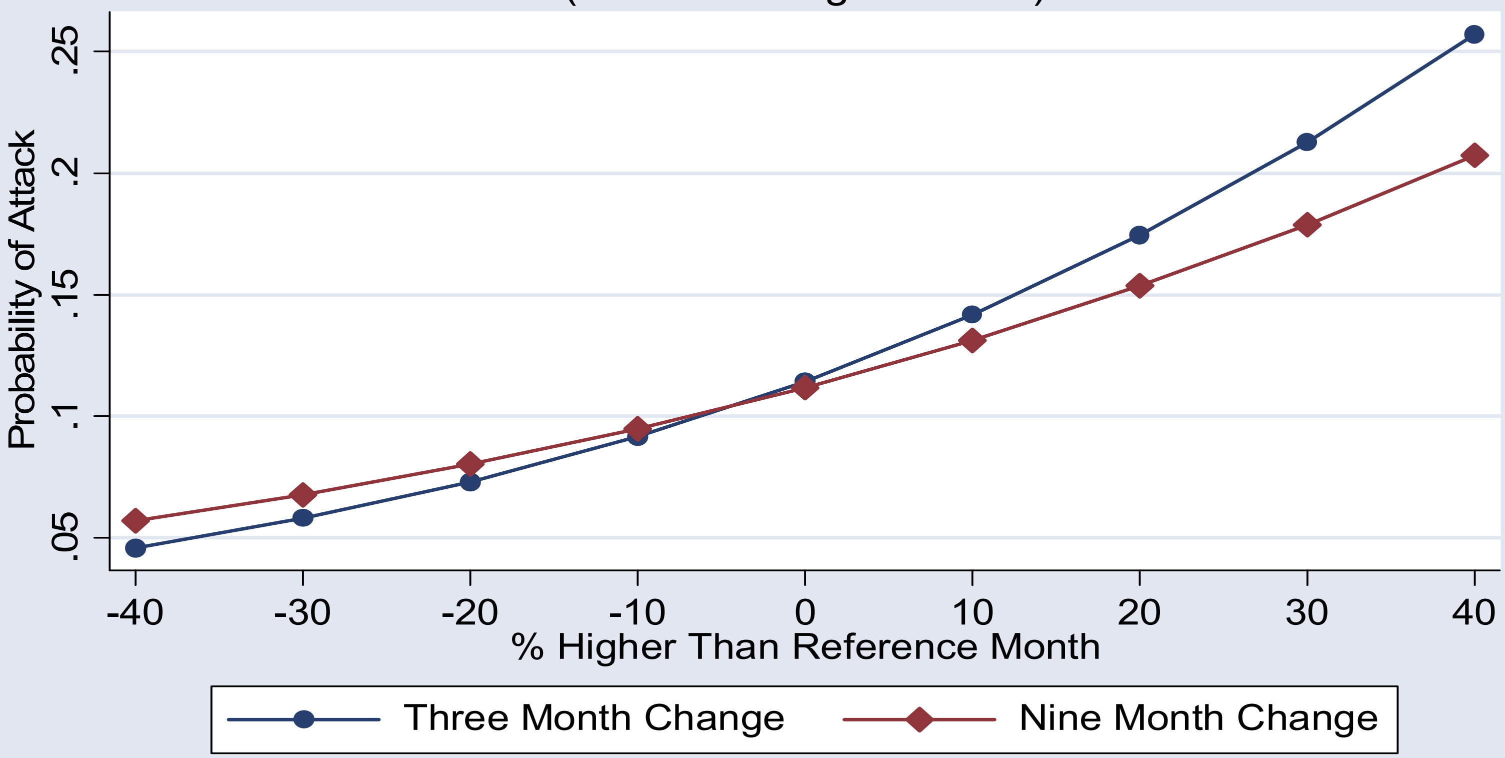

Figure 1.4
Predicted Probability of Terrorist Attack "Next Month"
Using Persian Gulf Crude Prices 3 and 9 Month Ago
(Based on Logit Results, Normalized Prices)

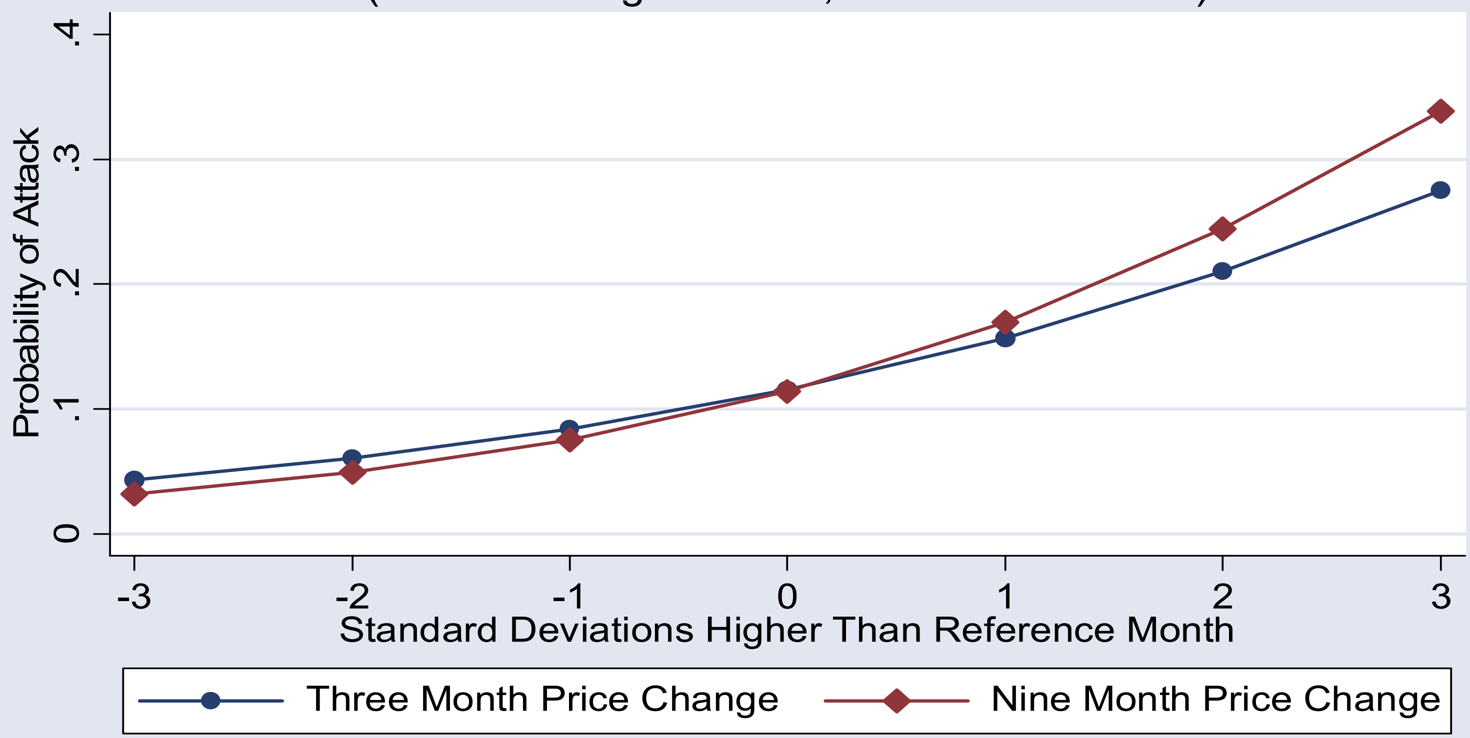

Figure 1.5

Implied Probability of Attack and Actual Terrorist Event Dates

Based on Logit of Persian Gulf Crude Prices changes at 9-Month Horizon

January 1974 - May 2003

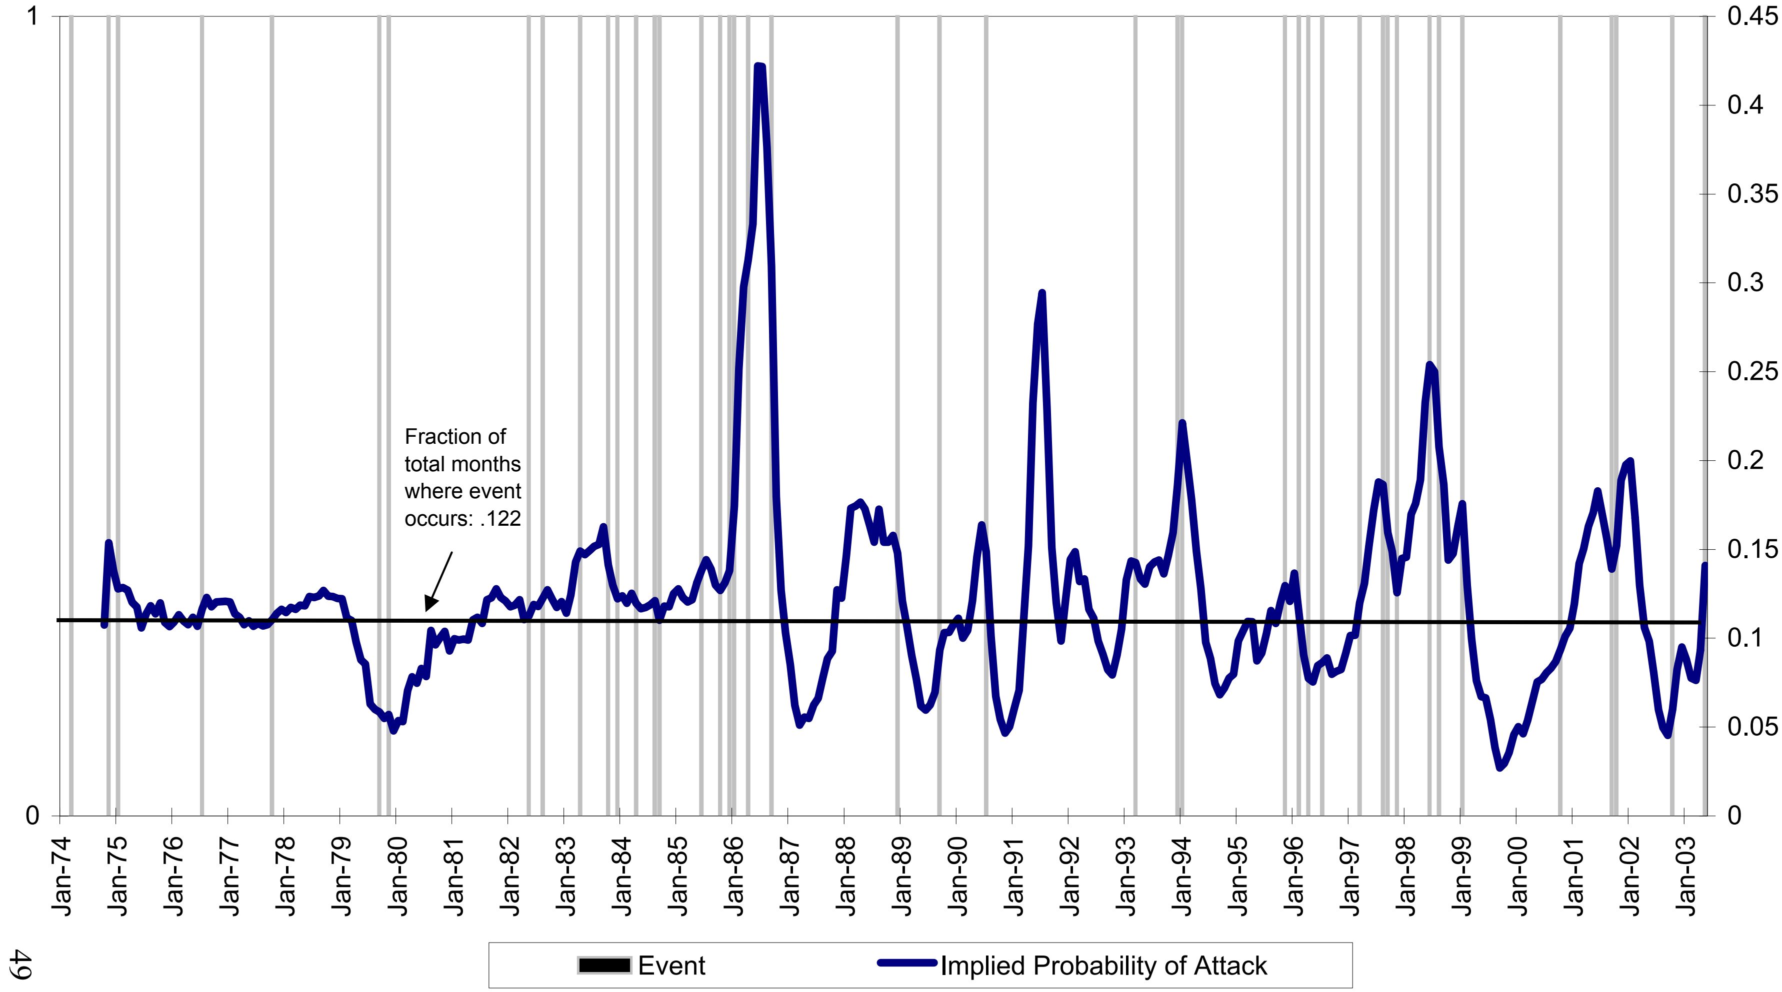

Figure 1.6
Saudi Arabia
Nominal Persian Gulf Crude Oil Prices vs. Nominal GDP per capita
1974-2002

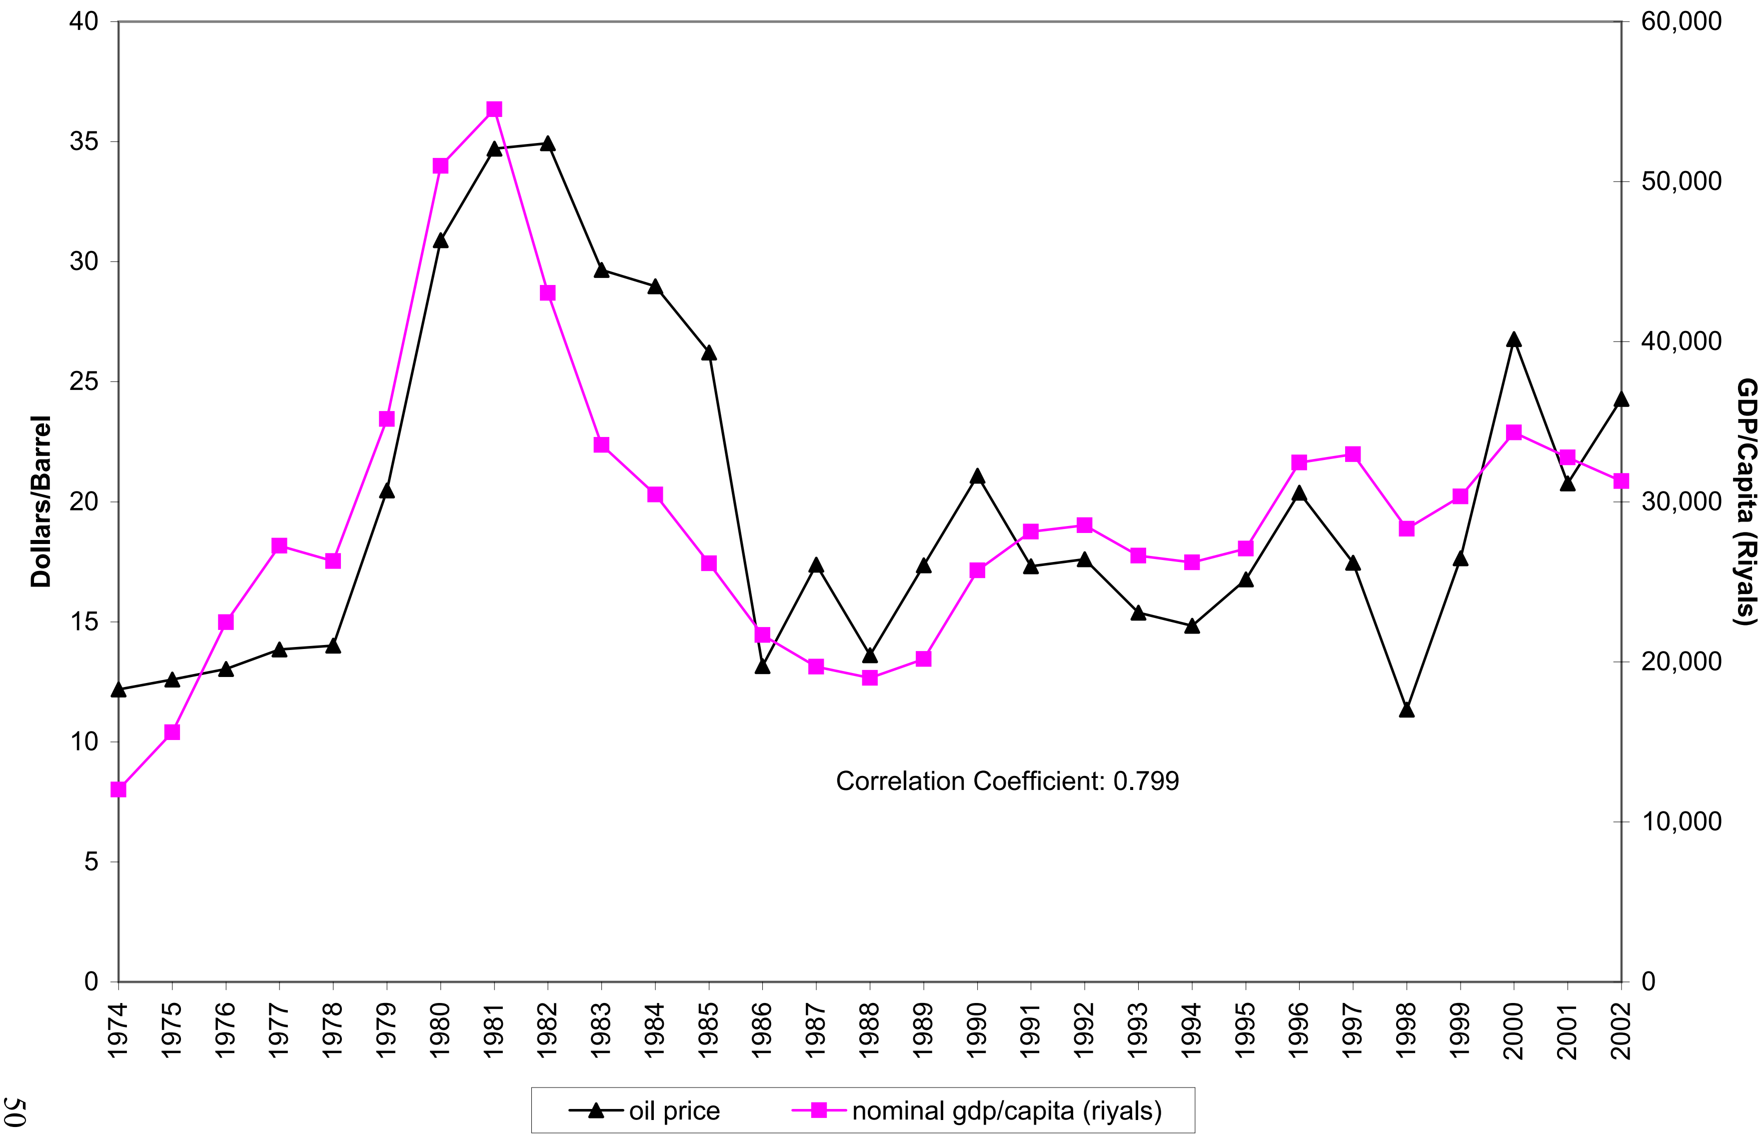

Figure 1.7
Saudi Arabia
Nominal Persian Gulf Crude Oil Prices vs. Total Government Expenditures
1974-2002

