

Identification and Inference Using Event Studies*

Refet S. Gürkaynak[†] Jonathan H. Wright[‡]

First Draft: September 2012

This version: March 9, 2013

Abstract

We discuss the use of event studies in macroeconomics and finance, arguing that many important macro-finance questions can only be answered using event studies with high-frequency financial market data. We provide a broad picture of the use of event studies, along with their limitations. As examples, we study financial markets' responses to specific events that help address questions such as the slope of bond demand functions and the efficacy of central bank liquidity programs. We also study the change in financial market responses to news in payrolls and unemployment in response to former Fed Chairman Greenspan's statement that payrolls are more informative.

JEL Classification: E43, E52, E58, G12, G14.

Keywords: Event Study, Bond Markets, High-Frequency Data, Identification, TAF

*Prepared for the 2012 Money, Macro, and Finance conference in Dublin. We are grateful to conference participants and Rob Shimer for very helpful comments. All errors are our sole responsibility.

[†]Department of Economics, Bilkent University, 06800 Ankara, Turkey; refet@bilkent.edu.tr.

[‡]Department of Economics, Johns Hopkins University, Baltimore MD 21218; wrightj@jhu.edu.

1 Introduction

Applied economists have a harder job than natural scientists because natural scientists get to do controlled experiments. Controlled experiments can give clear-cut evidence on, for example, the benefit of giving a patient a particular treatment rather than a placebo. Alas, we can seldom do anything like this in economics. But there are some examples where fate gives economists a pseudo-natural experiment. The controlled release of economic news is an important example. In the minutes before the news is released, its content is unknown to financial market participants. After the news comes out, with efficient markets, the information content is quickly impounded into asset prices (French and Roll (1986)). The jump in asset prices in a small window around the announcement reflects the causal impact of the news, and likely little else. Hence, the lumpy manner in which news is released to the public is a source of identification, as discussed more formally in Faust et al. (2007). This identification approach is known as the event study methodology and is the focus of this paper.

An event study looks at the reaction of the financial markets to news. This is of course a task that journalists undertake every day; the event study methodology does this systematically. It requires two key ingredients. First the researcher must be able to measure the “news” — the part of the event that changes the information set of market participants. Then, the researcher has to be able to measure the effect of this news on a relevant asset price. The task is best undertaken with very high-frequency data, because in a small enough window around the news announcement, nothing other than the announcement should be affecting asset prices. This is what makes it resemble a controlled experiment. With longer windows, other shocks will matter as well.

In this paper we provide a review of the event study methodology and give some new examples that shed light on important topical questions such as the

effects of central bank liquidity programs.

2 Event Studies: An Overview

There are two reasons why we use event studies. The first reason is because we are interested in the financial markets' reactions in their own right. An example of this is the study of the effectiveness of central banks' large scale asset purchases and other liquidity provisioning measures during and after the recent financial crisis. The financial markets' reaction to these announcements can be used to measure their effects. The second reason why we use event studies has to do with shortcomings of the available macroeconomic data and econometric techniques.

By and large macroeconomics is in a state of observational equivalence, where a great number of disparate models fit a small number of macroeconomic stylized facts roughly equally well. Taylor and Wieland (2012) show that models with very different microfoundations and policy implications fit the current canon of macroeconomic stylized facts about as well as each other. Hence macroeconomic facts and standard macroeconomic data sets cannot be used to distinguish between the models.

The literature has moved in two directions to overcome this problem; both involve looking at more data. The first direction is to use microeconomic data to make statements about macroeconomic models. This is why we see many recent studies using individual household and firm decision data in assessing macroeconomic models. The second direction, which is the focus of the current paper, is to look at financial markets both because financial markets offer a very rich set of data and because by the nature of financial markets these data are very amenable to manipulation in ways that allow backing out market participants' expectations about the future.

A simple example here is whether inflation targeting works. Inflation targeting is supposed to anchor inflation expectations. Inflation itself will still be subject to shocks and will deviate from the target. Nonetheless, with an inflation target, the central bank ought to ensure that such deviations are temporary. It is hard to answer the question of whether inflation targeting worked in the countries that adopted this regime because one needs a long time to see whether the nature of the inflation process has indeed changed. If inflation targeting, for example, is to bring inflation down to its target level within three years, one would need 30 years of data to have just 10 independent observations. Financial market data, on the other hand, provide an almost continuous reading of inflation-related expectations at various maturities in the future, and these can be analyzed in event studies to see how news about inflation today affects expectations of inflation in the distant future.