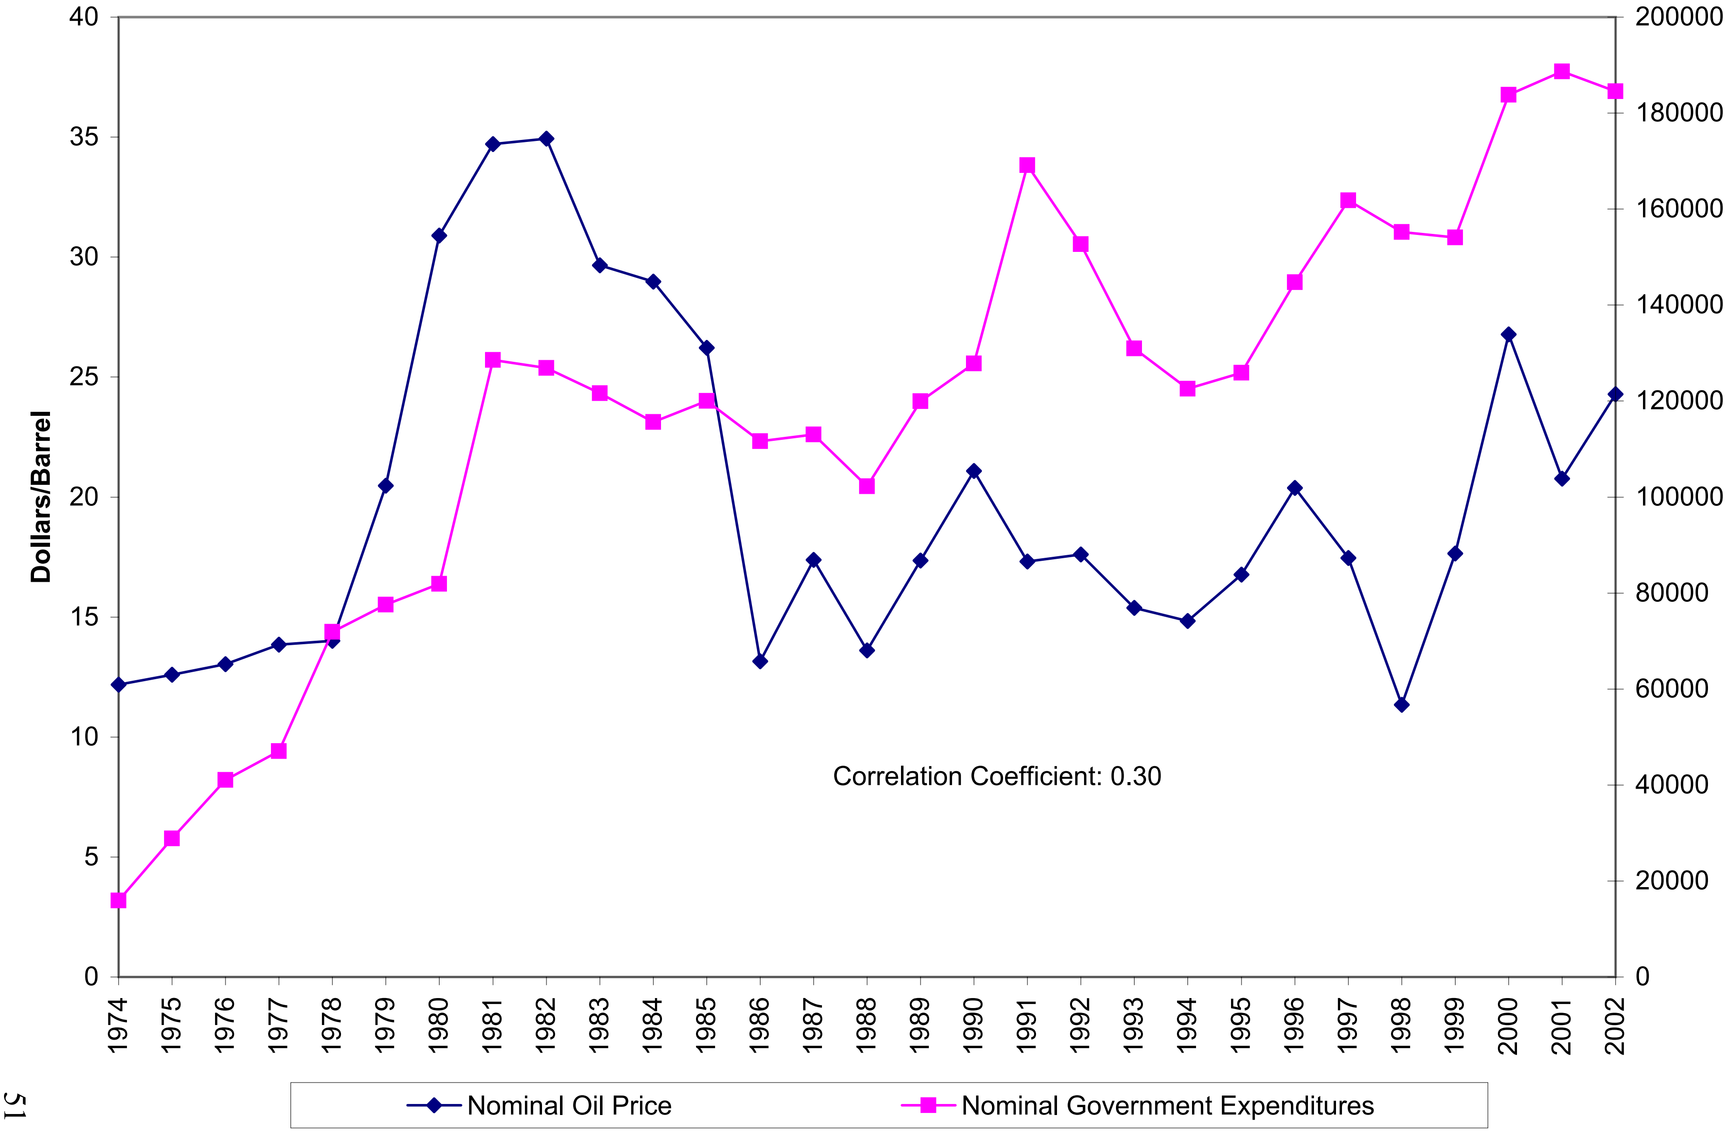

References

Atran, Scott. 2003. Genesis of Suicide Terrorism. *Science.* Vol. 299.

Barsky, Robert B, and Lutz Kilian. 2004. Oil and the Macroeconomy Since the 1970s. *Journal of Economic Perspectives,* forthcoming.

Becker, Gary S. 1968. Crime and Punishment: An Economic Approach. *Journal of Political Economy*. University of Chicago.

Blomberg, S. Brock, and Gregory D. Hess. 2002. The Temporal Links between Conflict and Economic Activity. *Journal of Conflict Resolution*. Volume 46:1 pp. 74-90.

Blomberg, S. Brock, Gregory D. Hess, and Akila Weerpana. (August 2002 version). Terrorism From Within: An Economic Model of Terrorism. Claremont Colleges Working paper.

Ehrlich, P.R. & Jianguo Liu. 2002. Some Roots of Terrorism. *Population and Environment.* Volume 24:2: 183-192.

Energy Information Administration. *Monthly Energy Review*.

Grossman, Herschel I. 1991. A General Equilibrium Model of Insurrections. *American Economic Review*. Volume 81:4.

Hamilton, James D. 2003. What is an Oil Shock? *Journal of Econometrics.* Vol. 113, April.

Institute for the Analysis of Global Security. 2004(?) Fueling Terror. <http://www.iags.org/fuelingterror.html>

Krueger, Alan & Maleckova, Jitka. 2002. Education, Poverty, Political Violence and Terrorism: Is There a Causal Connection? NBER Working Paper (9074).

Long, J. Scott, and Jeremy Freese. 2003. *Regression Models for Categorical Dependent Variables Using Stata*. College Station, TX: Stata Press.

Looney, Robert. 2006. The Mirage of Terrorist Financing: The Case of Islamic Charities. *Strategic Insights*. Volume V, No. 3. March 2006.

Mickolus, Edward F. 1989. *International Terrorism in the 1980s, A Chronology of Events*. Ames: Iowa State University Press.

Neumeyer, Eric. 2003. Good Policy Can Lower Violent Crime: Evidence from a Cross-National Panel of Homicide Rates, 1980–97. *Journal of Peace Research*. Volume 40:6. pp. 619–640

Sandler, Todd, John T. Tschirhart, and Caley, Jon (1983). A Theoretical Analysis of Transnational Terrorism. *The American Political Science Review*. Volume 77:1.

Shafritz, Jay, E.F. Gibbons, and Gregory Scott. 1991. *Almanac of Modern Terrorism*. Facts on File. NY.

United States Department of State. *Patterns of Global Terrorism*. Years 1991-2001.

Chapter 2

One for Some or One for All? Taylor Rules and Interregional Heterogeneity

Abstract

We document a very robust empirical phenomenon: both the US Federal Reserve and the European Central Bank appear to set interest rates partly in response to regional disparities in unemployment rates. This result is exceedingly robust – particularly for the US – even after controlling for a wide variety of factors, including the central bank's information set and a battery of explanatory variables. Furthermore, including measures of inter-regional unemployment dispersion in Taylor rule estimates greatly improves the identification of the central banks' responses to aggregate inflation and unemployment rates. Moreover, inclusion of the variance of unemployment across regions brings each bank's policies with respect to macroeconomic *aggregates* into alignment with each other. We propose three models in which central bank policymaking is influenced by disparities across regions. Testing specific implications of these models suggest that each bank's approach to policy may differ in fundamental ways.

1 Introduction

The traditional approach to monetary policy research focuses on aggregate measures of economic variables. Inflation, output, unemployment, productivity, and other components of these policy models are typically considered across the economy as a whole. As the monetary policy instrument – typically the interest rate chosen by the central bank – applies uniformly across the economy, perhaps the tendency to focus on aggregate measures of economic performance is a natural one. Certainly, written statements from the central banks themselves do not contradict this approach. Reviewing the minutes from a typical U.S. Federal Reserve Open Market Committee (FOMC) meeting, for example, there is essentially no mention of differences in economic performance across regions of the country. The Federal Reserve Act itself further dictates that the FOMC's policies should be made "with regard to their bearing upon the *general* credit situation of the country" (Federal Reserve Act, section 12Ac, emphasis added), and that the Fed's objective is to support "the economy's long run potential to increase production, so as to promote effectively the goals of maximum employment, stable prices, and moderate long-term interest rates" (FRA, section 2A).

Despite this, to conclude that inter-regional variation in economic activity is irrelevant to monetary policy *in practice* may be premature. Certainly, such variation is in the policy-makers' information set. In the weeks prior to each FOMC meeting, the Federal Reserve Board releases the Beige Book, an informal survey of economic trends in different parts of the country. Much of the content of the Beige Book's summary compares the state of the economy across the differing regions.

Furthermore, at any time, five of the members of the FOMC represent reserve banks scattered across the nation. That the Federal Reserve Act requires representation from different regions suggests that, at least initially, getting regional perspectives on the country's economic performance was important.

Further evidence on this matter can also be found in transcripts of FOMC meetings themselves. While the brunt of each of these meetings is focused on economic aggregates, at some point in the proceedings the president of each member bank makes a short statement about the status of his region's economy. Topics in these presentations can range widely; from the outlook for certain industries, to housing prices, to labor market tightness. The FOMC would likely not devote any time to sharing region-specific data if only aggregate information is of importance.

The European Central Bank is unfortunately less transparent than the Fed. The ECB does not publish transcripts or minutes from the meetings of its Executive Board. This makes it more difficult to determine the extent to which regional variation is considered in ECB policy. However, given that the Eurozone is comprised of sovereign nations, it is not difficult to imagine that stronger national identity may lead to considerable attention being given to localized economic performance. Dixit (2000) goes so far as to propose that the threat of exiting the Eurozone may lead to ECB policy that accommodates the needs of individual members.

The question as to whether regional differences play a role in monetary policy is not purely an academic one. It is particularly important in light of the recent tendency of countries to surrender their national central banks in favor of multinational ones. In these cases, nations have agreed to give up independent monetary policies in exchange for increased exchange rate stability with major trading partners and as a mechanism to credibly lower inflationary expectations. Yet as central banks gain control of monetary policy over more heterogeneous economic entities, the pressure to respond to or accommodate off-cycle regions may rise. For example, with Central and East European countries being considered for admission to the Euro-Zone, the degree of heterogeneity in regional welfare would likely increase for the European Central Bank. Additionally, the formation of a central bank and common currency for the entire continent of Africa is on the horizon.

In light of these concerns, we first turn to the data to see if there is any evidence of central bank sensitivity to such heterogeneity. Our empirical analysis is in the spirit of Taylor (1993), who found that the Federal Reserve's interest rate could be adequately represented via a simple rule in which interest rates change mechanically with various aggregate variables, empirical work has treated Taylor rules as a baseline for modeling a central bank's behavior.[1] Our own empirical model is essentially an augmented Taylor rule, which incorporates inter-regional variation in addition to aggregate measures of unemployment and inflation. For our baseline model we use GMM to estimate Taylor rules for the U.S. Federal Reserve and European Central Bank and test whether various measures of inter-regional unemployment rate dispersion offers any additional explanatory power after taking account of expectations of future aggregate inflation and unemployment rates.

[1] See Clarida, Gali, and Gertler (2000) for such an example focusing on whether the behavior of US Federal Reserve has changed over time.

Our empirical results strongly reject the null that these central banks do not respond to the inter-regional dispersion of unemployment rates. The results hold for various measures of dispersion, such as gaps between high and low unemployment regions or the weighted variance of unemployment rates across regions for each period. Appealingly, inclusion of the dispersion measures actually brings the *aggregate* policies of both the Fed and ECB into closer alignment with one another.

We are concerned with confirming this new result. We begin by verifying that our findings are not a statistical anomaly. For this purpose, we consider whether our results are robust to the inclusion of a variety of leading indicators and other economic variables that could be relevant to the central banks' expectations for aggregate inflation and unemployment, such as stock prices, oil prices, and the PPI, and find that our results are largely robust to the inclusion of such measures. Another concern is that our unemployment dispersion measures could help forecast future values of inflation and unemployment. In this case, finding significant coefficients on the former could simply reflect a failure to adequately capture the central banks' expectations of future inflation, output growth, or unemployment. We address this concern in two ways. First, we show that there is little evidence that our measures of regional heterogeneity are useful predictors of future values of the aggregate variables. Second, we reproduce our estimates using Green Book forecasts of future aggregate variables from the Federal Reserve. Thus, we can control for the central bank's information set.

Even here, we continue to find an independent role for our inter-regional dispersion measures.

As we are unable to explain away the result on a statistical basis, we seek an alternative explanation to that proposed by our theory. We propose three theoretical models that may shed light on this surprising empirical finding. In the first, we explore the implications of a central bank that dislikes disparity in welfare across nations. In the second, we propose that each region has its own non-linear Phillips curve, whose slope is dependent on the magnitude of a region-specific shock. Lastly, we consider institutional factors that may contribute to this result. In particular, that certain regions may be overweighted in the policy decisions of the banks, which contributes to the finding that regional differences matter. For all these models, we evaluate the theoretical implications in light of the data to determine whether the models seem to be appropriate descriptions of the observed policy.

The paper is organized as follows. Section 2 contains the baseline empirical results that suggest that interregional heterogeneity does matter for policy. Because the result is so surprising, we make numerous attempts to confirm that the result is not a fluke, and that our analysis fails to account for the central bank's full information set. We are ultimately unable to make this result disappear. Section 3 offers three new theoretical models in which interregional heterogeneity is relevant to the policymaker's decisions. Here, we also empirically test specific predictions made by each model. Section 4 summarizes and concludes.

2 Baseline Results

This section presents the baseline results of an empirical analysis that examines whether the Federal Reserve or the European Central Bank tend to respond to inter-regional heterogeneity. The results indicate that the dispersion of unemployment across regions is indeed an important predictor of Fed and ECB policy. One of the more appealing features of the analysis is that including an unemployment dispersion variable actually clarifies each central bank's responses to *aggregate* inflation and unemployment.

2.1 Estimation strategy and data

Given that interest rates are the primary tool used by these central banks in achieving their goals, policy-makers' decisions are naturally modeled by an interest rate rule of the type proposed by Taylor (1993)

$$i_t = \phi_\pi E_t \pi_{t+j} + \phi_{ue} E_t ue_{t+j} + \sum_{i=1}^{T} \rho_i i_{t-i} + \varepsilon_t \qquad (2.1)$$

Such a rule implies that interest rates rise by ϕ_π (ϕ_{ue}) basis points on impact when expectations of inflation (unemployment) rise by one percentage point. Lagged interest rate terms are included to allow for interest-smoothing. The i.i.d. (by assumption) error term ε_t represents monetary policy shocks. This specification assumes that interest rates are set in response to current expectations of future values of the independent variables, capturing the well-known fact that monetary policy acts with a lag, forcing policy-makers to be forward-looking.

Because expectations are not always available, our primary estimation strategy will employ the additional assumption of rational expectations for policy-makers, such that $E_t x_{t+j}=x_{t+j}+v_{t+j}$ where v_{t+j} is unforecastable using time t information. Substituting this into equation (2.1) yields

$$i_t = \phi_\pi \pi_{t+j} + \phi_{ue} ue_{t+j} + \sum_{i=1}^{T} \rho_i i_{t-i} + \zeta_t \tag{2.2}$$

where ζ_t consists of the monetary policy shock and the sum of rational expectations errors. Thus, following from our assumptions, $E_{t-j}\zeta_t=0$ for all $j\geq 1$. Because future values of inflation and unemployment can be expected to be correlated with time-t monetary policy shocks, we follow Clarida, Gali, and Gertler (2000) and estimate the parameters of equation (2.2) by GMM.

We use monthly data from January 1982 to September 2005 for the U.S. and from January 1999 to October 2006 for the Euro-Zone. For the US, we use the effective federal funds rate as our primary measure of interest rates, the 12-month log percentage change in the CPI for inflation and the BLS series for aggregate unemployment rates. For the Euro-Zone, we use the interbank overnight rate for our interest rate series, and harmonized aggregate inflation and unemployment rates. We use a six-month forecast horizon, though the results are insensitive to this assumption.

Throughout the paper, we allude to the inter-regional dispersion of unemployment rates. In practice, we measure this dispersion in a variety of ways. The first is to compute the variance of the distribution of unemployment rates each month, weighted by the population share of each region (var(UE)). We define regions in the US as each of the fifty states plus the District of Columbia. For the

ECB, each region is one of the eleven member states. In addition, for the ECB, we subtract the mean unemployment rate of each country between 1990 and 1998 before taking the variance to account for the fact that European countries have vastly different structural UE rates.[2] Our second measure is to take the difference between the unemployment rates of the 90th and 10th percentiles of the time-*t* distribution of regional unemployment rates (UEP9010). The second is a narrower band: the difference between the 75th and 25th percentiles (UEP7525).[3]

Figure 2.1 plots interest rates, inflation, and aggregate unemployment for the US over our time sample, while Figure 2 plots our three measures of regional heterogeneity in UE rates. The first thing worth noting is that these measures are broadly similar (all cross-correlations exceed 0.9). All three track the aggregate unemployment rate, by rising in recessions and falling as aggregate unemployment falls. This property fails in two instances. The first is the 2000 recession: as aggregate unemployment rose with the recession, none of the dispersion measures changed in this time period, indicating that the last recession was borne similarly by all states. The second occurs in 1986, when the price of oil fell dramatically. This led to sizeable increases in unemployment in oil-producing states, which shows up as an upsurge in dispersion of unemployment rates with no commensurate increase in aggregate unemployment. Figures 3 and 4 plot aggregate variables and our measures of unemployment dispersion respectively for the Euro-Zone. Again, the three

[2] All results for the US are robust to removing mean UE rates from each state.
[3] For the ECB, we again remove the average UE rate of each country from 1990 to 1998 before calculating the latter two dispersion measures.

dispersion measures are broadly similar. Unlike with the US, these series are uncorrelated with aggregate unemployment.

To test whether these measures affect central bank decision-making, we augment equation (2.2) with a measure of regional dispersion of unemployment rates

$$i_t = \phi_\pi \pi_{t+j} + \phi_{ue} ue_{t+j} + \sum_{i=1}^{T} \rho_i i_{t-i} + \beta D_t + \zeta_t \tag{2.3}$$

where the null is that $\beta=0$ and D_t is a measure of regional unemployment dispersion. As instruments, we will consistently use six lags of the RHS.[4]

2.2 *Estimation Results*

Our baseline results are presented in Table 2.2.1. Consider first the results for the US (panel A), which exclude dispersion measures. The coefficient on future inflation is positive and highly significant, but, surprisingly, the coefficient on aggregate unemployment is only marginally statistically different from zero. The coefficients on lagged interest rate imply an important amount of interest smoothing (the sum of the coefficients is 0.98), yielding a long-run response to inflation of about 4, consistent with the post-1982 estimates of Clarida, Gali and Gertler (2000). For the ECB (panel B), the degree of interest smoothing is once again high. However, unlike with the US, our estimates point to the ECB responding strongly to unemployment rates but not the inflation rate.

Including our measures of inter-regional dispersion does two things. First, for both the U.S. and the ECB, it yields estimates of the Taylor rule that are more consistent with central banks responding to both aggregate inflation and

[4] We use a Newey-West weighting matrix with a truncation at 6 .

unemployment rates. For the US, this means the response to unemployment becomes negative and highly statistically significant, while for the ECB the response to aggregate inflation becomes statistically significant. Second, each of the measures of UE dispersion enters the regression with a non-zero coefficient. For the US, the positive coefficient implies that *as the degree of heterogeneity in UE rates increases across regions, the central bank tends to raise interest rates*. For the ECB, the sign is reversed, indicating that increases in dispersion lead to interest rate decreases.[5]

This surprising set of results is remarkably robust and difficult to reconcile with standard models of monetary-policy decision making. In the rest of this section, we explore the robustness of the result before turning to potential explanations in subsequent sections.

2.3 Robustness

As a first step to investigating the robustness of these results, Figure 5 provides a scatterplot of the US (weighted) cross-sectional variance of state unemployment rates against the orthogonalized component of interest rates. A clear positive relationship exists regardless of the outliers that appear. (These outliers are almost exclusively from 1982.) Not surprisingly given this scatter plot, we have found the positive relationship between the degree of heterogeneity in states' UE rates and interest rates to be quite robust to sub-sample analysis.[6] Unfortunately, given the short time

[5] We also found similar results when looking at the dispersion of inflation rates, as well as for dispersion of real exchange rates for the ECB. However, the inflation results were much more sensitive for the US, reflecting the fact that regional price level data is only at a monthly frequency for the four Census Bureau divisions. The RER series for the ECB were very highly correlated with the UE dispersion measures, so we focus exclusively on the unemployment measures here.
[6] Sub-sample results available from authors upon request.

sample available since the inception of the ECB, no time-sample verification can be provided for the ECB.

A second type of robustness check is to consider an alternative estimation approach. Because GMM estimation may fare poorly in short samples, we estimated equations (2.2) and (2.3) by 2SLS.[7] The results were qualitatively unchanged. A third issue to be concerned about is whether our choice of interest rates is the correct one. For example, the U.S. Federal Reserve chooses a target for the Federal Funds rate (FFR), from which the effective FFR may differ, sometimes for reasons unrelated to monetary policy.[8] We reproduced the results of Table 2.1 using the target FFR for the Fed and the refinancing rate for the ECB and found nearly identical results.[9] We also found that our results are insensitive to using the GDP deflator or the non-farm business deflator to calculate inflation, rather than the CPI.

Another potential issue is the failure to correctly condition for the central bank's information set. To see this, suppose that the central bank's forecasts of future inflation and unemployment rates contain much more information than is embodied in our instruments. In this case, if dispersion measures are useful predictors of future aggregate measures, then they may show up as significant predictors in a Taylor rule simply because they are capturing elements of the central bank's information set that we are not controlling for.[10] We address this issue in three ways.

---

[7] See Hansen, Heaton, and Yaron (1996) and Christiano and den Haan (1996) for discussions of how GMM estimators fare in short time-samples. These 2SLS results are available from authors upon request..

[8] See Romer and Romer (2004).

[9] These results are also available upon request.

[10] See Orphanides (2001) for a discussion of the importance of controlling for the central bank's real time expectations.

The first approach is to determine whether our dispersion measures appear to contain useful information for predicting future aggregate variables. Figure 6 plots the dynamic cross-covariances of the weighted variance of regional unemployment rates each month against leads and lags of aggregate inflation and unemployment rates for the U.S. and the ECB over the time periods used for the empirical analysis.[11] For both the U.S and the Euro-Area., there is little evidence that regional dispersion in unemployment rates tends to lead aggregate measures, instead the opposite appears to be true. Both high inflation and high aggregate unemployment rates tend to be followed by increased dispersion of unemployment rates across states. Granger Causality tests, presented in Table 2.3, tell a similar story. For the US, we cannot reject the null that the cross-sectional variance of unemployment rates does not granger-cause aggregate inflation and unemployment. For the Euro-Area, we can reject the null for aggregate unemployment but not for inflation.

A second approach is to augment our Taylor rules with variables that are well-known leading indicators. Our test consists of augmenting equation (2.3) with leading indicators such that:

$$i_t = \phi_\pi \pi_{t+j} + \phi_{ue} ue_{t+j} + \sum_{i=1}^{T} \rho_i i_{t-i} + \beta_1 D_t + \beta_2 LI_t + \zeta_t \qquad (2.4)$$

where D_t is one of our measures of regional dispersion of unemployment rates while LI_t is the leading indicator added to the regression. The results are presented in Table 2.4, using the weighted variance of cross-sectional unemployment rates each month

[11] Use of our other measures of regional dispersion of unemployment rates yields nearly identical results.

as our measure of dispersion.[12] As forward-looking variables we use consumer confidence indicators, stock prices, WDI oil prices, PPI inflation, and industrial production.[13] We also consider the case of a time trend. For both the US and the ECB, none of variables eliminates the influence of the cross-sectional variance of unemployment rates. Again, there is little evidence that the measures of regional dispersion of unemployment rates are capturing forward-looking behavior or information that is inadequately modeled.

The third approach to control for the central bank's information set makes use of the real-time forecasts of future variables that the central bank relied on to make their decisions instead of using ex-post realized values of these variables. For the US Federal Reserve, Green Book forecasts are available for much of our time sample and provide expectations for the Fed, at the time of each meeting, of inflation, output growth, and unemployment in the future. Unfortunately, this restricts the time sample to January 1982 until December 2000. In addition, forecasts are available for each meeting only, so our time frequency is that of the meetings of the Board of Governors every six weeks. This leaves us with 152 observations, with which we estimate the following equation

$$i_t = \phi_\pi E_t\left[\sum_{j=0}^{2}\pi_{Q_j}\right]/3 + \phi_{gy} E_t\left[\sum_{j=0}^{2} gy_{Q_j}\right]/3 + \phi_{ue} E_t\left[\sum_{j=0}^{2} ue_{Q_j}\right]/3 + \rho i_{t-1} + \beta_1 D_{t-1} + \varepsilon_t \tag{2.5}$$

---

[12] The results are very similar using the other measures of regional dispersion of unemployment rates.

[13] For the US, these are specifically the University of Michigan Consumer Sentiment Index, the Dow Jones Industrial Average, the spot price of WTI oil, and the PPI-all commodities index. For the Euro-Zone, we use the Consumer Confidence Indicator of the European Commission Consumer Survey, the DAX German stock price index, and the Euro-Zone PPI all industries excluding construction index.

where i_t is the new Target FFR chosen at each meeting. The expectations terms are the average Green Book Forecasts of expected inflation, output growth, or unemployment over the current quarter and the subsequent two quarters. The lagged interest rate term is the Target FFR chosen at the previous meeting, and D_{t-2} is the measure of cross-sectional heterogeneity in unemployment rates two months prior to that in which the meeting occurs.[14] We include expectations of output growth because these appear to play an important role in affecting interest rate decisions. We use the two-month lag in the dispersion measure to ensure that these series are orthogonal to the error term, which captures monetary policy shocks, as well as to reflect the time lag involved in the release of unemployment rates. Because all of the RHS variables are determined prior to the decision about the new Target FFR, equation (2.5) can be estimated by OLS.

The results are presented in Table 2.5. The first column presents baseline results excluding unemployment dispersion measures. The results are quite consistent with our priors about how the central bank sets interest rates. The coefficients on inflation and output growth are positive and statistically significant, while that on unemployment is negative and also statistically different from zero. Adding our measures of dispersion has little effect on the coefficients on expectations of future aggregate inflation and output growth. However, as in Table 2.1, the estimated response to expected unemployment becomes larger (in absolute value), more than doubling when the weighted variance of state unemployment rates is added to the regression. The coefficients on measures of unemployment dispersion are still

[14] We include only one lag of the target FFR because higher order lags are all insignificant and have no effect on other coefficients.

positive, as they were in the estimates presented in Table 2.1. However, when using the percentile differences of the distribution of unemployment rates across states, the coefficient is only weakly (at the 10% level) significantly different from zero. The variance measure, however, remains positive and statistically different from zero at the 1% level, indicating that even after controlling for the central bank's expectations (and also losing many observations), the variance of state unemployment rates continues to be an important predictor of the Federal Reserve's interest rate policy.[15]

3 Possible Theoretical Explanations and Tests

Why would one be surprised to find that central bankers respond to dispersion of unemployment rates across regions? Consider the following simple representation of the central banker's problem, based on Barro and Gordon (1983): the central banker has a loss function over aggregate inflation and aggregate unemployment

$$L = \frac{1}{2}(\pi - \pi^*)^2 + \frac{\lambda}{2}(u_a - \overline{u}_a)^2 \tag{3.1}$$

where π is the inflation rate, common to all regions, and $u_a - \overline{u}_a$ is the deviation of aggregate unemployment u_a from the natural rate of unemployment $\overline{u}_a$. Aggregate unemployment is defined as a weighted average over regional unemployment rates u_i such that $u_a = \sum_{i=1}^{N} \omega_i u_i$, and the aggregate natural rate of unemployment is defined in an equivalent manner. The weights ω_i are the population share of each region.

[15] Because the ECB does not release forecasts used for each meeting, we cannot replicate this analysis for the ECB.

Suppose that each region is subject to an expectations-augmented Phillips Curve tradeoff between inflation and unemployment

$$u_i - \bar{u}_i = -\alpha\left(\pi - \pi^e\right) + \varepsilon_i \tag{3.2}$$

where ε_i is a region-specific shock. Assuming the central bank can choose the inflation rate, the optimal policy is given by

$$\pi^{opt} = \frac{\alpha\lambda}{1+\alpha^2\lambda}(\varepsilon_a + \alpha\pi^e) \tag{3.3}$$

This optimal policy accommodates both shocks to inflation expectations and the (weighted) average shock to the Phillips Curve. Clearly, optimal policy is independent of all regional considerations once one conditions on aggregate measures.

This feature of the model does not just reflect the fact that the loss function is over aggregate variables. Suppose the central banker seeks to minimize a weighted sum of regional loss functions

$$L = \sum_{i=1}^{N} \omega_i L_i = \sum_{i=1}^{N} \omega_i \left(\frac{1}{2}\pi^2 - \frac{\lambda}{2}\left(u_i - \bar{u}_i\right)^2 \right) \tag{3.4}$$

Such a loss function implies that certain regions may have a disproportionate effect on the aggregate loss function if they are experiencing large shocks. Yet the optimal policy in this context is identical to (3.3). That is, it is independent of regional considerations once aggregate variables have been taken into account. We will refer to this feature of the model as *aggregation equivalence*, though it must be emphasized that this is observationally equivalent to the well-established certainty equivalence principle.

In the context of interregional heterogeneity, three assumptions lead to this aggregation equivalence: the quadratic nature of the loss function, the linearity of the

Phillips Curve, and the equality of the weights in the loss function and population weights used to generate aggregate variables. In the subsequent sections, we relax each of these assumptions in turn and consider the implications for optimal monetary policy. The key question we are interested in is whether relaxing these measures can provide an explanation for our empirical findings.

3.1 Central Bank Preference for Homogeneity

We first consider breaking the quadratic nature of the loss function. Woodford (2003) demonstrates how one can derive a loss function similar to those considered here by taking a second-order approximation to the utility function of the representative consumer.[16] Taking a second-order approximation to the utility function imposes certainty equivalence on the consumer and the central bank, since it is the third derivative of the utility function that breaks certainty equivalence. In general, higher-order approximations to the utility function would break both certainty and aggregation equivalence on the part of consumers and the central bank.

Rather than derive a higher-order approximation to the utility function, we follow a more tractable approach and consider the policy implications of a central bank that dislikes dispersion in regional welfare, in addition to wanting to maximize aggregate welfare. The model that follows is appealing in that it nests the standard models of monetary policy-making that ignore regional differences in welfare. The goal is to identify how the central bank should respond to regional dispersion in unemployment rates if it did care about minimizing such dispersion.

---

[16] A key difference is that Woodford's loss function includes the deviation of output from the flexible-price level of output rather than the unemployment rate.

Suppose that the central banker's loss function is now given by:

$$L_a = \sum_i \omega_i L_i + \frac{\kappa}{2}\sum_i \omega_i \left[L_i - \sum_i \omega_i L_i \right]^2 \qquad (3.5)$$

The term κ represents the weight the central bank places on its dislike of heterogeneity. When κ=0, the model collapses back to the standard case, with optimal policy given by (3.3). When κ>0, the central banker wants to maximize aggregate welfare without imposing disproportionate losses on any single region. We can show that interregional differences matter in this context – in particular, the second and third moments are relevant for policy. Specifically:

Proposition 1: The optimal choice of inflation for the central bank (π^*) that minimizes (3.5) subject to (3.2) for all regions is of the form

$$\pi^* = \pi^{opt} f\left(\sigma_u^2\right) + skew\left(u_i - \bar{u}_i\right) g\left(\sigma_u^2\right)$$

where $\sigma_u^2 \equiv \sum_i \omega_i \left[\left(u_i - \bar{u}_i\right) - \left(u_a - \bar{u}_a\right)\right]^2$ is the weighted variance of cross-sectional unemployment rates, $u_a \equiv \sum_i \omega_i u_i$, $\bar{u}_a \equiv \sum_i \omega_i \bar{u}_i$, and f and g are continuous functions of this variance.

Proof: See Appendix 1.

This proposition shows that the optimal policy depends on the first three moments of the distribution of shocks across regions. Because the second and third moments of the distribution of shocks are the same as the second and third moments of the distribution of unemployment rates (around their natural levels), the optimal policy augments that of (3.3) with functions only of the variance and skew of regional unemployment rates. In addition, we can show:

Corollary 1: a) $f\left(\sigma_u^2\right) \geq 1$, with equality when κ=0 or $\sigma_u^2 = 0$.

b) $\dfrac{df\left(\sigma_u^2\right)}{d\sigma_u^2} > 0$.

c) $g\left(\sigma_u^2\right) \leq 0$, with equality when κ=0 or $\sigma_u^2 = 0$.

d) $\dfrac{dg\left(\sigma_u^2\right)}{d\sigma_u^2} < 0$.

Proof: See Appendix 1.

The first result in the corollary establishes that if one objective of the central bank is to minimize the differences in aggregate welfare across states or regions, it must respond *more strongly* to the determinants of optimal policy (aggregate shock and expectations of inflation) when the variance of unemployment rates is nonzero. Suppose all states experience a common shock that tends to increase aggregate unemployment with no change in the variance of regional unemployment rates. In the case with $\kappa=0$, the central bank chooses to raise inflation to offset some of the increase in aggregate unemployment, because the loss function is quadratic in both inflation and unemployment. If the cross-sectional variance of unemployment rates is positive, then some states have higher unemployment rates than the average state. These states suffer disproportionately from the increase in unemployment from the shock, again because of the quadratic nature of the loss function. Thus, if $\kappa>0$ so that the central bank has the additional goal of avoiding imposing disproportionate welfare losses on an single region, the central bank must raise inflation more than it would otherwise to accommodate the disproportionate loss suffered by high unemployment states. The second result of the corollary establishes that as the variance of unemployment rates rises, this phenomenon becomes increasingly important as larger fractions of states have disproportionately large welfare losses.

The third result indicates that the coefficient on the skew of the distribution of unemployment rates must be positive (when $\kappa=0$ and the variance of unemployment rates is nonzero). The skew captures the asymmetry of the distribution. When it is positive, there is a fat tail of high unemployment rates, whereas when it is negative there is a fat tail of low unemployment states. States in the fat tails tend to suffer disproportionate welfare losses and therefore have to be accommodated by the central bank when $\kappa>0$. The fourth element indicates that, holding the skew constant, an increase in the variance diminishes the response of optimal policy to the skew of the unemployment distribution.

These results are similar to the theoretical arguments laid out by Dixit (2000) and Fuchs and Lippi (2006). They each consider the problem of an aggregate central bank trying to maximize aggregate welfare subject to the constraint that the regional members find it optimal to stay in the monetary union. They find that the central bank should respond disproportionately to regions for which the participation constraint is binding. Thus, aggregate policy is affected by regional concerns above and beyond those embodied in aggregate variables. The approach considered in this paper naturally yields a similar conclusion but can be applied to both the ECB and the Fed, whereas the notion of states exiting the monetary union is clearly inapplicable to the US.

The baseline model is a straightforward approach to evaluating whether monetary policy seems to be influenced by interregional heterogeneity. The model under discussion here – in which the Fed has a preference for less heterogeneity of welfare – has additional features that are not included in the baseline case. Particularly, this model indicates that the skew of the distribution is important, and that the actual policy will be linked to the "optimal" policy absent of any preference for lower dispersion. Unfortunately, the optimal policy is highly nonlinear in the variance. To better map this into our empirical approach, we take a first-order approximation to the optimal policy π^*:

$$\pi^* \approx c + \phi_{opt}\pi^{opt} + \phi_{\mathrm{var}(ue)}\,\mathrm{var}(ue) + \phi_{skew(ue)}skew(ue) \qquad (3.6)$$

where $\phi_{opt} > 1$, $\phi_{\mathrm{var}(ue)} = f'\left(\overline{\sigma_u^2}\right)\overline{\pi^{opt}} + g'\left(\overline{\sigma_u^2}\right)\overline{skew}$, and $\phi_{skew(ue)} = g\left(\overline{\sigma_u^2}\right) > 0$. This expression implies that the central bank should allow for higher inflation when the skew of unemployment rates is high. The response to higher variance is more ambiguous since *g'* is negative. If the average skew of the distribution in negative, then higher variance of unemployment rates should be accommodated with higher inflation. If the average skew is positive, the sign of the response is ambiguous, but generally non-zero.

We thus augment our baseline model to incorporate these features. We treat the following as the desired interest rate rule when $\kappa=0$:

$$i_t^{des} = \phi_\pi \pi_{t+6} + \phi_{ue} ue_{t+6} + v_{t+6} \qquad (3.7)$$

and use the following estimating equation:

$$i_t = c + i_t^{des} + \sum_j \rho_j i_{t-j} + \beta_{\mathrm{var}} \operatorname{var}(ue_{i,t}) + \beta_{skew} skew(ue_{i,t}) + \varepsilon_t \quad (3.8)$$

to estimate c, ϕ_π, ϕ_{ue}, ρ_j, β_{var}, and β_{skew}, jointly. This specification is appealing because it naturally nests the baseline scenario and allows us to test the specific hypothesis predicted by the theory.

The null hypothesis (κ=0) is that $\beta_{var}=\beta_{skew}=0$. For the US, the mean skew of regional unemployment rates is positive, so the theory predicts β_{var} has an ambiguous, but generally nonzero coefficient. For the Euro-Area, the mean skew is negative, so the theory implies higher variance should lead to lower interest rates. For both the Fed and the ECB, the theory predicts that if κ>0, then $\beta_{skew}<0$.

The results are presented in Table 2.6. For the US, the results are consistent with the theory. Once we estimate equation (3.8) with both the variance and the skew, the coefficient on the variance is non-zero while that on the skew is negative. For the ECB, the results are more mixed. The coefficient on the variance is negative, as implied by the theory. However, the estimated coefficient on the skew of unemployment rates is positive, which directly contradicts the prediction of the model. Thus, one might conclude that the data are consistent with the predictions of the model for the US, but less so for the Euro-Area.

3.2 The Phillips Fan

A second necessary condition for aggregation equivalence to hold is a linear Phillips Curve (PC). One shortcoming of a linear Phillips curve is that the tradeoff between inflation and unemployment is unchanging. In this section, we consider an alternative

Phillips Curve representation such that when there is a shock to the economy, the inflation-unemployment tradeoff actually changes. A bad (positive shock to the PC) shock, for instance, will tend to increase unemployment. However, with increased slack in the economy, it may be the case that looser monetary policy can be more effective than before. That is, it becomes less costly in terms of inflation to lower unemployment in the face of a bad shock. Conversely, if there is a good shock (negative shock to the PC) to the economy and markets are tight, higher inflation does very little to reduce unemployment. In fact, a policy of lowering inflation will now increase unemployment by less than before, and may be desirable.

The simplest way to model this idea is to assume that there is a linear relationship between the shock, ε, and the slope of the Phillips curve, α. In particular,

$$\alpha = (C + \gamma\varepsilon) \tag{3.9}$$

where C is a constant, and $\gamma \geq 0$ is the degree to which the slope changes in response to a shock. To prevent the Phillips curve from sloping upwards, we further assume $(C+\gamma\varepsilon) \geq 0$. Note that the standard case with a constant slope returns if $\gamma=0$.

With this new version of α, the Phillips curve is now defined as

$$u - \overline{u} = -(C + \gamma\varepsilon)\left(\pi - \pi^{e}\right) + \varepsilon \tag{3.10}$$

A bad shock (that is, where ε is positive) simultaneously raises unemployment and makes it cheaper (in terms of inflation) to reduce unemployment. While the Phillips curve continues to be linear in inflation, the slope changes with different realizations of the shock, as depicted in Figure 2.A.

Figure A: The Phillips Fan

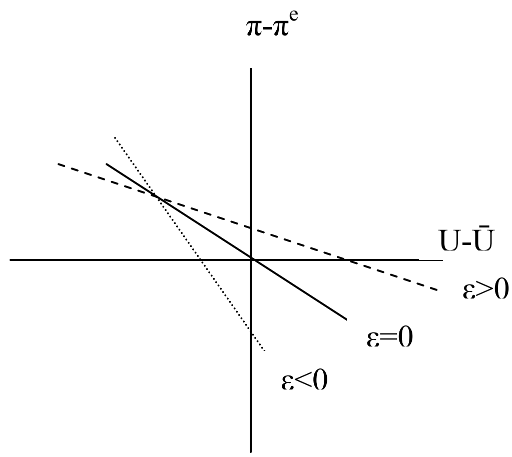

In this environment, we refer to the set of all possible Phillips curves as the "Phillips Fan." One can easily show that Phillips Fan pivots through a single point in the fourth quadrant of the graph. This is at the point (-C/γ, 1/γ). Note that at a realization to the left of this point, a bad shock actually makes feasible outcomes of both lower unemployment and lower inflation. We rule out this "free lunch" possibility by assuming $u-\overline{u} \geq$-C/γ, and π-$\pi^e \leq$1/γ. Equivalently, and more practically (as we will see), (1- $\gamma(\pi-\pi^e)) \geq 0$.

Now consider the same loss function for the central bank as (3.4). We show in Appendix 2 that:

$$\frac{d\pi}{d\sigma_\varepsilon^2} = \frac{\lambda\gamma(1-\gamma(\pi-\pi^e))}{1+\lambda\gamma^2\sigma_\varepsilon^2+\lambda(C+\gamma\varepsilon_a)^2} \tag{3.11}$$

This term is positive, given the no free lunch assumption of above (that is, $(1-\gamma(\pi-\pi^e))\geq 0$). This indicates that as the variance of the shocks increases, so does the optimal level of inflation. The regressions of Table 2.1 suggest that the Phillips Fan may not be a good representation of the United States, as the results imply that the

U.S. tries to lower inflation (raise interest rates) when there is a variance-increasing shock. For Europe, however, the results are consistent with the Phillips Fan story.

One can further find the direct relationship between the variance of unemployment and optimal inflation. We show in Appendix 2 that:

$$\frac{d\pi}{d\sigma_u^2} = \frac{\lambda\gamma}{(1-\gamma(\pi-\pi^e))(1-\lambda\gamma\sigma_\varepsilon^2+\lambda(C+\gamma\varepsilon_a)^2)} \tag{3.12}$$

The sign of (3.12) is determined by the sign of the term $(1-\lambda\gamma\sigma_\varepsilon^2+\lambda(C+\gamma\varepsilon_a)^2)$. Since by assumption $\lambda \geq 0$, this term will be negative only if $\frac{1}{\lambda} < \gamma\sigma_\varepsilon^2 - (C+\gamma\varepsilon_a)^2$. Thus, while the central bank unambiguously responds to increased variance of the shock by raising inflation, the net result of the shock and the change in policy leads to an ambiguous relationship between inflation and the variance of unemployment. This suggests that central banks with loss functions that place a relatively high weight on unemployment (large λ) will tend to demonstrate a pattern of lower inflation when the interregional variance of unemployment increases. Also, economies that display a relatively high level of interregional heterogeneity (large σ_ε^2) will also tend to show a negative correlation.

In these cases, as the variance of the shock rises – that is, as the Phillips Fan spreads open – the monetary authority seeks higher inflation. This drives down overall unemployment, and also drives down the variance of unemployment as we travel up the spines of the fan. We can see this latter result when we think of the variance of unemployment as the distance between the spines of the fan. Figure 2.B shows a hypothetical case where this distance has decreased following an increase in the variance of the shocks and the ensuing monetary policy adjustment. Neither the

US nor EU shows that tendency, as both show a positive correlation between inflation and the variance of unemployment (roughly 0.34). That said, this is an interesting theoretical possibility, that may be relevant for monetary regimes that place a large value on deviations of output from the natural level, or if nations with a great degree of heterogeneity attempt to enter into a monetary union.

Figure 2.B: Hypothetical Response to a Variance-Increasing Shock.

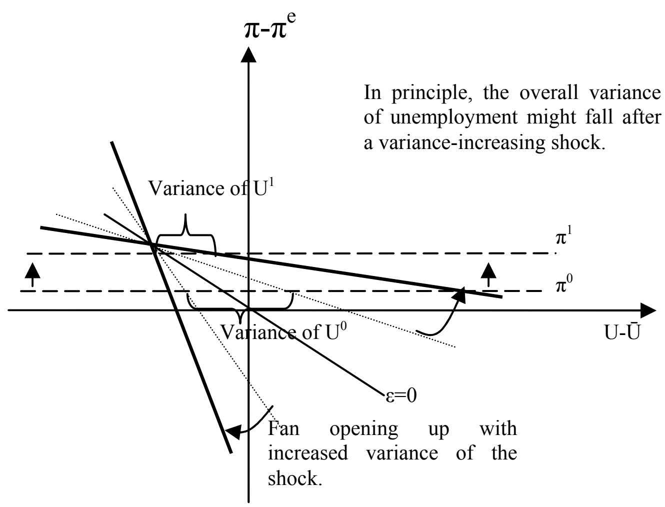

Another test for the relevance of the Phillips Fan as a description of policy involves how the central bank responds to shocks in states with high unemployment vs. low unemployment (relative to their natural levels). We show in Appendix 2 that:

$$\frac{d\pi}{d\varepsilon_i} = \frac{\lambda[2\gamma(u_i - \overline{u}_i) + C]}{1 + \lambda \sum_i \omega_i (C + \gamma\varepsilon_i)^2} \tag{3.20}$$

Since the denominator is the same for all regions, the impact on policy of a shock to region *i* is an increasing function of the difference between unemployment and the natural level. That is, regions where unemployment is high compared to the natural level will see a stronger response if that region experiences a shock.[17] This is due to the increasing effectiveness of monetary policy after a bad shock hits: after a region experiences a bad shock, the central bank gets more bang from a monetary expansion than it could before. With the tradeoff between unemployment and inflation shrinking, the new optimal level is with higher inflation than before.