When the object of interest is the behavior of financial markets, a high-frequency event study is the natural tool of analysis as this method allows for fairly straightforward identification. A quarterly VAR may be sufficient to study the relationship between GDP or inflation and monetary policy. But if another financial variable (say, stock prices) is added to the mix, there exists no ordering of the variables that could reasonably identify the monetary policy effect on the asset price. As both monetary policy and asset prices are forward-looking, they react to each other within the quarter. In contrast, assuming that monetary policy does not react to financial market developments on the *day* of the monetary policy announcement, which is very reasonable, allows us to give a causal interpretation to any correlation that can be observed between policy actions and asset price reactions on the days of monetary policy announcements. This is the essence of the event study methodology.

3 Methodology

The modeling setup is a system of two simultaneous structural equations that relate an asset price and a macro variable or policy action to each other:

$$\Delta a_t = \alpha \Delta x_t + z_t + \varepsilon_t \quad (1)$$

$$\Delta x_t = \beta \Delta a_t + \gamma z_t + \eta_t \quad (2)$$

where Δa_t the change in the asset price, Δx_t is the news (the surprise in the data release or policy action), z_t is a vector of other variables affecting both the asset price and the news and ε_t and η_t are uncorrelated error terms. The parameter of interest is α but it is not identified because both variables depend on each other.

Here is where high-frequency data can help. When one looks at a short enough window, say a day or even an hour around an event of interest, it becomes easy to argue that the variance of z_t and ε_t are small relative to the variance of the shock to the news, η_t . Then, equation (1), can simply be estimated by an OLS regression of Δa_t on Δx_t . The equation is estimated over windows that include an announcement. Estimation of equation (1) may include a constant, but it is typically not significant, and sometimes dropped. White (1980) standard errors are generally used to allow for heteroskedasticity, but no allowance for serial correlation is made, as asset returns in successive announcement windows exhibit minimal autocorrelation.

The argument for estimating equation (1) by OLS is actually quite a bit stronger, however. The news that comes out is typically something that by construction can only depend on *lagged* asset returns. For example, the FOMC decision is made late on the morning of the meeting day and then released at 2:15. No asset price movement that afternoon could ever change the decision

(nor could it change expectations made as of before the meeting). Similarly, the employment report that comes out on the first Friday of the month pertains to the labor market in the previous month. No asset price shocks on the day of the release can ever affect it. So in many contexts, assuming that there is no other shock before the announcement but within the event study window affecting the agents' expectations about the news (an assumption that will almost certainly hold in short windows), we can replace equation (2) with:

$$\Delta x_t = \beta \Delta a_{t-j} + \gamma z_{t-j} + \eta_t$$

where j denotes some lag length that depends on the particular context. Then, equation (1) can safely be estimated by OLS, even if the variances of z_t and ε_t are not necessarily small relative to the variance of the shock to the news.

3.1 Identification through Heteroskedasticity

Sometimes, though, one may wish to estimate equations (1) and (2) and the researcher may feel that there is some risk of endogeneity bias from estimating (1) alone by OLS. In this case, there is an elegant generalization of the event study approach considered by Rigobon (2003) and Rigobon and Sack (2003, 2004, 2005) which depends on the observation that the variance of ε_t is the same at all times but the variance of the news is higher on news release times compared to non-news windows. The heteroskedasticity caused by the news timing can then be used to separate the effect of news on asset prices from usual background noise which affects both variables. This is called heteroskedasticity-based identification. Rigobon and Sack have used this to study the effect of monetary policy on asset prices and since then many studies have used the same idea in many different contexts including Ehrmann, Fratzscher and Rigobon (2011) on cross-country and cross-market interlinkages and Duran et al. (2012)

on financial market responses to monetary policy in an emerging market. A useful feature of this approach is that it does not rely on actually measuring the news—it is much easier to know *when* the variance of the news is higher than to actually be able to measure the news.