To test this, one can break our percentile-based dispersion terms into two components: UEP9010=UEP90-UEP10=(UEP90-UE)+(UE-UEP10). The first component is the deviation of high unemployment states (here the 90th percentile) from the mean, while the second is the deviation of low unemployment states from the mean. Applying this decomposition to our empirical specification yields:

$$i_t = c + i_t^{des} + \sum_j \rho_j i_{t-j} + \beta_1 (ue_{90,t} - ue_t) + \beta_2 (ue_t - ue_{10,t}) + \varepsilon_t \qquad (3.21)$$

where the null of $\beta_1=-\beta_2$ is that the central bank responds in the same way to higher unemployment on the part of high and low unemployment states. Note that by (3.20), the Phillips Fan implies that $\beta_1<-\beta_2$. That is, the central bank responds more aggressively to shocks to high unemployment than low unemployment regions.

Table 2.7 presents empirical results from applying this decomposition of dispersion to both the 90-10th and 75-25th percentile differences for the US and ECB. For the US, the first thing to note is that β_1 is not significantly different from zero. However, the coefficient on low-unemployment states, β_2, is positive and highly statistically significant. This says that when low-unemployment states get even lower

[17] Scrutiny of (3.18) reveals that if a region has unemployment sufficiently below the natural level, then a bad shock to that state actually lowers the optimal inflation level. This unusual result is due to the nonlinear properties of the Phillips fan.

unemployment rates, the central bank raises interest rates more than what is implied by the change in aggregate unemployment rate. *In other words, the US Federal Reserve appears to respond disproportionately to shocks to low-unemployment states*, which directly contradicts the predictions of the Phillips Fan. The Wald tests confirm that we can strongly reject the null of equal responses to high and low unemployment states, but the direction of the failure is opposite that implied by the model. As mentioned above, since the US pattern also seems to deviate from the policy suggested by equation (3.11), the Phillips Fan may not be appropriate description for US monetary policy.

The story may be different for Europe, however. For the ECB, the coefficient on high-unemployment states is negative and highly statistically significant, implying that when high-unemployment states see their unemployment rates rise relative to the mean, the central bank lowers interest rates disproportionately. For low-unemployment states, the evidence is more mixed. Using the 90-10th percentile decomposition, the ECB's response to changes in a given country's unemployment rates is the same regardless of whether that country unemployment rate is low or high (relative to the natural level). However, when using the 75th-25th decomposition, the response to changes in unemployment in the high-unemployment country is indeed stronger than for the low. Thus, the results for the ECB are consistent with the theory in one case, 75th-25th, but not in the other, 90th-10th.

3.3 Regional Voting Shares and Interest Rate Decisions

As a final attempt to account for the apparent importance of unemployment dispersion in monetary policy, we examine the representation of different regions in the voting decisions of each central bank. That the Fed and the ECB tend to respond differently to regional unemployment gaps raises the possibility that the central banks may, by institutional construction, weigh regions differently, and this is showing up through our dispersion terms.

To see how this could be, suppose that the loss function (3.4) is replaced with the following loss function:

$$L = \sum_{i=1}^{N} \Omega_i L_i = \sum_{i=1}^{N} \Omega_i \left(\frac{1}{2}\pi^2 - \frac{\lambda}{2}\left(u_i - \overline{u_i}\right)^2 \right) \tag{3.22}$$

where Ω_i may be different than the population weight ω_i for some i. In this case, one can readily show that the optimal policy, given regional PCs as in (3.2) is given by

$$\pi^* = \pi^{opt} + \frac{\lambda\alpha}{1+\lambda\alpha^2} \sum_{i=1}^{N} \Omega_i \left[\left(u_i - \overline{u_i}\right) - \left(u_a - \overline{u_a}\right) \right]. \tag{3.23}$$

The optimal policy π^* is equal to the typical optimal policy π^{opt} plus an extra term that reflects the deviation of each region's unemployment rate from the aggregate level, weighted by that region's influence in the aggregate loss function. A high unemployment state with disproportionately high voting power will induce the central bank to tolerate higher inflation than would be optimal if voting weights reflected population weights.[18]

[18] Aksoy et al (2002) similarly argue that if some voting members focus on regional concerns, aggregate interest rate decisions can be sub-optimal because of majority voting.

Why would one believe that the central bank could weigh regions differently than implied by their population weights? Meade and Sheets (2005) providence that voting members of the Federal Open Market Committee (FOMC) respond disproportionately to their region of origin in their voting decisions.[19] Meade and Sheets conclude that this is unlikely to affect interest rate decisions since regional differences would cancel out across voting members. However, if voting weights are not equal to population weights, then regional biases may not cancel out, leading to policy function like (3.23).

In the ECB, interest rate decisions are made by the Governing Council which consists of the twelve governors of national central banks and six members of the Executive Board, of which four are typically from the "big" countries: Germany, France, Italy, and Spain. A majority vote decides interest rate policy, but no minutes or records of voting patterns are released. In contrast, the Federal Reserve sets interest rates by a majority vote among seven Board members and five (rotating) regional Bank Presidents (Dominguez, 2006).

We first consider whether the composition of voting members of the FOMC and Governing Council of the ECB are representative of the population shares accounted for by each region. For the US, we use data from Meade and Sheets (2005), which provides the voting decisions of each member of the FOMC for each meeting since 1982 out to 2000. Each voting member is assigned a region of origin, including Board members. Table 2.8 presents the fraction of votes accounted for by

---

[19] Heinemann and Huefner (2004) provide similar evidence for the ECB.

members of each Federal Reserve District. For the ECB, we present a similar breakdown by country from January 1999 to September 2005.

For the US, the Northeast, i.e. New England and New York, appear to be the most heavily over-represented regions in FOMC meetings. Both districts 1 and 2 of the Federal Reserve have been the source of a disproportionate amount of voting Board members relative to their share of the population. In addition, because the New York Fed always has a vote at FOMC meetings, it also accounts for a disproportionate share of voting done by regional presidents. On the other hand, the Southeast (Atlanta-based) and the West (San Francisco-based) are the most under-represented in voting decisions relative to their share of the population. The Southeast is particularly unaccounted for in terms of voting Board members. All other districts have accounted for a share of votes approximately equal to their share of the population. Note that there appears to be no close relationship between voting representation in FOMC meetings and the average difference between regional unemployment and the aggregate unemployment rate.[20]

For the ECB, France, Germany, and Italy are heavily underrepresented in the voting decisions of the Governing Council of the ECB, despite each of them consistently occupying a seat on the Executive Board, in addition to their representation via their national central banks. Most dramatically, while Germany accounts for over 25 percent of the Euro-Zone's population, it only accounts for ten percent of votes cast. Instead, the smaller countries, which each receive at least one

[20] Because some districts account for parts of states, we had to arbitrarily place some states entirely within some districts and therefore exclude them from districts of which they are partially part of. This was done only for the purposes of calculating unemployment rates per district. See Appendix 2 for a complete description of which districts includes which states.

vote are over-represented relative to their share of the population. Luxembourg, for example, accounts for little over one-tenth of a percent of the population but has cast about six percent of all votes in meetings of the Governing Council. Interestingly, the over-represented "small" countries of the ECB have had much lower unemployment rates than those of France, Germany, and Italy (Finland is the only exception).

To determine whether our finding that dispersion measures of regional unemployment rates affect interest rates is due to the over- and under-representation of regions in central banks' decision-making procedures, we construct a measure of weighted regional unemployment gaps, where the weights are given by the voting representation of each region, as implied by the optimal policy function (3.23). For the US, we take the fraction of votes associated with each region on any given meeting of the FOMC and multiply these fractions by the difference between the (one-month lagged) unemployment rates of each region from the (one-month lagged) aggregate unemployment rate. This yields a series with frequency given by Fed meetings. For the ECB, we apply the same procedure using contemporaneous values of unemployment rates. Because the ECB Governing Council meets monthly, this is a monthly series.

We then use the following equation to test whether including these measures eliminates the predictive power of the dispersion measures used in the previous sections

$$i_t = \phi_\pi E_t \pi_{t+j} + \phi_{gy} E_t gy_{t+j} + \phi_{ue} E_t ue_{t+j} + \sum_{i=1}^{T} \rho_i i_{t-i} + \beta_1 D_t + \beta_2 D_t^{voting} + \varepsilon_t \tag{3.24}$$

where D_t one of our measures of the cross-sectional dispersion of regional unemployment rates and D_t^{voting} is our new measure of dispersion using voting shares of each region. For the US, we estimate equation (3.24) using Green Book forecasts of inflation, output growth, and aggregate unemployment by OLS from 1982:01 to 2000:12 at the frequency of Fed meetings. Both the dispersion measures are lagged one period to ensure orthogonality of RHS variables to error term. For the ECB, we use ex-post values of inflation and aggregate unemployment and estimate (3.24) by GMM, as in Section 2, at a monthly frequency.[21]

The results are presented in Table 2.9. The main feature of the results is that accounting for voting weights has almost no effect on the original results. For the US, including voting weights does not alter the significance of our dispersion measures. In addition, the voting dispersion terms appear to have no predictive power for interest rate decisions. For the ECB, only in the case of the difference between the 90th and 10th percentiles of the unemployment distribution are our results affected. In particular, the coefficient on this dispersion term becomes insignificantly different from zero. For the other two dispersion measures, our baseline results continue to hold. The voting term is only significantly different from zero (at the 10% level) when using the 75th-25th dispersion measure. In addition, the positive coefficient on this term is counterintuitive. Overall, it appears that our measures of dispersion are not capturing institutional biases for and against various regions. While regional concerns may have an effect on the individual decisions of voting members, as argued by Meade and Sheets (2005), these appear to have no aggregate effect on interest

[21] We have to use different approaches because Fed meetings are not held monthly and we do not have ECB forecasts of future aggregate variables. ECB estimates are insensitive to including industrial production.

rates and cannot explain why interest rates appear to systematically respond to the regional dispersion of unemployment rates each period.

4 Conclusion

This paper presents robust evidence that interest rate decisions of policy-makers are systematically affected by the distribution of unemployment rates across regions. We find that this phenomenon appears to hold both for the US Federal Reserve and the European Central Bank. While the signs on the coefficient for unemployment dispersion is different for each bank, an appealing feature of augmenting estimates of the Taylor rule with these dispersion measures is that they help more clearly identify the response of the central bank to aggregate inflation and unemployment, and make the Fed and ECB aggregate monetary policies appear to be more qualitatively consistent with one another.

We show that this result does not appear to be a statistical anomaly, particularly for the US. Even after controlling for the Fed's expectations via Green Book forecasts, or attempting to adjust for imbalanced regional representation on the boards, we continue to find a strong predictive role for the cross-state variance of unemployment rates.

This result is surprising because according to traditional theory and the legal foundations of these institutions, regional concerns should not affect interest rates decisions other than by their effect on macroeconomic aggregate variables. Due to "aggregation equivalence," even attempts that initially separate differences in unemployment across regions tend to collapse to the standard result in which only aggregates matter for policy. We have made three attempts to break this aggregation equivalence: based on the central bank disliking dispersion, nonlinear Phillips Curves, and accounting for regional representation in voting decisions.

The results for these theoretical explanations are mixed. The results for the US are consistent with the central bank placing a penalty of regional differences in welfare, but this is not the case for the ECB. On the other hand, non-linear Phillips curves, as approximated by our Phillips Fan approach, can partially account for the ECB lowering interest rates when the variance of the shocks across regions increases. Also, for the ECB, the differential response of the shocks to high vs. low unemployment states is consistent with the conjecture of the Phillips Fan.. Finally, we find no strong evidence that regional misrepresentation in voting decisions could explain our empirical results for either the Fed or the ECB.

One conclusion that is consistent with these results is that the monetary policymakers in the US and ECB have a fundamental difference in approach. The Fed may be more concerned with reconciling welfare differences across regions, while the ECB may cognizant of the differential impact of regional shocks on the overall effectiveness of monetary policy. Further analysis, particularly as the Euro-zone matures, could shed more light on the relevance of heterogeneity for policy.

The remarkable robustness of our results for the US is striking and not readily reconciled with simple deviations from a static model of optimal monetary policy. One tantalizing piece of evidence is the fact that the statistical significance of unemployment dispersion terms for the US seems to come from a disproportionate sensitivity of the central bank to low-unemployment states. This could arise in a dynamic "bottleneck" model of inflation. If low unemployment rates tend to lead to production bottlenecks that place significant pressures on inflation, a central bank should raise interest rates disproportionately when low-unemployment states see their unemployment rates fall further. Such a model would deviate from linearized New Keynesian models of inflation dynamics and would require taking seriously nonlinear features of the economy.

Table 2.1: Does Regional Variation in UE Affect Interest Rates?

	Panel A: United States					Panel B: Euro-Zone			
	Baseline	***(1)***	***(2)***	***(3)***		***Baseline***	***(1)***	***(2)***	***(3)***
c	0.00	0.33***	0.03	-0.02	c	1.17***	3.00***	1.96***	2.79***
	(0.09)	(0.09)	(0.07)	(0.08)		(0.25)	(0.26)	(0.35)	(0.26)
ϕ_{π}	0.08***	0.08***	0.08***	0.10***	ϕ_{π}	0.03	0.15***	0.21***	0.15***
	(0.03)	(0.02)	(0.02)	(0.02)		(0.03)	(0.03)	(0.06)	(0.04)
ϕ_{ue}	-0.02*	-0.10***	-0.07***	-0.05***	ϕ_{ue}	-0.13***	-0.30***	-0.16***	-0.27***
	(0.01)	(0.02)	(0.01)	(0.01)		(0.03)	(0.02)	(0.02)	(0.02)
ρ_1	1.32***	1.25***	1.26***	1.27***	ρ_1	0.95***	0.83***	0.88***	0.84***
	(0.06)	(0.05)	(0.05)	(0.05)		(0.01)	(0.02)	(0.03)	(0.02)
ρ_2	-0.34***	-0.30***	-0.30***	-0.31***	ρ_2				
	(0.06)	(0.04)	(0.05)	(0.05)					
Dispersion Measure:		*var(UE)*	*UEP9010*	*UEP7525*	*Dispersion Measure:*		*var(UE)*	*UEP9010*	*UEP7525*
β		0.12***	0.10***	0.13***	β		-0.05***	-0.08***	-0.10***
		(0.02)	(0.02)	(0.02)			(0.01)	(0.02)	(0.01)
Sample	January 1982 - September 2005				*Sample*	January 1999 - October 2006			

Note: All estimates done by GMM with Newey-West weighting matrix, standard errors in parentheses. Instruments include 6 lags of each variable (inflation, unemployment, interest rates, and additional variables when included). Dependent variable is the interest rate, while ϕ_{π}, ϕ_{ue}, and ρ_i are coefficients on 6-month ahead inflation, 6-month ahead unemployment, and i lags of the interest rate respectively. Statistical significance at the 1%, 5%, and 10% levels are indicated by a ***, **, and * respectively.

Table 2.2: Decomposing the Variance of the Endogenous Component of Interest Rates

Panel A: United States

	Fraction of Var(Endog) due to		
	var(agg)	***var(D)***	***cov(agg,D)***
Dispersion Measure			
UEP90-UEP10	87%	113%	-100%
UEP75-UEP25	60%	69%	-29%
var(UE)	114%	167%	-181%

Panel B: Euro-Zone

	Fraction of Var(Endog) due to		
	var(agg)	***var(D)***	***cov(agg,D)***
Dispersion Measure			
UEP90-UEP10	81%	38%	-19%
UEP75-UEP25	79%	44%	-22%
var(UE)	96%	51%	-46%

Note: The table presents decompositions of the variance of the endogenous component of interest rates as defined in equation (5), but normalized by the variance of the endogenous interest rate component. Var(*agg*), var(*D*), and cov(*agg,D*) are the variance of the endogenous component of interest rates due to aggregate inflation and unemployment, the variance of the measure of regional unemployment dispersion, and the covariance of the two respectively. Each expression is multiplied by relevant constants from equation (5). The estimated coefficients used come from the results of Table 1.

Table 2.3: Granger Causality Tests

	F-Statistic	*p-value*
United States		
π does not Granger-Cause var(UE)	1.68*	0.07
var(UE) does not Granger-Cause π	1.47	0.13
UE does not Granger-Cause var(UE)	1.48	0.13
var(UE) does not Granger-Cause UE	0.92	0.53
Euro-Zone		
π does not Granger-Cause var(UE)	0.47	0.92
var(UE) does not Granger-Cause π	0.62	0.82
UE does not Granger-Cause var(UE)	1.58	0.12
var(UE) does not Granger-Cause UE	2.26**	0.02

Note: Granger Causality tests done with 12 lags. The var(UE) series is the weighted variance of regional unemployment rates each month. Data is from 1982:01 to 2005:09 for US and 1999:01 to 2006:10 for Euro-Zone. Statistical significance at the 1%, 5%, and 10% levels are indicated by a ***, **, and * respectively.

Table 2.4: Robustness to Including Other Variables

	Panel A: United States					
Added Variable:	***Consumer Confidence***	***Stock Prices***	***Oil Prices***	***PPI Inflation***	***Industrial Production***	***time trend***
c	-2.98***	0.52	0.34***	0.29***	0.14*	0.44**
	(1.15)	(0.56)	(0.11)	(0.11)	(0.08)	(0.20)
ϕ_π	0.10***	0.08***	0.08***	0.13***	0.07***	0.08***
	(0.02)	(0.02)	(0.02)	(0.03)	(0.02)	(0.02)
ϕ_{ue}	-0.03	-0.10***	-0.09***	-0.10***	-0.06***	-0.10***
	(0.03)	(0.02)	(0.01)	(0.02)	(0.02)	(0.02)
ρ_1	1.25***	1.24***	1.27***	1.26***	1.23***	1.26***
	(0.04)	(0.04)	(0.05)	(0.05)	(0.04)	(0.05)
ρ_2	-0.29***	-0.29***	-0.31***	-0.32***	-0.26***	-0.30***
	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)	(0.05)
β_1	0.09***	0.12***	0.11***	0.13***	0.09***	0.11***
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
β_2	0.65***	-0.02	-0.01	-0.01**	0.01***	-0.0004
	(0.22)	(0.05)	(0.03)	(0.005)	(0.004)	(0.0004)
Sample	January 1982 - September 2005					
	Panel B: Euro-Zone					
Added Variable:	***Consumer Confidence***	***Stock Prices***	***Oil Prices***	***PPI Inflation***	***Industrial Production***	***time trend***
c	1.72***	1.27**	1.85***	2.49***	1.73***	2.88***
	(0.22)	(0.60)	(0.27)	(0.27)	(0.30)	(0.26)
ϕ_π	0.19***	0.17***	0.03	0.10***	0.10***	0.13***
	(0.03)	(0.02)	(0.03)	(0.03)	(0.02)	(0.02)
ϕ_{ue}	-0.16***	-0.23***	-0.24***	-0.23***	-0.17***	-0.29***
	(0.02)	(0.03)	(0.02)	(0.03)	(0.03)	(0.02)
ρ_1	0.83***	0.84***	0.87***	0.84***	0.88***	0.84***
	(0.01)	(0.01)	(0.02)	(0.01)	(0.02)	(0.02)
β_1	-0.03***	-0.04***	-0.08***	-0.06***	-0.03***	-0.04***
	(0.004)	(0.005)	(0.01)	(0.004)	(0.006)	(0.01)
β_2	0.01***	0.12***	0.28***	0.03***	0.03***	0.000
	(0.002)	(0.04)	(0.05)	(0.004)	(0.003)	(0.001)
Sample	January 1999 - October 2006					

Note: All estimates done by GMM with Newey-West weighting matrix, standard errors in parentheses. Instruments include 6 lags of each variable (inflation, unemployment, interest rates, the weighted variance of regional unemployment rates and the additional variable included). Dependent variable is the Effective Federal Funds Rate, while ϕ_π, ϕ_{ue}, and ρ_i are coefficients on 6-month ahead inflation, 6-month ahead unemployment, and i lags of the interest rate respectively. β_1 and β_2 are the coefficients on the dispersion measure and the additional variable respectively. Statistical significance at the 1%, 5%, and 10% levels are indicated by a ***, **, and * respectively.

Table 2.5: Using Green-Book Forecasts

	United States			
	Baseline	***(1)***	***(2)***	***(3)***
c	0.09	0.74**	0.18	0.20
	(0.19)	(0.36)	(0.20)	(0.24)
ϕ_{π}	0.31***	0.33***	0.31***	0.31***
	(0.11)	(0.11)	(0.10)	(0.11)
ϕ_{gy}	0.18***	0.12***	0.15***	0.15***
	(0.04)	(0.04)	(0.05)	(0.04)
ϕ_{ue}	-0.11***	-0.24***	-0.18***	-0.15***
	(0.04)	(0.07)	(0.06)	(0.05)
ρ	0.88***	0.85***	0.86***	0.86***
	(0.04)	(0.05)	(0.05)	(0.05)
Dispersion Measure:		*var(UE)*	*UEP9010*	*UEP7525*
β		0.19***	0.13*	0.12*
		(0.07)	(0.07)	(0.07)
Sample	Fed Meetings from Jan. 1982 - Dec. 2000			

Note: All estimates done by OLS with Newey-West HAC standard errors in parentheses. Dependent variable is the target FFR chosen at each meeting, while ϕ_{π}, ϕ_{gy}, *and* ϕ_{ue} are coefficients on Green-Book forecasts of average inflation, output growth, and unemployment over current quarter through next two quarters. ρ is the coefficient on the target FFR from the previous meeting. Statistical significance at the 1%, 5%, and 10% levels are indicated by a ***, **, and * respectively.