News releases are long and contain a great deal of information, including sectoral detail and in some cases revisions to previously released data. The headline surprise will not be the only news to come out of the announcement, hence methods not directly measuring the surprise may be unable to differentiate between asset price responses to different components of the news. Gürkaynak, Kısacıkoglu, and Wright (2012) discuss this issue, and propose a remedy, which is related to identification through heteroskedasticity. We do not go into this methodological debate in this paper but rather talk about the economics to be learned from the use of event studies regardless of whether they are carried out via OLS or via heteroskedasticity-based identification.

3.2 Caveats and Limitations

For every question to be answered using an event study, there are several crucial caveats and limitations:

1. The data must come out in a controlled manner and at a precise time. If, for example, the news leaks ahead of time, then the event study methodology cannot be used.
2. We must have good measures of the expectations. Failing to condition on the news, that is, not isolating the expected component of any release, will lead to an attenuation bias in econometric work as forward-looking financial markets will have responded to the expected part already. It is therefore important that the news capture a change in the information set of market participants. That is why the literature often uses the phrase the

“surprise” component of news or data releases—a surprise by definition captures a change in the information set.

3. We can only measure the effects of the news on asset prices (and on expectations that are embedded in those asset prices) that are available at high frequency.
4. Event studies provide information on market participants’ expectations, but not on actual outcomes. For example, using an event study we may learn whether market participants *think* that a newly announced policy will help bring inflation to a certain level but not whether the policy will *actually* deliver that outcome.
5. Finally, event studies provide information on market participants’ beliefs under risk neutrality.¹ News could affect risk premia as well—the effects on expectations and risk premia cannot readily be disentangled.

Notwithstanding these caveats, there are many cases where the structure of news announcements allow us to measure casual effects of news on financial market participants’ beliefs, assuming rational expectations and unchanging risk pricing .

3.3 The Window Size

Generally, the window around the data release that researchers use should be as small as possible. Indeed, the widespread availability of intraday data has been central in recent empirical work on event studies, as the evidence suggests that the jump in conditional mean following a news announcement is done within

¹In fact, the condition is much weaker. Event studies work well under risk aversion as long as risk pricing does not change in response to the news release. As the asset price in (1) is differenced, the levels of any risk premia are differenced out. Changes in premia not systematically related to the surprise will be in ε_t .

about 10 minutes (Andersen et al. (2003)). There are two reasons why it is generally better to use small windows to estimate the reactions to news:

1. It reduces the danger of any feedback in the estimation of equation (1).
2. Even if the window is small enough that there is no concern about endogeneity in a regression of Δa_t on Δx_t , a smaller window will reduce the error variance while leaving the variance of the right-hand-side variable unchanged. Improving the signal-to-noise ratio will then make the coefficient more precisely estimated.

In assessing the effects of news announcements on major markets, it is now typical to use roughly 20 minute windows bracketing release times (five minutes before the release to 15 minutes after). Any gains to using smaller windows are likely to be slight, and there are potential downsides to making the windows smaller still. In some markets, illiquid trading or even lack of data availability might motivate using larger windows. Bauer (2012), for example, argues that TIPS incorporate news somewhat more slowly than nominal bonds and hence the measurement window must be large enough to capture TIPS updating.

4 News and Markets for Event Studies

In this section we provide a short and very selective survey of some of the key questions addressed using event studies and the news and markets used to do so.

The event study literature has its roots in the seminal works of Royle (1982), Cutler, Poterba and Summers (1989), Ederington and Lee (1993), Fleming and Remolona (1997, 1999) and Kuttner (2001). The literature generally uses survey expectations of data releases such as inflation, unemployment and the GDP

growth rate together with the actual released values of these variables to calculate surprises on the days of the releases. One can obtain a long time series of survey expectations for data releases from sources such as Action Economics (formerly Money Market Services) and Bloomberg. These survey expectations are surely not perfect measures of expectations (they are at least somewhat stale), but nonetheless have good properties as expectations proxies. Balduzzi, Elton and Green (2001) find that they pass simple forecast efficiency tests and outperform simple benchmarks. While survey expectations of the target federal funds rate to be announced at the next FOMC meeting exist, it is generally thought to be preferable to measure these expectations from federal funds futures contracts instead (Kuttner (2001)).