Table 2.6: Testing whether Central Banks Minimize Dispersion of Regional Losses

	Panel A: United States *Regional UE in Levels*			Panel B: Euro-Zone *Regional UE in Levels*		
	(1)	*(2)*	*(3)*	*(1)*	*(2)*	*(3)*
c	0.33***	0.15*	0.30***	3.00***	2.77***	2.99***
	(0.09)	(0.08)	(0.07)	(0.26)	(0.29)	(0.23)
ϕ_{π}	0.08***	0.07***	0.09***	0.15***	0.11***	0.15***
	(0.02)	(0.02)	(0.02)	(0.03)	(0.04)	(0.02)
ϕ_{ue}	-0.10***	-0.04***	-0.10***	-0.30***	-0.29***	-0.31***
	(0.02)	(0.01)	(0.01)	(0.02)	(0.03)	(0.02)
ρ_1	1.25***	1.29***	1.24***	0.83***	0.86***	0.84***
	(0.05)	(0.05)	(0.05)	(0.02)	(0.02)	(0.02)
ρ_2	-0.30***	-0.33***	-0.28***			
	(0.04)	(0.05)	(0.05)			
$\beta_{var(ue)}$	0.12***		0.16***	-0.05***		-0.03***
	(0.02)		(0.03)	(0.01)		(0.01)
$\beta_{skew(ue)}$		0.04***	-0.03*		0.014***	0.008***
		(0.01)	(0.01)		(0.002)	(0.003)
Sample	Jan 1982 - Sep 2005			Jan 1999 - Oct 2006		

Note: The table presents estimates of equation (13) in the text. Estimates are done by GMM with New-West weighting matrix, with a truncation of 6 lags. The dependent variable is the interest rate. ϕ_{π}, ϕ_{ue}, and ρ_j (for j=1 or 2) are the responses of the central bank to expected inflation, expected unemployment and lag j of the interest rate respectively. β_1, β_2, and β_3 are the responses to the interaction term, the level of the variance of regional unemployment rates each month, and the skew of unemployment rates each month. We allow for the variance and skew measures to be taken across regional unemployment rates as reported (in levels) and across regional demeaned unemployment rates. Statistical significance at the 1%, 5%, and 10% levels are indicated by ***, **, and * respectively. Standard errors of estimates are in parentheses below coefficients.

Table 2.7: Decomposing the Dispersion of Regional UE Rates

	United States		Euro-Zone	
UE Percentiles:	***90th and 10th***	***75th and 25th***	***90th and 10th***	***75th and 25th***
c	0.08	-0.06	1.69***	2.58***
	(0.06)	(0.07)	(0.26)	(0.29)
ϕ_{π}	0.08***	0.11***	0.07	0.13***
	(0.02)	(0.02)	(0.05)	(0.03)
ϕ_{ue}	-0.10***	-0.11***	-0.24***	-0.26***
	(0.02)	(0.02)	(0.02)	(0.03)
ρ_1	1.27***	1.23***	0.88***	0.86***
	(0.05)	(0.05)	(0.02)	(0.02)
ρ_2	-0.30***	-0.25***		
	(0.05)	(0.05)		
β_1	0.02	-0.10	-0.10***	-0.11***
	(0.05)	(0.07)	(0.02)	(0.02)
β_2	0.18***	0.37***	0.10**	-0.05
	(0.05)	(0.08)	(0.04)	(0.04)
Wald ($\beta_1 = \beta_2$)	32.6***	42.4***	0.01	15.8***
Sample	Jan 1982 - Sep 2005		Jan 1999 - Oct 2006	

Note: This table presents estimates of equation (14) in text. β_1 is the coefficient on the difference between the 90th or 75th percentile of the regional unemployment distribution and the mean unemployment rate. β_2 is the same using the 10th or 25th percentiles. Estimates done by GMM with Newey-West weighting matrix (6 lags). Standard errors in parentheses below coefficients. Wald is the Wald test statistic of the restriction that $\beta_1 = \beta_2$. Statistical significance at the 1%, 5%, and 10% levels are indicated by ***, **, and * respectively.

Table 2.8: Regional Representation in Interest-Rate Decision-Making

Panel A: US Federal Reserve Bank Districts

District:	***1***	***2***	***3***	***4***	***5***	***6***	***7***	***8***	***9***	***10***	***11***	***12***
Share of Board Member Votes	0.12	0.15	0.09	0.00	0.14	0.00	0.13	0.07	0.00	0.10	0.11	0.08
Share of Regional Pres. Votes	0.06	0.20	0.07	0.10	0.07	0.07	0.10	0.06	0.06	0.07	0.07	0.07
Share of Total Votes	0.10	0.17	0.08	0.04	0.11	0.03	0.12	0.07	0.03	0.08	0.09	0.08
Share of Population	0.05	0.07	0.08	0.04	0.09	0.14	0.13	0.04	0.03	0.05	0.07	0.19
Mean UE gap	-1.1	0.1	-0.1	0.4	-0.7	0.2	0.3	0.1	-1.4	-0.9	0.3	0.5

Panel B: Members of European Central Bank

Country:	***AU***	***BE***	***FI***	***FR***	***DE***	***GR***	***IR***	***IT***	***LX***	***ND***	***PR***	***ES***
Share of Total Votes	0.08	0.06	0.09	0.10	0.11	0.10	0.06	0.11	0.06	0.10	0.06	0.11
Share of Population	0.03	0.03	0.02	0.19	0.27	0.04	0.01	0.19	0.00	0.05	0.03	0.13
Mean UE gap	-4.3	-0.8	0.7	0.8	-0.5	1.8	-4.0	0.5	-5.2	-5.1	-3.2	2.4

Note: US Federal Reserve Bank Districts are based in Boston (1), New York (2), Philadelphia (3), Cleveland (4), Richmond (5), Atlanta (6), Chicago (7), St. Louis (8), Minneapolis (9), Kansas City (10), Dallas (11), and San Francisco (12). Members of the ECB are Austria (AU), Belgium (BE), Finland (FI), France (FR), Germany (DE), Greece (GR, since Jan. 1 2001), Ireland (IR), Italy (IT), Luxembourg (LX), Netherlands (ND), Portugal (PR), and Spain (ES). Data for FOMC votes is from Meade and Sheets (2005) and contains votes by Board members and Regional Presidents. For the ECB, votes are of members of the Governing Council, which include Executive Board members and Presidents of each national central bank. Mean UE gaps are average difference between regional unemployment rates and the aggregate unemployment rates, from 1982:01 to 2005:09 for US and from 1999:01 to 2006:10 for Euro-Zone.

Table 2.9: Does Regional Representation in Interest-Rate Decisions Matter?

	Panel A: United States			Panel B: Euro-Zone		
	(1)	*(2)*	*(3)*	*(1)*	*(2)*	*(3)*
c	0.81**	0.25	0.29	2.86***	1.92***	2.78***
	(0.40)	(0.31)	(0.31)	(0.53)	(0.47)	(0.52)
ϕ_{π}	0.34***	0.31***	0.32***	0.11**	0.10*	0.08**
	(0.12)	(0.11)	(0.11)	(0.04)	(0.05)	(0.04)
ϕ_{gy}	0.11**	0.14**	0.14***			
	(0.05)	(0.06)	(0.05)			
ϕ_{ue}	-0.25***	-0.18***	-0.16***	-0.27***	-0.20***	-0.23***
	(0.08)	(0.07)	(0.06)	(0.05)	(0.06)	(0.04)
ρ	0.85***	0.86***	0.86***	0.85***	0.91***	0.85***
	(0.05)	(0.05)	(0.04)	(0.04)	(0.04)	(0.04)
Dispersion measure:	*var(UE)*	*UEP9010*	*UEP7525*	*var(UE)*	*UEP9010*	*UEP7525*
β_{disp}	0.18***	0.12***	0.09	-0.05***	-0.03	-0.11***
	(0.06)	(0.04)	(0.07)	(0.01)	(0.03)	(0.04)
β_{voting}	-0.22	-0.16	-0.23	0.10	-0.13	0.20*
	(0.49)	(0.47)	(0.51)	(0.10)	(0.15)	(0.11)
Sample	Jan 1982: Sep 2005 (FOMC meetings)			Jan 1999: Oct 2006 (monthly)		

Note: Estimates for the US are done by OLS using GreenBook Forecasts of future inflation (ϕ_{π}), output growth (ϕ_{gy}), and unemployment (ϕ_{ue}) with interest rates measured by the target FFR on data with frequency of FOMC meetings. Estimates for ECB are done by 2SLS with 6-month ahead values of inflation and unemployment with interest rates measured by interbank overnight rate. β_1 is the coefficient on each measure of the cross-sectional regional dispersion of unemployment rates. β_2 is the coefficient on the weighted sum of differences between regional and aggregate unemployment rates, where the weights are the voting share of each region in the interest-rate decision process that period. All standard errors are Newey-West HAC. Statistical significance at the 1%, 5%, and 10% levels are indicated by ***, **, and * respectively.

Figure 2.1: US Aggregate Variables

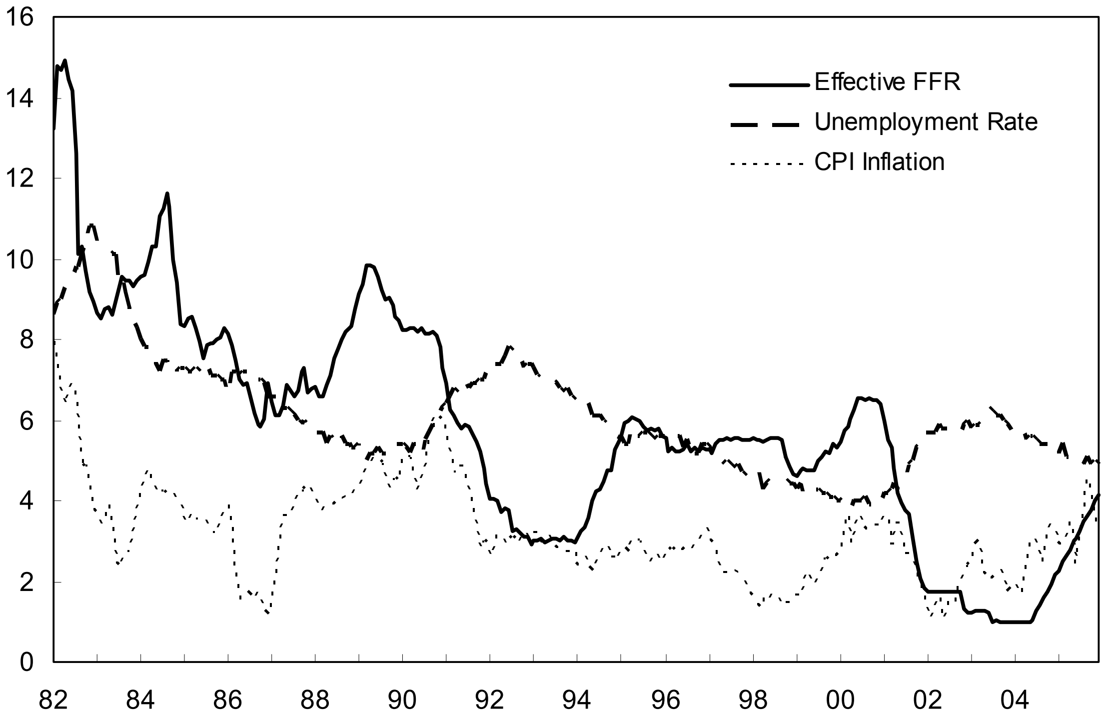

Figure 2.2: Measures of the Regional Dispersion of Unemployment Rates

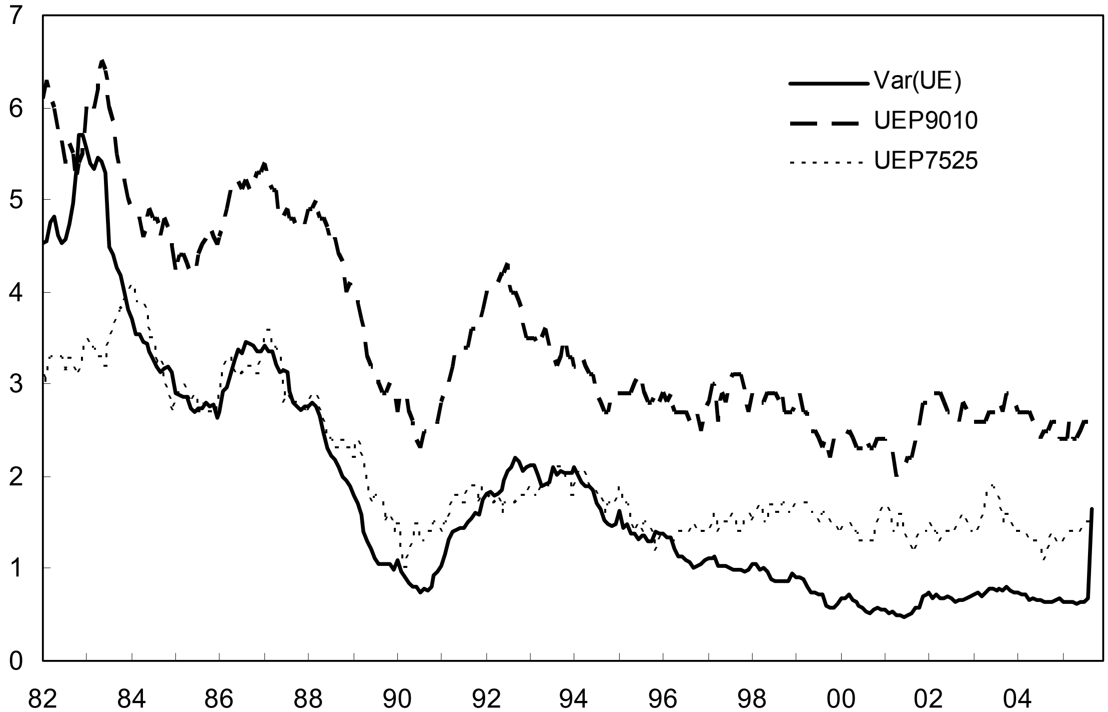

Figure 2.3: European Aggregate Variables

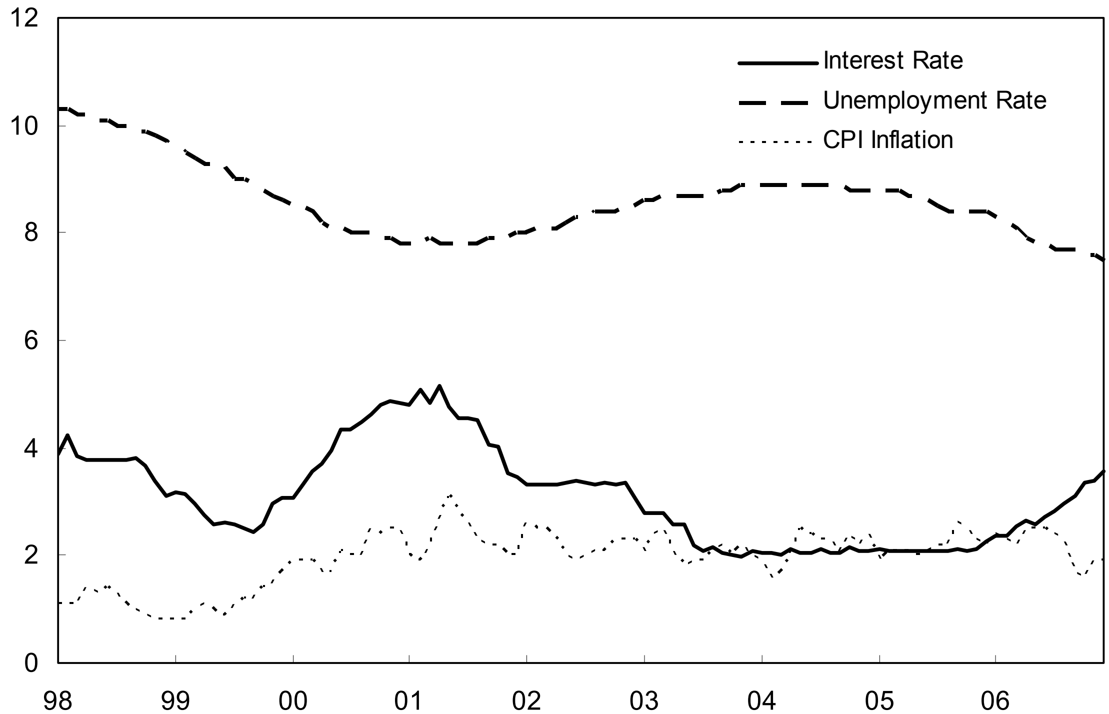

Figure 2.4: Measures of Regional Dispersion of Unemployment Rates in Euro-Area

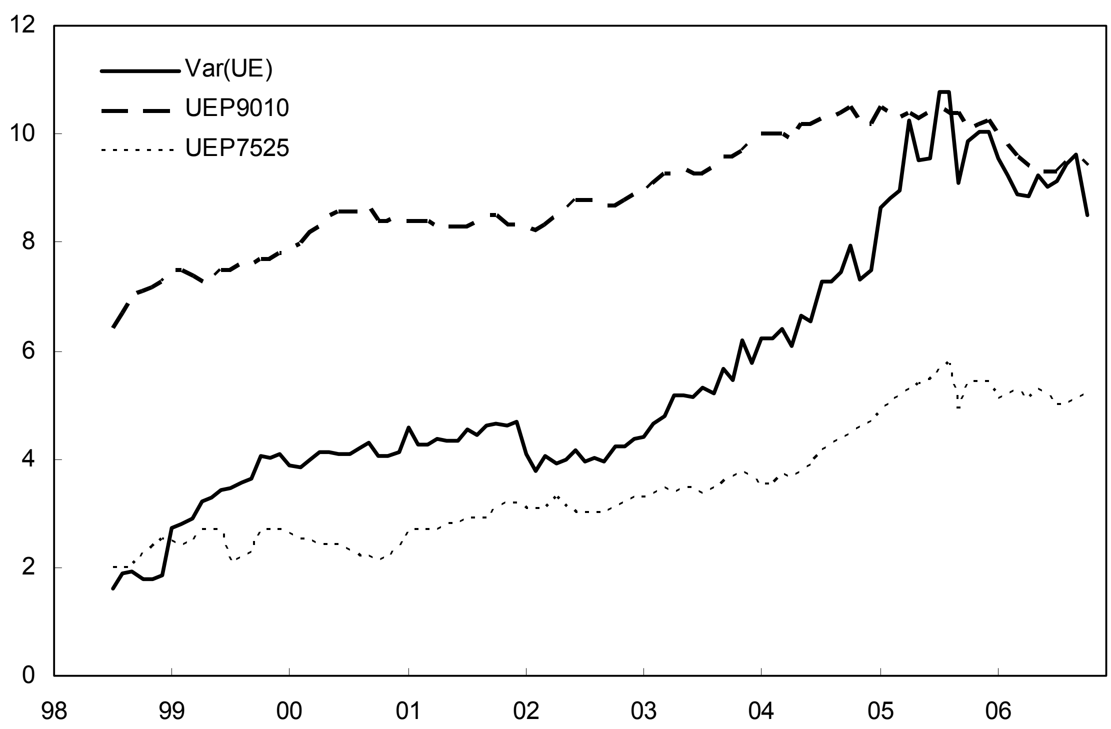

Figure 2.5: Scatter Plot of US Unemployment Dispersion and Interest Rate Prediction Error Excluding Dispersion Measure.

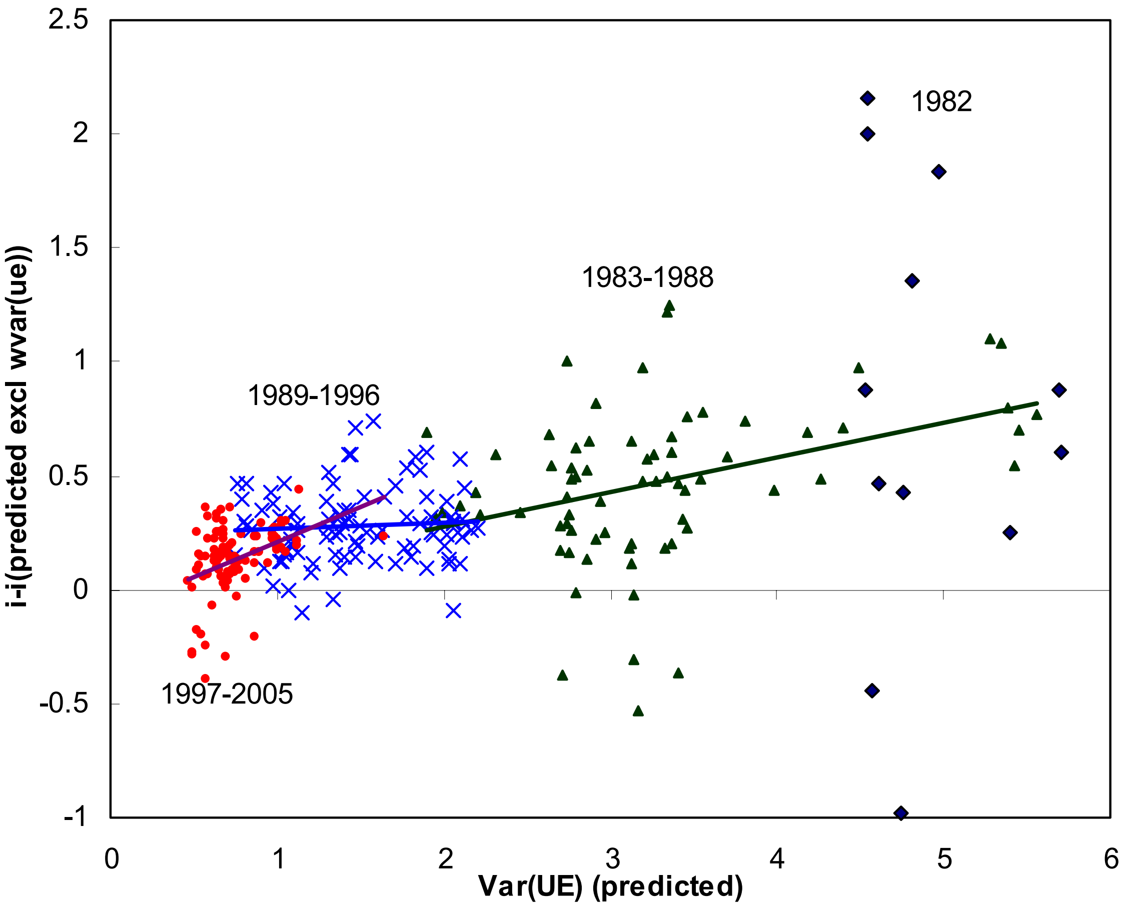

Note: Figure presents a scatter of plot of predicted cross-sectional variance of unemployment rates (based on projection of instruments) against difference between actual interest rates and predicted interest rates using non-dispersion (predicted) RHS variables of equation (2.3) in text with coefficients from Table 2.1. Scatter plot is broken into four time periods, indicated by different markers. Each period, excluding 1982, has trend line plotted.

Figure 2.6: Dynamic Cross-Correlation of Cross-Sectional Variance of Unemployment Rates with Inflation and Aggregate Unemployment

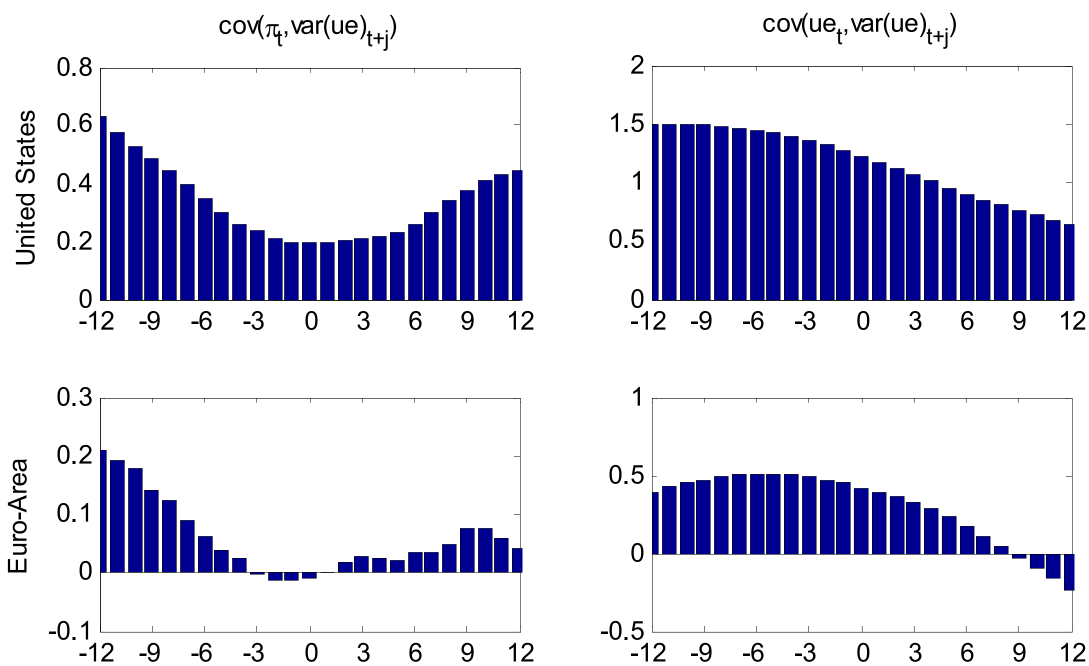

Appendix 1

Proof of Proposition 1 and Corollary 1:

Note first that minimization of (3.5) over inflation yields the optimality condition

$$\sum_i \omega_i \frac{dL_i}{d\pi} + \kappa \sum_i \omega_i \left[L_i - \sum_j \omega_j L_j \right] \left[\frac{dL_i}{d\pi} - \sum_j \omega_j \frac{dL_j}{d\pi} \right] = 0 \qquad \text{(A1.1)}$$

Combining (3.1) and (3.2), the first bracketed term can be represented as:

$$L_i - \sum_j \omega_j L_j = \frac{\lambda}{2}\left(\varepsilon_i^2 - \sum_j \omega_j \varepsilon_j^2 \right) - \alpha\lambda\left(\pi - \pi^e\right)\left(\varepsilon_i - \varepsilon_a\right) \qquad \text{(A1.2)}$$

Differentiating (3.1) and combining with (3.2) we can also show the second bracketed term to be:

$$\frac{dL_i}{d\pi} - \sum_j \omega_j \frac{dL_j}{d\pi} = -\alpha\lambda\left(\varepsilon_i - \varepsilon_a\right) \qquad \text{(A1.3)}$$

Differentiating (3.1) and substituting the result, along with (A1.2) and (A1.3), into (A1.1) and rearranging yields, in terms of the optimal inflation level π^*:

$$\left(1+\alpha^2\lambda\right)\pi^* - \alpha\lambda\varepsilon_a - \alpha^2\lambda\pi^e = \kappa\alpha\lambda^2 \left[\sum_i \omega_i \left[\frac{\left(\varepsilon_i - \varepsilon_a\right)}{2}\left(\varepsilon_i^2 - \sum_j \omega_j \varepsilon_j^2 \right) - \alpha\left(\pi^* - \pi^e\right)\left(\varepsilon_i - \varepsilon_a\right)^2 \right] \right] \qquad \text{(A1.4)}$$

which simplifies to:

$$\left(1+\alpha^2\lambda\right)\pi^* - \alpha\lambda\varepsilon_a - \alpha^2\lambda\pi^e = \kappa\alpha\lambda^2 \left[\frac{1}{2}\left(\sum_i \omega_i \varepsilon_i^3 - \varepsilon_a \sum_i \omega_i \varepsilon_i^2 \right) - \alpha\left(\pi^* - \pi^e\right)\sum_i \omega_i \left(\varepsilon_i - \varepsilon_a\right)^2 \right] \qquad \text{(A1.5)}$$

Now note that the weighted variance of the observed regional shocks is given by

$$\operatorname{var}\left(\varepsilon_i\right) = \sum_i \omega_i \left(\varepsilon_i - \varepsilon_a\right)^2 = \sum_i \omega_i \varepsilon_i^2 - \varepsilon_a^2 \qquad \text{(A1.6)}$$

and the weighted skew of these observed shocks is

$$\begin{aligned} skew(\varepsilon_i) = \sum_i \omega_i (\varepsilon_i - \varepsilon_a)^3 &= \sum_i \omega_i \varepsilon_i^3 - 3\varepsilon_a \sum_i \omega_i \varepsilon_i^2 + 2\varepsilon_a^3 \\ &= \sum_i \omega_i \varepsilon_i^3 - \varepsilon_a^3 - 3\varepsilon_a \operatorname{var}(\varepsilon_i) \end{aligned} \tag{A1.7}$$

where the last equality makes use of (A1.6).

Substituting both (A1.6) and (A1.7) into equation (A1.5) yields:

$$\begin{aligned} \left(1+\alpha^2\lambda+\alpha^2\lambda^2\kappa \operatorname{var}(\varepsilon_i)\right)\pi^* = \alpha\lambda\left(\varepsilon_a+\alpha\pi^e\right)+\frac{\kappa\alpha\lambda^2}{2}\left[skew(\varepsilon_i)+2\varepsilon_a \operatorname{var}(\varepsilon_i)\right] \\ +\alpha^2\lambda^2\kappa \operatorname{var}(\varepsilon_i)\pi^e \end{aligned}$$

(A1.8)

Defining $\Psi \equiv 1+\kappa\lambda \operatorname{var}(\varepsilon_i)$, we can rewrite (A1.8) as:

$$\left(1+\alpha^2\lambda\Psi\right)\pi^* = \Psi\alpha\lambda\left(\varepsilon_a+\alpha\pi^e\right)+\frac{\alpha\lambda^2\kappa}{2} skew(\varepsilon_i) \tag{A1.9}$$

Note that $\sigma_u^2 \equiv \sum_i \omega_i \left[(u_i-\bar{u}_i)-(u_a-\bar{u}_a)\right]^2 = \sum_i \omega_i \left[\varepsilon_i-\varepsilon_a\right]^2 \equiv \operatorname{var}(\varepsilon_i)$

and $\quad skew(u_i-\bar{u}_i) \equiv \sum_i \omega_i \left[(u_i-\bar{u}_i)-(u_a-\bar{u}_a)\right]^3 = \sum_i \omega_i \left[\varepsilon_i-\varepsilon_a\right]^3 \equiv skew(\varepsilon_i)$

then defining $f\left(\sigma_u^2\right) \equiv \dfrac{\Psi(1+\alpha^2\lambda)}{1+\alpha^2\lambda\Psi}$ and $g\left(\sigma_u^2\right) \equiv \dfrac{\alpha\lambda^2\kappa}{2\left(1+\alpha^2\lambda\Psi\right)}$ yields

$$\pi^* = \pi^{opt} f\left(\sigma_u^2\right) + skew(u_i-\bar{u}_i)\, g\left(\sigma_u^2\right)$$

Note that $\Psi = 1+\kappa\lambda \operatorname{var}(\varepsilon_i) \geq 1 \Rightarrow f\left(\sigma_u^2\right) \geq 1$ and $g\left(\sigma_u^2\right) \geq 0$.

Finally, also note that:

$$\frac{df}{d\sigma_u^2} = \left[\frac{1+\alpha^2\lambda}{1+\alpha^2\lambda\Psi} - \frac{\Psi\alpha^2\lambda\left(1+\alpha^2\lambda\right)}{\left(1+\alpha^2\lambda\Psi\right)^2} \right] \frac{d\Psi}{d\sigma_u^2}$$

$$= \left[\frac{1+\alpha^2\lambda}{1+\alpha^2\lambda\Psi} \right] \frac{d\Psi}{d\sigma_u^2} > 0$$

while $\frac{dg}{d\sigma_u^2} < 0$ since $\frac{d\Psi}{d\sigma_u^2} > 0$.

Appendix 2

Now consider the same loss function for the central bank as (3.4). Substituting the Phillips Fan of (3.10) into the loss function and taking the first order conditions with respect to the policy π yields:

$$0 = \pi + \lambda[\sum_i \omega_i[(-(C+\gamma\varepsilon_i)(\pi-\pi^e)+\varepsilon_i)(-(C+\gamma\varepsilon_i))] \qquad \text{(A2.1)}$$

Expanding and taking the summation where possible, this can be rewritten as:

$$0 = \pi + \lambda[(C^2 + 2C\gamma\varepsilon_a + \gamma^2 \sum_i \omega_i\varepsilon_i^2)(\pi-\pi^e) - C\varepsilon_a - \gamma\sum_i \omega_i\varepsilon_i^2] \qquad \text{(A2.2)}$$

By definition, the weighted variance of the shocks is given by $\sigma_\varepsilon^2 = \sum\omega_i\varepsilon_i^2 - \varepsilon_a^2$, or

$$\sum_i \omega_i\varepsilon_i^2 = \varepsilon_a^2 + \sigma_\varepsilon^2 \qquad \text{(A2.3)}$$

Plugging (A2.3) into (A2.2) and simplifying yields:

$$0 = \pi + \lambda[((C+\gamma\varepsilon_a)^2 + \gamma^2\sigma_\varepsilon^2)(\pi-\pi^e) - \varepsilon_a(C+\gamma\varepsilon_a) - \gamma\sigma_\varepsilon^2] \qquad \text{(A2.4)}$$

Taking the total derivative of this term with respect to π and σ_ε^2, one can show:

$$\frac{d\pi}{d\sigma_\varepsilon^2} = \frac{\lambda\gamma(1-\gamma(\pi-\pi^e))}{1+\lambda\gamma^2\sigma_\varepsilon^2 + \lambda(C+\gamma\varepsilon_a)^2} \qquad \text{(A2.5)}$$

This term is positive, given the no free lunch assumption of above (that is, $(1-\gamma(\pi-\pi^e))\geq 0$).

Now, for the Phillips Fan given in (3.10), the variance of unemployment around the natural level is:

$$\sigma_u^2 = \sigma_\varepsilon^2(1-\gamma(\pi-\pi^e))^2 \qquad \text{(A2.6)}$$

The total derivative of this equation is:

$$d\sigma_u^2 = (1-\gamma(\pi-\pi^e))^2 d\sigma_\varepsilon^2 - 2\gamma\sigma_\varepsilon^2(1-\gamma(\pi-\pi^e))d\pi \qquad \text{(A2.7)}$$

Solving (A2.5) for $d\sigma_\varepsilon^2$, and plugging it into (A2.7) leads to an expression for the correlation between the variance of unemployment and the level of inflation:

$$\frac{d\pi}{d\sigma_u^2} = \frac{\lambda\gamma}{(1-\gamma(\pi-\pi^e))(1-\lambda\gamma\sigma_\varepsilon^2+\lambda(C+\gamma\varepsilon_a)^2)} \tag{A2.8}$$

Lastly, starting with the first order condition (A2.1), take the total derivative with respect to π and ε_i. This yields:

$$0 = d\pi + \lambda[[\sum_i \omega_i (C+\gamma\varepsilon_i)^2 d\pi] + [2\gamma(C+\gamma\varepsilon_i)(\pi-\pi^e) - C - 2\gamma\varepsilon_i]d\varepsilon_i] \tag{A2.9}$$

Noting that the expression for the Phillips Fan (3.10) implies $(C+\gamma\varepsilon_i)(\pi-\pi^e) = -(u_i-\overline{u}_i)+\varepsilon_i$, we can plug this into (A2.9) and cancel out the $2\gamma\varepsilon_i$ term. Rearranging, we arrive at:

$$\frac{d\pi}{d\varepsilon_i} = \frac{\lambda[2\gamma(u_i-\overline{u}_i)+C]}{1+\lambda\sum_i \omega_i (C+\gamma\varepsilon_i)^2} \tag{A2.10}$$

Appendix 3

Data Details for Voting Measures

<u>A- States associated with each Fed District</u>

District 1: Connecticut, Maine, Massachusetts, New Hampshire, Rhode Island, and Vermont.

District 2: New York

District 3: Pennsylvania, Delaware, New Jersey.

District 4: Ohio.

District 5: Washington DC, Maryland, Virginia, North and South Carolina, West Virginia.

District 6: Alabama, Florida, Georgia, Louisiana, Mississippi, and Tennessee.

District 7: Iowa, Illinois, Indiana, Michigan, and Wisconsin.

District 8: Arkansas, Missouri, and Kentucky.

District 9: Minnesota, Montana, North and South Dakota.

District 10: Colorado, Kansas, Nebraska, Oklahoma, Wyoming, and New Mexico.

District 11: Texas.

District 12: California, Arizona, Utah, Nevada, Oregon, Washington, Idaho, Alaska, and Hawaii.

Because districts typically include parts of states, this division only approximately captures the division of states across districts. It was necessary to divide districts into states because employment data by month is only available at the state level.

B- ECB members of the Governing Council.

Every nation has one representative through the head of its central bank. Greece joined in January 2001. In addition, the following were members of the Executive Board and had voting rights in interest rate decisions:

President: Duisenberg (ND) from Jan. 1999 to Oct. 2003. Replaced by Trichet (FR) in Nov. 2003, to present.

Vice-President: Noyer (FR) from June 1998 to May 2002. Replaced by Papademos (GR) in June 2002, to present.

Members:

Solans (ES) from June 1998 to May 2004. Replaced by Gonzalez-Paramo (ES) June 2004 to present.

Hamalainen (FI) from June 1998 to May 2003. Replaced by Tumpel-Gugerell (AU) June 2003 to present.

Issing (DE) from June 1998 to May 2006. Replaced by Stark (DE) June 2006 to present.

Padoa-Schioppa from June 1998 to May 2005. Replaced by Bini Smagghi (IT) June 2005 to present.

References

Aksoy, Yunus, Paul De Grauwe, and Hans Dewachter. 2002. Do Asymmetries Matter for Monetary Policy? *European Economic Review.* Vol. 46: 443-469.

Barro, Robert J. and David B. Gordon. 1983. A Positive Theory of Monetary Policy in a Natural Rate Model. *Journal of Political Economy.* Vol. 91: 589-610.

Clarida, Richard, Jordi Gali, and Mark Gertler. 1998. Monetary Policy Rules in Practice: Some International Evidence. *European Economic Review.* Vol. 42 No. 6: 1033-1067.

Clarida, Richard, Jordi Gali, and Mark Gertler. 2000. Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory. *Quarterly Journal of Economics.* Vol. 115 No. 1: 147-180.

Christiano, Lawrence J and Wouter J den Haan. 1996. Small-Sample Properties of GMM for Business Cycle Analysis. *Journal of Economic and Business Statistics*. Vol. 14 No. 3: 309-327.

Dominguez, Kathryn. 2006. The European Central Bank, the Euro and Global Financial Markets. Forthcoming in *Journal of Economic Perspectives*.

Dixit, Avinash. 2000. A Repeated Game Model of Monetary Union. *The Economic Journal.* Vol. 110: 759-780.

Fuchs, William and Francesco Lippi. 2006. Monetary Union with Voluntary Participation. *Review of Economic Studies*. Vol. 73. Pp. 437-457.

Hansen, Lars Peter, John Heaton, and Amir Yaron. 1996. Finite-Sample Properties of Some Alternative GMM Estimators. *Journal of Economic and Business Statistics*. Vol. 14 No. 3: 262-280.

Heinemann, Friedrich and Felix P. Huefner. 2004. Is the View from The Euro-Tower Purely European? National Divergence and ECB Interest Rate Policy. *Scottish Journal of Political Economy.* Vol. 51, No.4: 544-558.

Meade, Ellen E. and D. Nathan Sheets. 2005. Regional Influences on FOMC Voting Patterns. *Journal of Money, Credit, and Banking*. Vol. 37, No. 4: 661-677.

Orphanides, Athanasios. 2001. Monetary Policy Rules Based on Real-Time Data. American Economic Review, Vol. 91 No. 4. Pp. 964-985.

Romer, Christina D. and David H. Romer. 2004. A New Measure of Monetary Shocks: Derivation and Implications. *American Economic Review.* Vol. 94 No. 4: 1055-1084.

Taylor, John B. 1993. Discretion versus Policy Rules in Practice. *Carnegie Rochester Conference Series on Public Policy*. Vol 39: 195-214.

Woodford, Michael. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy.* Princeton: Princeton University Press.

Chapter 3

Emotions in the Utility Function

Abstract

In this essay, I suggest that simple economics models can often be successfully used in evaluating the role of emotional factors in decision making. Particularly, I consider the implications for agents strictly adhering to the principles of constrained optimization. I show that this "neoclassical" approach implies that emotions are a feature of agents' utility functions. I then present the basic theoretical apparatus of this model, and apply it to a wide variety of emotionally linked behaviors. I find that there are many examples that seem to indicate emotions are affecting choices in a manner consistent with some of the less-obvious predictions of this purely neoclassical framework. While it is unlikely that the standard approach can hope to account for all behaviors, economic studies of micro-level behavior can benefit from, at least at the outset, a benchmark model that incorporates emotions in a neoclassical way.

1 Introduction

Emotions are a fundamental feature of human experience. Unless we happen to be suffering a specific sort of brain damage, we are all familiar with what the Oxford English Dictionary calls the "strong feeling[s], such as joy, anger, or sadness" that arise, not through a process of reasoning, but as a seemingly unbidden response to either an external stimulus, or our own behavior.

Like all characteristics that develop through natural selection, emotions exist today because they, at least at some time, aided survival. Just as the giraffe with the longest neck could eat more (and thus become stronger) than others in his herd, at some point along the evolutionary path the ability for some proto-humans to have feelings as varied as fear, love, happiness, or guilt gave those individuals an advantage in passing along their genetic material. However, while it is clear how the giraffe's neck allows it to pick off the highest leaves, the physical manifestation of a given emotion is far more subtle, if it exists at all. Instead, an emotion will tend to benefit an individual *only* if it is able to modify the person's *behavior* in a manner that promotes the transmission of genes to the next generation.

It is this premise that underlines this essay: Human emotions exist, and they exist only because they are able to affect behavior. While the advantages conferred by emotions far back along the evolutionary path may no longer be advantages in modern times, there is no reason to suspect emotions have become irrelevant in decision making. Rather, our prior ought to be that emotions can and do continue to influence behavior. The question for the social scientist to try and answer, is how?

The field of Psychology has, perhaps naturally, taken the lead in the study of emotion. Economists have historically tended to ignore the role of emotions altogether. Even those who attempt to analyze emotions often view them as corrupting the decision-making process (Loewenstein 1996). Recent work by neuroscientists (Damasio 1994) suggest that in fact emotions can serve to actually improve decision-making, casting some doubt on the premise that cold-blooded (that is, non-emotional) evaluation of problems is the superior way for individuals to make decisions.

In this essay, I argue that standard Economics can in fact yield important insights on the relationship between emotion and behavior, and do so cheaply. The standard – or "neoclassical" – approach that individuals seek to maximize their utility is quite amenable to incorporating emotions. As we will see, a merger of constrained optimization with emotions implies a simple, yet provocative result: if they affect behavior, then emotions *must* be a component of the utility function.

This last point is, I believe, a subtle one that is overlooked in other discussions by economists considering how emotions influence behavior. While the economics literature rarely takes on emotions on a formal basis, when it does, one approach is to assume a) that agents are maximize utility, *and* b) that the emotions are in the utility function just like any other utility generating component (Elster, 1998, p. 64). I argue that the second assumption is unnecessary. Assuming agents are utility maximizers – and that the emotions are in fact important – implies that the emotions are a part of the utility function. On the other hand, if emotions are not in the utility function – and emotions are still important to behavior – then the agent cannot be optimizing. One goal of this essay is to establish this claim, and sketch out some of the features of neoclassical theory that incorporates emotions.

The second objective of this paper is to see how far the neoclassical approach can be pushed. To this end, I consider a number of real-world behaviors through the lens of a neoclassical model and attempt to gauge where the model succeeds or fails. Perhaps the most revealing approach here is the application of Young's theorem. I show that some of the non-obvious mathematical implications of the model seem to jibe with our common experiences in the world. In addition, I propose an interpretation of "cognitive dissonance" – a phenomenon commonly viewed as a failure by agents to form rational beliefs or expectations – in which the supposed "non-rational" beliefs are simply tools used for emotional gain, but do not reflect an agent's underlying beliefs.

2 Theoretical Principles

The theoretical implications of incorporating emotions into the standard decision-maker's problem are straightforward. In this section I highlight some of the basic price theory that underlies the arguments.

2.1 Emotions in the Utility Function

Here, I show that a neoclassical model that includes emotions does not need to make any assumptions that emotions are a component of the utility function. Rather, emotions *must* be in the utility function as a consequence of two key assumptions: utility maximization, and behavior modification. I state and informally prove this fact here:

Proposition 1: Ceteris peribus, if a utility maximizing agent with concave utility function U(C,.) in C alters her behavior as her emotional state changes (M), then emotion is an argument of her utility function. That is, utility is given by U(C,M,.).