These measures of surprises have been used to answer many different questions in the literature. Kuttner (2001) found that monetary policy has sizable effects on longer-term interest rates. Andersen et al. (2003) examined reactions of exchange rates to macroeconomic news and found significant relationships, which was a major breakthrough as exchange rates had previously been seen as unrelated to any macroeconomic fundamentals: past, present or future (Meese and Rogoff (1983)). Bernanke and Kuttner (2005) examined the effects of monetary policy shocks on stock prices. Boyd, Jagannathan and Hu (2005) found that higher-than-expected unemployment rate announcements cause stock prices to rise during expansions, but not during recessions. Unemployment news has two offsetting effects: on future dividends and on discount rates. The latter seems to be dominant most of the time, but not in recessions. Gürkaynak, Sack and Swanson (2005) looked at the long forward interest rates' response to data releases and monetary policy surprises to see whether inflation expectations are anchored. With anchored inflation expectations, surprises about the state of the economy today should be uninformative about inflation at steady-state, say ten

years' hence. That means that ten year forward interest rates should be unresponsive to macroeconomic developments today. However, in the US, far-ahead forward interest rates are responsive to data surprises, in ways consistent with market participants not knowing the inflation target and updating their beliefs about steady-state inflation based on what they learn today.

The introduction of inflation-indexed bonds has been a boon for macroeconomic research in general, and for the event study approach in particular. These markets, which existed in the UK since the early 1980s, began in the US in 1997. The difference between nominal and indexed bonds' interest rates, called inflation compensation or breakeven inflation, measures inflation expectations under risk neutrality. One can construct an inflation compensation term structure and look at the responses of the long end of that term structure directly, which gives a cleaner read on long-term inflation expectations.

Gürkaynak, Levin and Swanson (2010) do this for the US, UK, and Sweden and find that far-forward inflation compensation is unresponsive to data releases in Sweden and also in the UK after the independence of the Bank of England, but does respond to data releases in the US and also in the UK before central bank independence. This suggests that a credible inflation target does indeed succeed in anchoring long-run inflation expectations. Beechey et al. (2011) do a similar exercise for the euro area using inflation swaps instead of inflation compensation. Gelati et al. (2011) study whether inflation expectation responsiveness to data has changed as concerns about fiscal sustainability have mounted during the recent crisis, and find that it has not.

Another application of event studies is with political data. Snowberg et al. (2012) provide interesting examples of event studies involving political events, including the stock market reaction to changes in perceptions about who will win on an election night. A particularly striking political event study is the

work of Rigobon and Sack (2005), showing that increases in the probability of US entering war against Afghanistan in early 2003 led to lower stock prices and higher oil futures quotes. This paper also contains a nice application of heteroskedasticity-based identification. La Ferrara and Guidolin (2007), again using event studies—this time from Angola—show that war may lead to higher stock prices if there are entry barriers due to the conflict.

4.1 Quantitative Easing Event Studies

In the aftermath of the financial crisis, with short-term interest rates stuck at the effective zero lower bound, the Fed and the Bank of England have undertaken large scale asset purchases, commonly referred to as quantitative easing. The European Central Bank has also engaged in bond purchases, although within a quite different framework. The usefulness of these actions is a subject of vigorous debate. Low-frequency data will not help because of endogeneity issues. The recovery in most countries has been weak, yet that is clearly not evidence against the usefulness of unconventional policy. However, the event study methodology can be used to identify the effects of these and other unconventional policy actions on asset prices, which is at least an important part of the transmission mechanism to the broader economy.

Authors using the event study framework to examine the effects of quantitative easing announcements include Doh (2010), D’Amico and King (2012), D’Amico et al. (2012), Gagnon et al. (2011), Neely (2010), Krishnamurthy and Vissing-Jorgenson (2011), Joyce and Tong (2012), and Wright (2012). Also, Swanson (2011) reexamined Operation Twist from the 1960s using a modern event study perspective, and compared it to the unconventional monetary policies presently being employed by the Federal Reserve. This body of work collectively establishes very convincingly that unconventional monetary policy an-

nouncements had material, albeit perhaps small, effects on asset prices.