Proof: See figure 3.1. Utility is plotted as a function of C on these graphs. Suppose that the agent's emotional state changes from M to M', and the optimal level of C has changed. In all possible cases where this is so, the slope of the utility function has changed at the initial level of C. That is, U_C is changed as M changes. Equivalently, $U_{CM} \neq 0$. This cannot be true if $U_M=0$ at all levels of C. Thus, in general $U_M \neq 0$ and therefore emotion must be an argument of the utility function.

The *ceteris peribus* assumption is included to control for changes in income. A lottery winner might feel happier and also increase his consumption, but this does not necessarily imply that happiness changes the marginal utility of his consumption. An alternative way to eliminate income effects is to assume preferences are quasi-linear in C.

Figure 3.1 shows three possible scenarios in which a change in emotion corresponds to a change in behavior. In panel (A), utility is higher for all levels of C. In panel (B), utility is lower for all levels of C. And finally in panel (C), utility is higher for some levels of consumption and lower for others. Thus, we can see that regardless of what is the actual effect of the emotion on behavior, it is consistent with an emotion that is unambiguously "good" (i.e. one always raises utility), "bad," or "sometimes good and sometimes bad." That this can be utilized in the context of widely varied emotions is appealing.

There is a Corollary to Proposition 1: if a change in emotion leads to changes in behavior, but emotion is NOT in the utility function, then the agent cannot be acting in a utility maximizing way:

> Corollary 1: Ceteris Peribus, if an agent has a concave utility function U(C,.) and the emotion is not in her utility function, then if a change in emotion from M to M' induces a choice of C that was strictly dominated when the emotion was at level M, then the agent is not acting in a utility maximizing way.
>
> Proof: If emotion is not in the utility function, then U(C,.) is unchanged as we move from M to M'. Changing the level of C to a previously strictly dominated choice when the utility function is unchanged implies that the agent cannot be maximizing utility.

Corollary 1 is relevant when considering some attempts in the behavioral literature to incorporate "emotional" features into an agent's constraint set, rather than into the objective. Gul and Pesendorfer (2001) and Ozdenoren, Salant, and Silverman (2006), are two examples. They suggest that self-control and willpower, respectively, are best modeled as constraints. Note that the models in these cases are dynamic, while I am focusing on static ones. However, there is an implicit assumption in these papers that the proposed emotional constraints are necessary to prevent deviations from optimal consumption paths. That is, when these emotions are not in themselves valuable to utility, they influence behavior only when the behavior would otherwise be non-optimal in some sense, as is suggested by Corollary 1.

Panel (D) of Figure 3.1 also suggests that the sign of the cross partial derivative U_{CM} is linked to the interplay of emotions and behavior in a consistent way.

The result stated in the next Proposition turns out to be very important when we look for real-world instances of emotion operating from within the utility function.

> Proposition 2: Ceteris peribus, if $dC/dM>0$, then $U_{CM}>0$. Alternatively, if $dC/dM<0$, then $U_{CM}<0$.
>
> Proof: See figure 1. Suppose M'>M (that is, the emotional state M' is more "intense" than M), and $dC/dM>0$. Figure 1 displays all three cases in which this is possible. In each instance, the slope of the utility function is higher at the original level of consumption. That is, U_C increases as M increases. Equivalently, $U_{CM}>0$. The reverse is true if $dC/dM<0$.

Note that in these pictures there is no motivation for the shift from M to M'. This is for the sake of generality: the results hold regardless of whether the change in emotional state is exogenous or endogenous. That is, it does not matter if a change in emotional state is imposed from the outside (say, feeling sad if a family member dies) or is internally motivated (say, if one seeks psychotherapy to reduce depression).

With emotions included as an argument of the utility function, it is desirable that our model allow for emotions to be, at least in part, endogenously determined. However, if emotions are indeed endogenous our model will require more structure with respect to how the emotions are created. This leads to the next feature of the model.

2.2 The Production Function for an Emotion

Emotions are not tangible goods. There is no relationship between physical capital and labor that is going to churn out emotions. Nor can one purchase an emotion in the marketplace. At best, the market can only provide intermediate goods, such as psychoactive drugs or scary movies.

At the same time, emotions do not appear out of thin air. They are created in the mind – induced through the interplay of the environment, one's own actions, and the neurons in the brain. It seems reasonable to begin to think of these latter items as the "factors of production" for an emotion, and embed them within what I'll call the "production function for an emotion," or "EPF."

> Definition: Let M(X,S) be an agent's production function for an emotion M, where X is a vector representing the individual's actions, and S is a vector representing the state (variables that are outside the agent's control).

The vectors X and S are essentially the endogenous and exogenous factors, respectively, that contribute to one's emotional condition. Using the examples above, the action of going to visit a psychologist to reduce depression would be an action in X, whereas the death of a family member would be a change in S. In addition to exogenously induced emotions, the state could also include emotions carried forward from earlier periods. If one enters our time period of interest in a bad mood, for instance, this may impact how emotions are generated in the current period.

The production function M(X,S) may be very complicated. In addition to complexity, agents may not completely understand how their own emotions are produced. This would be particularly true for new experiences. Having never taken action X_0 while in state S_0, how are agents to know their emotional response to it? To the extent that agents incorrectly gauge how their emotions are formed, they may fail to optimize their utility.

This is a legitimate concern, but a concern that applies just as much to "non-emotional" economics. Even absent any emotional considerations, people often do not have all the information they need to make the utility maximizing decision. Furthermore, even when they do have the information, they are imperfect processors of that information. We are not all chess grandmasters, after all. A mis-specification or mis-evaluation of the EPF – and therefore failing to act optimally – is the same type of problem.

I abstract from this concern for the remainder of this piece, and will take as given that agents know their own production functions for emotions. While people may not in fact precisely know how their emotions are produced, they frequently will have a very good idea of what makes them feel happy, afraid, guilty etc. To the extent that they do make "errors," it may well be that they were anticipating a given decision would lead to a different emotion than was actually produced. These types of "emotional mistakes," especially if the agent makes the error consistently, could be an interesting avenue of research, but I will not discuss it further here.

While "production function for an emotion" is a new term, the idea that emotions can be defined in a mathematical form has sporadically appeared in the literature. Becker initially developed the notion that production functions can be applied to a wider variety of problems than the fabrication of goods. As for explicit production functions for emotions, Ackerlof and Dickens (1982) include a production function for fear in their paper on cognitive dissonance. Glaeser (2002), in a working paper version of his paper on the politics of hatred, proposes a production function for hatred. More recently, Kimball and Willis (2006) suggest a "household production function for happiness," which serves as the backbone of their work on utility and happiness.

Laibson (2001) suggests a model of addiction that has a feature similar in many ways to an EPF. He suggests that certain environmental cues can change the marginal utility of consumption for some goods; for instance, the smell of baking bread might make eating bread a more desirable activity. I view these cues as arguments in one's EPF. The smell of the bread may induce a craving for bread – but the degree of that craving is also determined by exogenous factors (such as, whether or not one has just finished eating, or if one knows the baker to be a bad man), or endogenous (such as, the individual may be on a diet, which makes the craving even worse, or his mind is simply occupied with other concerns and he ignores the smell.) The proposed model here is essentially a more general version of Laibson's.

2.3 Price Theory with Emotions in the Utility Function

Given the motivation of Proposition 1 and the notion of production functions for emotions, it is natural to now write utility in the general form U(C(X),M(X;S);S). The C, X, M and S may of course represent vectors of their respective variables. I treat C(X) very flexibly throughout the remainder of the essay, but it can perhaps best be thought of as generic type of consumption resulting from actions X. In a special case, X could represent consumption itself (if C(X)=X). The emotional state M is produced by an EPF as described above. The vector of state variables S influences utility both through the emotional channel and a more direct one. An example of the latter could be if your stomach is full, the utility from consuming a gourmet meal is lowered, regardless of the emotional impact that might result from such a meal.

Furthermore, the emotional punch of the gourmet meal might also be lowered if one has already eaten.

Notice that in principle one could rewrite this utility function as U(X;S). This moves us towards the normally presented form for utility that completely ignores any mention of emotion. Indeed, this is one appealing feature of this manner of incorporating emotions into the utility function: it is entirely consistent with standard economic arguments. All that has been done is added a layer of additional structure to the utility function.

In many applications, the added complexity may not be important. However, in studies of micro-level behavior, this extra dimension may be crucial. In fact, I believe a model that explicitly considers emotions should be the default choice for behavioral economics. At the very least, distinguishing the utility that is generated from specific emotions – and the intensity of those emotions – from the utility generated through other means may yield a much richer description of behavior.

To help develop our intuition, let's focus on the simplest case. We'll assume all the relevant functions are continuous, and our agent has only one decision to make: the magnitude of a single action x. The maximization problem for the agent in this model is to choose a feasible x that maximizes his utility.

$$\begin{aligned} &\underset{x}{Max}\ U(C(x), M(x;S); S) \\ &s.t.\ x \in \mathbb{X} \end{aligned} \tag{2.1}$$

Suppose our agent picks an interior value for x. Then, the first order condition is:

$$U_C \frac{dC}{dx} + U_M \frac{dM}{dx} = 0 \tag{2.2}$$

This implies the standard condition that the marginal rate of transformation (MRT) equals the marginal rate of substitution (MRS):

$$MRT = \frac{\frac{dC}{dx}}{\frac{dM}{dx}} = \frac{dC}{dM} = \frac{-U_M}{U_C} = MRS \tag{2.3}$$

An implication of (2.3) is that if C and M are both either increasing or decreasing in x – or equivalently, the MRT is positive – then U_C and U_M must have opposite signs at the optimum value of x. If the MRT is negative, then U_C and U_M must have the same sign.

This is reminiscent of a familiar idea in labor economics. If the wage rate is positive, then the marginal rate of transformation dC/dN will be positive (where N is number of hours worked). With the marginal utility of consumption assumed to be positive, it then must be the case that the marginal utility of labor is negative for the last hour of work. Even if the individual enjoys his work generally, if he is paid then it must be the case that the last hour of his labor has a negative marginal utility. The same kind of result holds here, with emotions replacing labor in the equation.

Of course, the MRT between emotions and consumption need not be positive. Indeed, the signs of all components of (2.3) will depend on the type of problem we are investigating. This is one area where this framework becomes useful. If we are examining a certain problem and have prior beliefs about the sign of some of the derivatives, then the first order conditions allow us to pin down the sign of the unknown terms. That is, equation (2.3) can help us identify unknown features of the EPF or the marginal utility of an emotion.

2.4 *Application: Guilt*

To illustrate the usefulness of this approach, let's turn our attention to a particular emotion: *guilt*. In this section, I explore several ways that the emotion of guilt may factor into rational decision making. Let's start with a simple example. One finding from social psychology is that people often do not steal from others, even if there is no risk of getting caught. Why not? So long as the marginal utility from the stolen good is positive, if there are no repercussions from stealing, we ought to steal whenever the opportunity arises. Of course, the marginal utility of consumption is not the only consideration when it comes to stealing. Another is how stealing makes us feel. And for most people, that feeling is guilt. How can guilt inhibit theft, in the context of the economic model?

Consider a petty theft, such as stealing office supplies from your employer. Let x represent the amount you steal. If you steal a positive amount, but restrain yourself from stealing everything you can, then the first order conditions described above will apply to you. Suppose the marginal utility of consumption is positive, the marginal utility of guilt is negative, and that our consumption rises the more we steal (either through consuming the pilfered office supplies or selling them). Then, our rational framework implies that the last office supply you steal makes you feel guilty. That is, the EPF for guilt has a positive first derivative.

Table 3.1: Marginal Effects of Stealing

dC/dx	U_C	U_M	**(Implied) dM/dx**
+	+	-	+

Our everyday experience makes this result seem almost trivial: of course stealing can make us feel guilty, and of course we don't like guilt. Absent guilt, or some other factor making theft undesirable, every office will have been ripped-off to the studs in the wall. That the economic model with guilt predicts this so cleanly is an appealing feature of the approach.

Let's now turn about a somewhat more interesting example: penance. Many people who feel guilty about something take actions to try and alleviate that guilt. Examples of penitent behavior range from the sacred (Catholic traditions of confessing sins and repeated prayer) to the banal (washing of hands). This latter example, dubbed the "Macbeth Effect," has recently been uncovered by an experimental team from the University of Toronto (Zhong and Liljenquist, 2006). They find that individuals who were induced to recall previous acts of immoral behavior are more likely to wash their hands than those who were not so compelled. The feeling of guilt seems to trigger an urge to cleanse oneself physically.

In this case, we let x represent the amount of penance performed. Let's also continue to assume $U_C>0$, $U_M<0$, and suppose that it is true that penance does indeed reduce the intensity of guilty feelings (that is, our EPF for guilt is decreasing in penance ($dM/dx<0$)). This implies that, on the margin, it must be the case that consumption is lowered by penance. This too jibes with our everyday intuition. Whether the time is spent on one's knees in church or bathing 5 times a day, the act of penance necessarily implies the sacrifice of one's consumption. If it did not require such a sacrifice, then whenever we feel the slightest amount of guilt we would seek to alleviate it through penance at every possible. This type of behavior is not often observed.

Table 3.2: Marginal Effects of Penance

(Implied) dC/dx	U_C	U_M	dM/dx
-	**+**	**-**	**-**

An implication of this result is that, all else equal, those whose consumption is highly impacted by penitent acts (perhaps through foregone income) will be less penitent than those for whom the opportunity cost of their time is low. One way this might manifest itself is, all things equal, if those with a lot of unoccupied time (say, the elderly or unemployed) are more likely to engage in guilt-lessening behavior than other groups. There may be data that support this conjecture.

For a final case, consider charitable giving. I focus on charity, because thinking about altruism more generally can awaken a nest of philosophical problems. This is largely due to a philosophical dimension of the problem. Some would argue that any act that is truly altruistic must decrease the giver's utility, relative to other feasible actions. Otherwise, the act wouldn't be altruistic, but self-serving. Describing altruism this way means that by definition an altruistic agent cannot be utility maximizing.

Giving money to a charity unambiguously lowers one's consumption.[1] In the context of our simple model, this implies that it can only be optimal to donate if there is a compensating emotional return. What sorts of emotional benefits are to be had

---

[1] This assumes that the person does not receive sufficient economic benefits from giving. For instance, a billionaire might support a political candidate that would cut his taxes by an amount that exceeds the donation.

from charitable giving? One explanation for charitable giving is that induces a "warm glow" feeling. (See, for example, Andreoni 1990.) From the name "warm glow," we could guess that it is a desirable emotion. If charitable giving does indeed induce this emotion, then the first order conditions in fact demand this be true: it must be the case that $U_M > 0$.

Table 3.3: Marginal Effects of Charitable Giving (Warm Glow)

dC/dx	U_C	**(Implied) U_M**	dM/dx
-	+	+	+

Another possible channel may exist if our agent believes he has an obligation to be charitable. In this case, a positive amount of giving may be optimal, even if he receives no other rewards. This obligation may be a result of social pressures (say, there is a general consensus that the well-off should help support the needy, or if a political/religious organization to which one belongs encourages its members to contribute to certain causes), or a deep-seated sense of personal morality (say, to support a relative in financial distress). Regardless of the source of this sense of obligation, failure to live up to one's obligations may induce feelings of guilt. Charitable giving, to the extent that it satisfies the perceived obligation, lowers that guilt. In this case, our table of derivatives will look very similar to the case of penitent behavior.

Table 3.4: Marginal Effects of Charitable Giving (Guilt)

dC/dx	U_C	U_M	**(Implied) dM/dx**
-	+	-	-

3 Young's Theorem

Young's Theorem is a well known result from calculus that states that the mixed cross partial derivatives of a given function are equal to each other, regardless of the order in which the derivatives are taken. If we are interested in a function U(C,M), say, then by Young's Theorem it must be the case that $U_{CM}=U_{MC}$.[2] In words, the change in the marginal utility of *consumption* when an *emotion* changes is exactly equal to the change in the marginal utility of the *emotion* when *consumption* changes. This simple fact, when combined with Proposition 2, turns out to yield non-obvious – and in principle testable – implications for our framework of economic agents acting optimally over their emotions.

The way to apply this insight in practice is to first look for correlations between our "consumption" variable and a particular emotion. Via Proposition 2, that correlation will imply a sign for the mixed cross-partial derivatives of the utility function. Then, we can examine both orderings of the cross-partial derivative – by first treating one component of utility function as exogenous and the other as endogenous, and then try the reverse direction – and see if they both make sense. If the implications of the cross-partials are not supported by the evidence, this suggests that the agent cannot be acting in a utility maximizing way.

[2] Young's theorem is a very general result, and does not rely on any assumptions about continuity or differentiability.

What do I mean by “evidence” here? One type of evidence is just common experience. What this approach lacks, unfortunately, is rigor. Not only is it possible our interpretation of the events is incorrect, but in real life it can be difficult to separate the endogenous from the exogenous. What Young’s Theorem offers is a causal line of thinking. If such and such event is exogenous, then behavior adjusts thusly. Observations from everyday life are not so clean, as what is exogenous and what is endogenous can be difficult to determine.

However, in the controlled environment of a laboratory one can determine what is exogenous and what is endogenous. In principle all of the conjectures implied by Young’s Theorem can be addressed through experimentation. If, say, an experiment shows that triggering emotion A induces action B, a clever experimenter can reverse the stimuli. That is, she can “force” the subject to take action B, and see if the subject then induces emotion A. Successfully executing such an experiment can directly test the implications of Young’s Theorem.

Let’s look at some examples of this technique.

3.1 Cheating and Self-Esteem

In an experimental setting, Aronson and Mattee (1968) found that individuals who have had their self-esteem artificially lowered are more likely to cheat at a subsequent game of cards. Other studies have found that individuals with low levels of serotonin – a hormone which is strongly correlated with self-esteem – are more likely to commit impulsive crimes (Masters and McGuire, 1994). Robert Wright (1994) interprets this from the perspective of evolutionary psychology: those who have been shunned by society (or feel that they are of little value to society) begin to believe that they need to break the rules to get access to resources.

Translating this into our Emotions in the Utility Function setting is straightforward. If C is cheating, and M is self esteem, then the experimental results imply that $dC/dM<0$. Therefore, by Proposition 2, $U_{CM}=U_{MC}<0$. The first of these cross partials derivatives suggests that if an individual has high self-esteem, he gets less utility from cheating. This seems reasonable, as a genuinely confident person would seem to get little benefit from competing dishonestly in a contest, because he believes he will win anyway. On the other hand an opponent who feels inferior may believe he can only hope to win by breaking the rules, or at least be more willing to accept awards earned through questionable means.

We may be able to observe this tendency in a recent well-publicized case. In *Game of Shadows* (Fainaru-Wada and Williams, 2006), the writers tell the story of Barry Bonds – a tremendously successful baseball player. During the 1998 season, Bonds was overcome by feelings of jealousy with respect to the popular sluggers Mark McGwire and Sammy Sosa. Bonds had an intense desire to be considered the best player by others, and anything short of that enraged him. In order to catch up, he (allegedly) began to intensively use performance enhancing drugs. While Bonds, who already had established Hall of Fame credentials, would hardly seem a candidate for low self-esteem, the evidence does suggest that he was experiencing feelings of inferiority, and this apparently made cheating a more appealing option. This outcome seems consistent with the experimental findings cited above.

However, does the reverse implication suggested by Young's theorem also hold? It should be the case that if one cheats, the marginal utility of low self-esteem rises. So, if a person is obliged to cheat (or is already cheating), does he also tend to behave in ways that lower his self-esteem? This is a more difficult direction to unravel, as it is harder to identify if a person is acting in a way that may lower his self esteem than it is to identify cheating.

Let's continue with the case of Barry Bonds and see if there are hints of this tendency. When information on his steroid usage came to light, he did not publicly respond with self-righteous indignation. More commonly he prevailed on a woe-is-me response. In one televised interview, he famously broke down in tears "like a broken man" (Marriotti, 2006). "They can take me down. I don't really care. I never cared," said Bonds in a tearful tirade. "Baseball, if they want to take me down, go right ahead, take it. Anyone who ever knows me knows ... I don't care. But there are so many other people who depend on me to stay strong." One interpretation of Bonds's conduct is that through the act of portraying himself to the public as being helpless, and conceding defeat to his "enemies," he was engaging in a self-humbling behavior. In the context of Young's theorem, this type of self-esteem-lowering behavior becomes more desirable when an agent is cheating.

Of course, one does not want to generalize from the exceedingly complex – and difficult to interpret – case of a single person. I use the Barry Bonds example here only because it is a particularly vivid demonstration of the application of Young's theorem. The point is, if this interaction between dishonest behavior and self esteem is a result of the agent's making utility maximizing choices, Young's theorem provides clear predictions for *both* directions of the interaction. While one direction has been tested experimentally – that those with (exogenously) low self-esteem tend to be willing to behave dishonestly – the other has not.

Consider an experiment in which a subject is either induced to cheat, or otherwise unfairly wins a contest. Is the subject indeed more likely to lean towards behaviors that reduce his self esteem than those in a control group? This tendency might be manifest itself in a variety of ways, such as making self-deprecating statements, or self-identification with individuals of dubious moral character.

Another virtue of such an experiment is that, by gauging which "esteem-lowering" activity (if any) the cheating agent engages in, it may more precisely identify what "self-esteem" is referring to in this context. Is it that the person feels inferior relative to some internal personal standard? If so, perhaps this could be revealed in an experiment where the cheating agent gets to choose whether or not to write a brief essay that offers him the opportunity to make humble declarations. Or, perhaps self-esteem is more closely related to the image of being a social outsider. In that case, does he make decisions that try to imply some connection between himself and those perceived as villains? Consider, for example, offering the cheating subject the choice to watch clips from either the films The Godfather or Superman. Are cheaters more apt to choose the mobster movie?

3.2 Wanting to Hate

Now consider a soldier fighting in a war. The more he hates the men on the other side of the line, the less hesitation he will have when trying to kill them, and, presumably, the more effective a soldier he will be. He may lose utility when he kills, and the emotion of hating the enemy may also lower his utility, but the cold hard fact is the more he hates the enemy, the better he is at killing them.

If this account is accurate, then translating the result into our framework is simple: $dC/dM > 0$, where C in this case represents the number of enemy killed, and M is the emotion of hatred. By Proposition 2, then, it must be true that $U_{CM}>0$: the marginal utility of killing is greater if one feels more intense hatred. Even if killing the enemy lowers utility overall ($U_C<0$), by hating them more, the negative effect of killing on utility is reduced. This seems like a natural result. The more you hate someone, the less you mind if that person dies – even if you happen to be the instrument of death.

By Young's Theorem, however, $U_{MC}>0$ as well. So it is equally true that if a soldier is obligated to kill, then his marginal utility of hatred increases. That is, he will have an incentive to increase his hatred for the enemy if he has to kill them. The literature on internal conflict resolution suggests that this is a common practice. As one writer puts it, "[p]sychologically, it is necessary to categorize one's enemy as sub-human in order to legitimize increased violence or justify the violation of basic human rights" (Maiese, 2003).

We would thus expect soldiers in the field to tend to want to raise their hatred of the enemy. This could manifest itself in behaviors in which they actively seek out negative information about the enemy. They may consume negative propaganda, or listen to other soldiers who are expressing their own hatred. Even if a soldier resists the incentive to try to actively hate the enemy more, being in an environment where he might need he might need to kill the enemy might make him resist hateful attitudes less than he would ordinarily.

3.3 *Sex and Desire*

An article in The New York Times (4/10/07) recently noted that there is a burgeoning interest in academic circles regarding research on human sexuality. Much of this work is, in essence, an attempt to identify components of the EPF for sexual arousal and desire. The research suggests that the EPFs for sexual desire can be very complicated, but there has been progress in pinpointing particular sensations or images that induce sexual desire.

It is probably not controversial to suggest that sexual activity and feelings of desire are positively correlated. This indicates that $dC/dM>0$, where C is the frequency of sexual activity, and M is desire. By our usual arguments, this implies $U_{CM}=U_{MC}>0$. The first of these cross-partial derivatives seems obvious. Feeling stronger desire raises the marginal utility of sex, if for no other reason than that sex can directly address the desires. Feeling desire in essence raises the incentive for having sex.

The reverse implication is also a natural one, if perhaps less obvious. If one is participating in sexual activity, then the marginal utility of desire is higher. That is, if one is having sex, one wants to feel desire. For instance, a person might be obliged to have sex in order to satisfies one's partner, or for purely for procreative reasons. In this case, it is not uncommon for the less interested partner to try and induce lustful feelings, perhaps through fantasy or role play. Arousal-seeking behavior of this sort is a natural consequence of the neoclassical model.

Much of the demand for pornography might be generated by a similar process. When alone, a person might desire sexual release in order to relax, or to aid falling asleep. If solitary sexual activity is initiated for such non-emotional reasons, the cross-partials still indicate that the individual still has an incentive to increase his level of desire as part of the process. Use of pornography may serve to raise the level of desire.

4 Cognitive Dissonance

Cognitive dissonance is a term used by psychologists to describe situations in which an individual holds incompatible cognitions. Typically, one of these cognitions is the agent's own action or behavior, while another is a belief related to that action. Leon Festinger (1957) is credited with originating cognitive dissonance theory. The primary thesis of this theory is that individuals who hold incompatible cognitions experience discomfort, and they will seek to reduce that discomfort. In the years following the development of this theory, a number of experiments have yielded results supporting this thesis.

Cognitive Dissonance theory is potentially troublesome for economics. Perhaps the cleanest way to demonstrate this is through one of Aesop's fables, called The Fox and the Grapes. One afternoon a hungry fox notices a bunch of grapes dangling high up on a fence. He wants to eat them, but try as he might, he just can't jump high enough to reach them. After a while, he gives up and sulks away, muttering to himself "the grapes are probably sour, anyway." The fox's attitude towards the grapes suggests the moral "it is easy to despise what you cannot get," and this tale has thus spawned the term "sour grapes."

The cognitive dissonance interpretation of the fable is this: the cognition that the hungry fox is not eating the grapes is inconsistent with the cognition that the grapes are good. To eliminate the dissonance, there is pressure to change one of his cognitions. He is powerless to change the former, but he may be – and in the story apparently is – able to change the latter, his belief about the quality of the grapes.[3]

In economics, applications of the neoclassical model usually insist that agents form beliefs based on an objective appraisal of the information available – that is, that the beliefs are "rational." In this story, the fox's belief about the quality of the grapes has apparently changed – he wouldn't have tried to get the grapes in the first place if he initially believed they were sour. However, there is no new information about the quality of the grapes that has led to this new belief. This seems to be a non-rational belief.

Economists have tackled this problem in a variety of ways. Brunnermeier and Parker (2005) endorse the possibility that agents choose a set of "optimal beliefs." These optimal beliefs may differ from rational beliefs, but in some cases the hedonic benefits of holding optimistic – though non-rational – beliefs about the future outweigh the costs of holding those beliefs (through any non-optimal choices based on those beliefs). I should note that these beliefs are *not* an argument in the utility function. Instead, the agent's consumption plan is changed in such a way as to confer immediate benefits to utility. Bénebau and Tirole (2002) make a similar conjecture with respect to self-confidence. While they do not emphasize the hedonic virtues of holding a false belief about one's own abilities, they suggest that having high self-confidence can enhance an agent's ability to perform a given task.

Approaches to the problem of cognitive dissonance by Ackerlof and Dickens

[3] Refinements of the dissonance theory suggest that the tendency to reduce dissonance is stronger when both cognitions are, to some extent, entered into voluntarily, which is not the case for the fox.

(1982) and Rabin (1992) are more commensurate with the approach endorsed here. Both of these papers take the position that dissonance imposes a direct psychic cost in the individual's objective. In Ackerlof and Dicken's case, they argue that fear effectively lowers the net payment to an individual performing a dangerous job. Rabin proposes that dissonance lowers the value of doing an "immoral" action, such as eating veal. He argues that the amount of veal consumed and the belief about the morality of eating the veal are determined at the same time, so as to maximize utility.

A feature that is common to all these approaches is that they all feature non-rational beliefs. This is not necessarily that bad. For the consistency of the neoclassical approach, which is more important: for beliefs to be rational or for behavior to maximize utility? I would argue that the latter is more important. Conditional on the fox's inability to get the grapes, he can either believe the grapes are good – and experience the negative emotional consequences of dissonance – or believe the grapes are not good, and not experience dissonance. The latter would seem to be the option that maximizes utility, and it would certainly appear that this option indicates the formation of non-rational beliefs.

However, I suggest that even in light of cognitive dissonance we ought to not be quite so quick to abandon the notion of rational beliefs as our benchmark case. To this end, take a moment to think about how the fox would react if – after giving up on the grapes and concluding that they are sour – the grapes were to fall off the fence. Would he turn up his nose at them, since he now supposedly believes they are sour, or would he change his belief and eat them? Since his beliefs seem to be so malleable to begin with, I suspect he would now decide that perhaps the grapes might not be so bad after all. What does this thought experiment say about the whole notion of a belief more generally? Consider the following conjecture.

Suppose are two different classes of beliefs. One set does reflect an objective interpretation of evidence. Call these one's "true beliefs." Because these beliefs are dictated by rational deliberation, there is no room for choice about these beliefs. A second class of beliefs we can call "pseudo-beliefs." These are beliefs that can be modified to suit emotional needs. The pseudo-beliefs can override the true beliefs if the combination of possible cognitions (action and pseudo-belief) yield more utility than the other pair of possible cognitions (action and true belief). True beliefs, while they may be concealed by a pseudo-belief, manifest themselves whenever the agent must make a choice involving that belief. Such an interpretation of beliefs allows the fox to turn up his nose at the grapes when they are unreachable on the fence, but immediately return to feast when they blow off. His true belief is that the grapes are (probably) not sour, but when he couldn't eat them, he consoled himself by adopting a pseudo-belief that they (probably) were sour.

When viewing beliefs in such a framework, the formation of behaviors and beliefs are essentially broken into two parts. The first part is an action taken based on true beliefs, the second part is choosing a pseudo-belief that impacts emotions in such a way as to raise utility. For instance, a person might really like veal and decide to eat lots of it, despite any qualms about the morality of eating it. Conditional on his eating the veal, he will suffer emotionally if he thinks of himself as being immoral. It is therefore optimal to induce compensating emotions by justifying to himself (or others) that he is not immoral. That is, he adopts a pseudo-belief that eating veal is not so harmful after all. In the laboratory, when his beliefs seem to change in response to his behavior, it is called reducing cognitive dissonance. In the framework of Emotions the Utility Function, the dissonance reduction boils down to making a utility maximizing choice of a belief – with emotions serving as the conduit between the two. The strict neoclassical model continues to be relevant, because at a deep level he is aware of his true beliefs about eating veal.

Studies of galvanic skin response (GSR) – a measurement of the electrical conductivity of skin – show hints that such a multi-layered belief system is not purely conjecture. It has been found that when an individual hears a recording of her own voice, GSR levels rise more than when she hears the voice of others. Interestingly, researchers have observed that if a subject fails at some contrived task (lowering self-esteem), she is more likely to deny that a recording of her voice is her own. However, while a subject may deny the voice is hers, her GSR reading indicates that in fact she does, at least on some level, recognize the voice as her own. Amazingly, the experiment then shows that when self-esteem levels of a subject are artificially raised, the subject is more likely to identify *other* people's voices as her own – even though her GSR level continues to give an accurate identification. (Cited in Trivers, 1985.)

This result suggests that at a deep level, agents can indeed hold a "true belief" about the world even when espousing a "psuedo-belief," with the divergence between the two being triggered by emotional factors. The neoclassical approach to studying emotions outlined here suggests that when we observe as reduction in cognitive dissonance, we may interpret this as a utility maximizing manipulation of beliefs.

To illustrate this, consider a common example of what one might call cognitive dissonance: the belief in heaven. Over 89% of Americans claim to believe in heaven, and 80% of those believe they are themselves going to heaven (ABC News Poll, 10/2005). It should then be surprising that people feel sad when a good person dies, or if they themselves are diagnosed with a fatal disease. If good people all go to the eternal paradise of heaven, then death would seem to be a wonderful thing. Yet this is not how death is thought of in most circles.

The above interpretation of cognitive dissonance could explain this apparent contradiction. There is one fundamental cognition that everyone possesses: I exist. What does death represent? The end of my existence. The belief in heaven may seek to eliminate the dissonance that these mutually exclusive ideas generate, by replacing the cognition that existence ends with the more emotionally satisfying cognition that existence continues in paradise.

As there is no verifiable evidence that heaven exists, the belief in heaven is unlikely to be a true belief as I have defined it. That heaven is actually a pseudo-belief is revealed when an emotion arousing event – such as a friend's death, or a terminal diagnosis – triggers those emotions that depend in part on the true belief. This interpretation suggests that people adopt the psuedo-belief in heaven because the dissonance resulting from holding the true belief may be too painful. That is, the model of this paper suggests that adopting a pseudo-belief may be viewed as the result of a utility maximizing strategy – where we take seriously the role of Emotions in the Utility Function..

5 Conclusion

This essay has been an attempt to justify the usage of strictly neoclassical economic models that include emotions as benchmarks for micro-level behaviors. The first objective was to show that such models – in which agents seek the optimal choice and emotions matter for behavior – necessarily imply that the emotions must be represented as an argument of the utility function. That is, for a neoclassical model, one need not assume that agents optimize *and* emotions are in the utility function when studying emotionally-linked actions. Only the former is required. This is a subtle point that has not been addressed in the literature. Being able to reduce the number of assumptions when evaluating such problems is a non-trivial realization.

Given that the conventional tools of the economics trade are most useful in the context of agents who are maximizing some objective, it would seem that it is in the interests of economists to push the neoclassical model with Emotions in the Utility Function to its limits. In addition to straightforward undergraduate-level price-theory, I have suggested that extensions of the model yield hypotheses for behavior in a variety of contexts that seem consistent with common experience. The application of Young's Theorem is particularly useful in this respect, as it can often suggest non-obvious implications that both seem reasonable and are in principle testable in the lab – a very useful feature. The model also provides the structure for a purely neoclassical interpretation of the phenomenon known as cognitive dissonance.

What this essay does not do is try to explain all emotionally-linked behavior in terms of constrained optimization. Such an attempt would likely fail, as emotions, particularly intense emotions, may well corrupt one's ability to make an optimal choice. However, I suggest that abandoning the standard economic framework should be the researcher's second option. Rather, one should try hard to evaluate behaviors in light of optimal choice. What this essay has attempted to convey is that carefully scrutinizing the role of Emotions in the Utility Function is an important first step in this process.

Figure 3.1 Hypothetical Responses of Consumption to a Change in Emotion

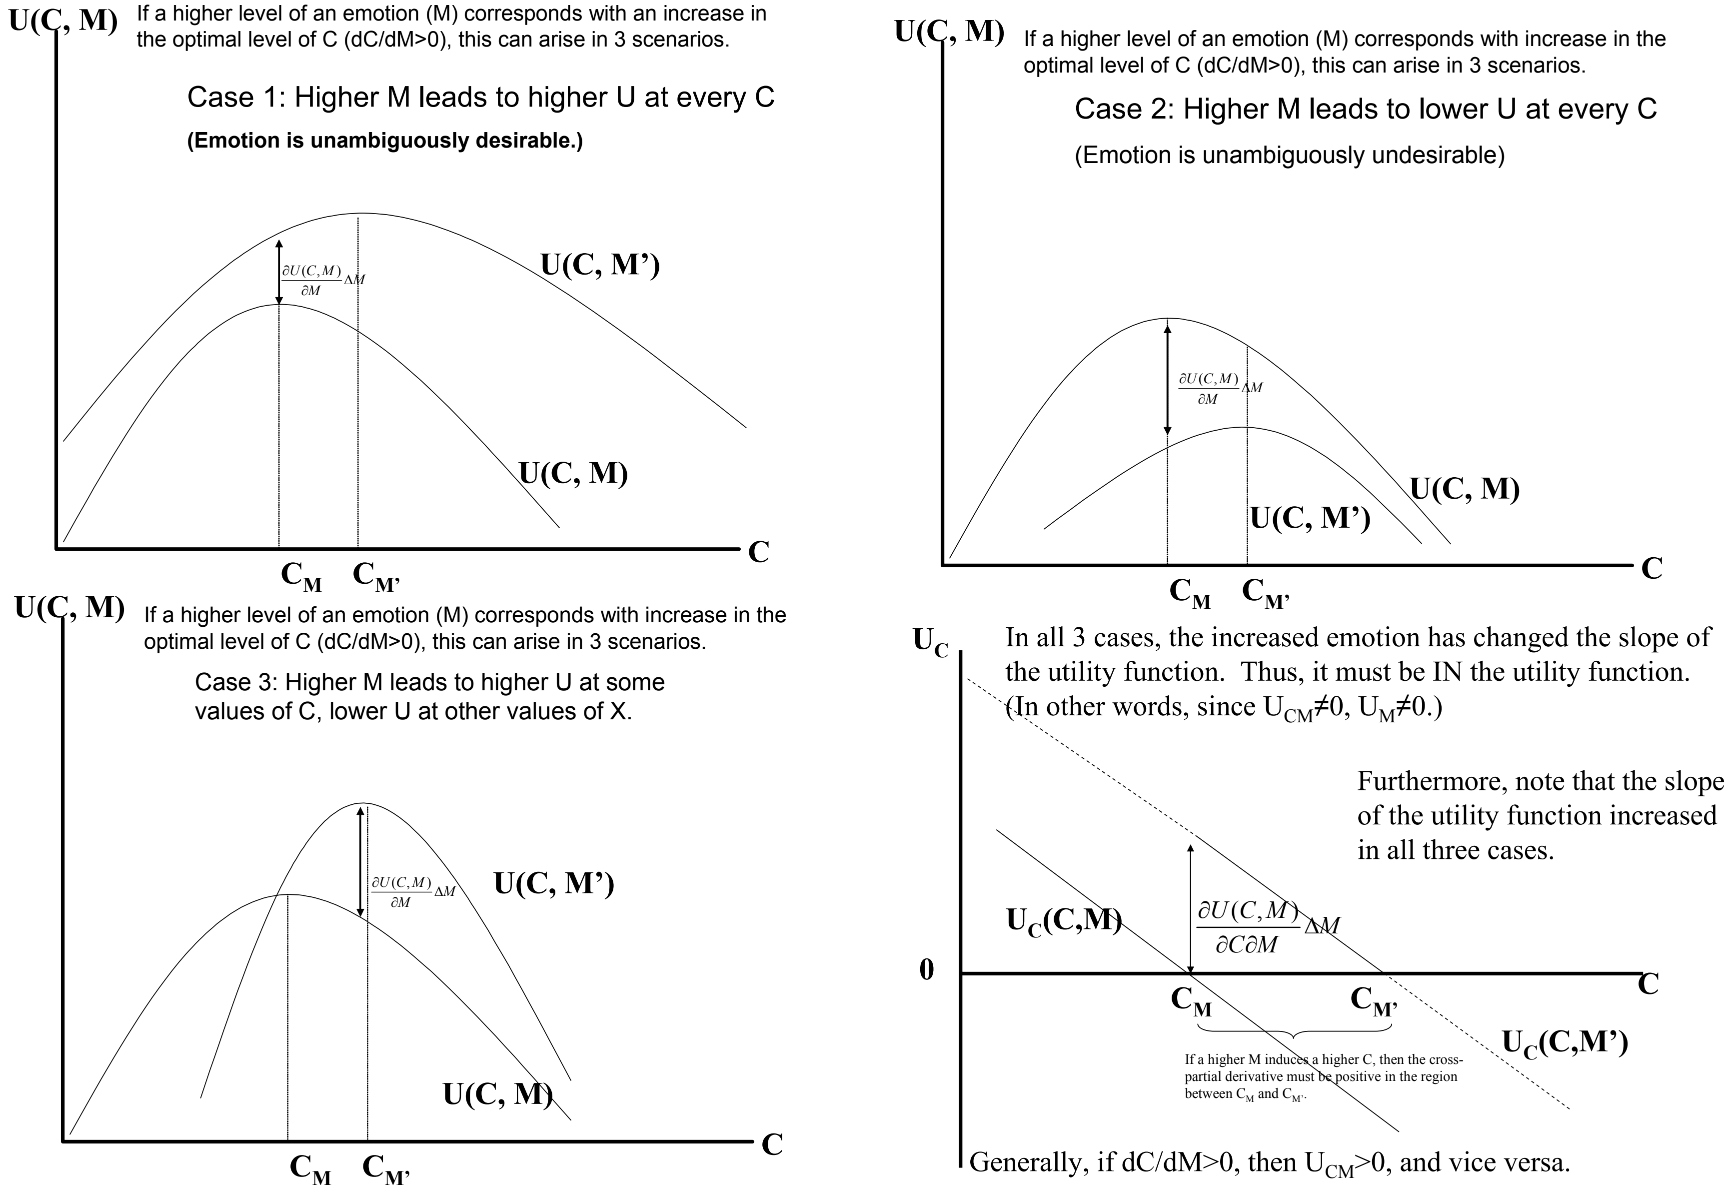

References

Ackerloff, George, and William T. Dickens. 1982. The Economic Consequences of Cognitive Dissonance. *The American Economic Review*. Volume 72, No. 3: 307-319.

Andreoni, James. 1990. Impure Altruism and Donations to Public Goods: A Theory of Warm Glow Giving. *The Economic Journal*. Volume 100: 464-477.

Angier, Natalie. 2007. Birds Do It. Bees Do It. People Seek the Keys To It. *The New York Times*. 4/10/2007.

Aronson, Elliot, and David R. Mettee. 1968. Dishonest Behavior as a Function of Differential Levels of Induced Self-Esteem. *Journal of Personality and Social Psychology*. Volume 9: 121-127.

Bénabou, Roland, and Jean Tirole. 2002. Self-Confidence and Personal Motivation. *The Quarterly Journal of Economics.* August 2002: 871-915.

Becker, Gary. 1976. *The Economic Approach to Human Behavior*. Chicago: U. of Chicago Press.

Becker, Gary, and Kevin M. Murphy. 1993. A Simple Theory of Advertising as a Good or Bad. *The Quarterly Journal of Economics.* November 1993: 941-964.

Brunnermeier, M, and Jonathan Parker. 2005. Optimal Expectations. *The American Economic Review*. Volume 95, No. 4: 1092-1118.

Camerer, Colin F, George Loewenstein, and Matthew Rabin, eds. 2004. *Advances in Behavioral Economics*. Russell Sage Foundation.

Damasio, Anthony. 1994. *Descartes Error: Emotion, Reason, and the Human Brain.* G.P. Putnam.

Elster, Jon. 1998. Emotions and Economic Theory. *Journal of Economic Literature.* Volume 36, No. 1: 47-74.

Fainaru-Wada, Mark, and Lance Williams. 2006. *Game of Shadows: Barry Bonds, BALCO, and the Steroids Scandal that Rocked Professional Sports*. Gotham.

Festinger, Leon. 1957. *A Theory of Cognitive Dissonance.* Evanston, IL: Row, Peterson.

Gilad, Benjamin, Stanley Kaish, and Peter D. Loeb. 1987. Cognitive Dissonance and Utility Maximization: A General Framework. *Journal of Economic Behavior and Organization*. Volume 8: 61-73.

Gul, Faruk, and Wolfgang Pesendorfer. 2001. Temptation and Self Control. *Econometrica.* Volume 69, No 6: 1403-1435.

Kimball, Miles, and Robert Willis. 2006. Utility and Happiness. Working Paper.

Laibson, David. 2001. A Cue-Theory of Consumption. *The Quarterly Journal of Economics.* February 2001: 81-119.

Loewenstein, George. 1996. Out of Control: Visceral Influences on Decision Making. *Organizational Behavior and Human Decision Processes*. Volume 65.

Maiese, Michelle. 2003. What it Means to Dehumanize. *Beyond Intractability*. Eds. Guy Burgess and Heidi Burgess. July 2003. Conflict Research Consortium, University of Colorado, Boulder, Colorado, USA. <http://www.beyondintractability.org/essay/dehumanization/>

Ozdenoren, Emre, Stephen Salant, and Daniel Silverman. 2006. Willpower and the Optimal Control of Visceral Urges. Working Paper.

Rabin, Matthew. 1994. Cognitive Dissonance and Social Change. *Journal of Economic Behavior and Organization.* Volume 23: 177-194.

______________ 2002. A Perspective on Psychology and Economics. UC Berkeley Working Paper E02-313.

Wright, Robert. 1994. *The Moral Animal: Why We Are the Way We Are, The New Science of Evolutionary Psychology.* Vintage Books.

Zhong, C. B., and Liljenquist, K. Washing Away Your Sins: Threatened Morality and Physical Cleansing. *Science*. Vol. 313, September 8, 2006: 1451-1452