The initial announcements of purchases of government bonds by the Fed and the Bank of England, both in March 2009, were complete surprises, which is very helpful from the event study perspective. They were followed by large drops in bond yields.

Since August 2010, the FOMC has generally signaled its intentions about quantitative easing plans to markets before the actual meeting announcements. This makes the event study methodology harder to use, because the methodology relies on the information coming out in a lumpy manner. But there was a modest surprise at the September 2012 meeting. At this meeting, the FOMC announced more purchases of mortgage backed securities (MBS), but not of Treasuries, whereas the expectation had apparently been from some mix of the two. Following the announcement, MBS yields fell while Treasury yields rose.

D'Amico and King (2012) and D'Amico et al. (2012) carefully study the cross-section of announcement effects on different bond yields (at the CUSIP level) with the objective of shedding light on the channels whereby quantitative easing announcements affect yields. One possibility is the duration channel: purchases of long-term Treasuries reduces the duration of the stock of bonds in the market, and drive up the price of long-duration assets. Another is the local supply channel: markets are segmented and investors have demand for Treasury securities of a particular maturity.² The FOMC announcement of August 10, 2010 was particularly useful in disentangling these two channels. At 2:15pm, the FOMC announced that principal payments from agency securities

²Another possibility is the signaling channel, whereby asset purchase announcements lead agents to conclude that the federal funds rate will remain at the zero lower bound for longer. Federal Reserve announcements of large-scale asset purchases had large effects on ten-year-ahead forward rates, and we find it difficult to imagine that the announcements had material effects on expectations of the federal funds rate at that horizon. On the other hand, the signaling channel is likely important when analyzing the effects of forward rate guidance, such as expressing the intention to keep the federal funds rate at zero through late 2014 (announced at the January 2012 FOMC meeting).

would be reinvested in Treasury securities. At 2:45pm, the New York Fed gave a breakdown of the maturities of the Treasury securities that would be purchased. These were to be concentrated in the 2-10 year sector. This staggered timing of the revelation of information may be unusual central banking practice, but it was very helpful to academic research as information about duration and local supply (which securities would be targeted) came out separately, allowing two event studies to identify both effects. D'Amico et al. (2012) found that both the duration and local supply channels were present, but that the latter seemed to be quantitatively more important.

The bottom line from event study analysis of the impact of quantitative easing on asset prices is that there is a remarkable degree of consensus. All the papers agree that quantitative easing announcements mattered, and they agree on the sign and the even the order of magnitude of the effect. That's quite rare in macroeconomics but is consistent with other event study evidence on the effects of changing supply on yields. However, from an asset pricing theory point of view, the effects are surprising. Large as these asset purchases are, they are a trivial share of global fixed income markets, and conventional finance theory would hold that their effects on interest rates should be negligible. A significant practical contribution of the event study methodology is that it has demonstrated that this view is unambiguously wrong, which is an important contribution to our empirical understanding of financial markets. At the same time, an important limitation of the event study approach to measuring the effects of quantitative easing is that it is harder to know what effects the announcements had on real economic activity.

An important question in quantitative easing event study analysis (and indeed in event studies more broadly) is that it is hard to know the persistence of the effects. On the one hand, the reactions that are documented in these event

studies are to *announcements* of future purchases, not to the actual purchases themselves, and so they cannot be dismissed as transitory microstructure effects. On the other hand simply totting up the announcement effects of news about asset purchases—as though they were all permanent—gives big effects. We are a little suspicious of this since in a preferred habitat model, there are natural mechanisms that might cause the effects to wear off over time. One is that quantitative easing may make the economy stronger than would otherwise have been the case, driving interest rates back up. A second is that quantitative easing may induce more corporate issuance of long-term bonds (Stein (2012)), which would in turn also tend to push long-term rates higher. A third is that frictions may make arbitrage capital slow-moving—and thus better able to offset the impact of preferred habitat investors in the long-run than in the short-run (Mitchell, Pederson and Pulvino (2007), Duffie (2010)). Wright (2012) uses a VAR methodology to estimate the persistence of quantitative easing effects and obtains wide confidence intervals, but his point estimates indicate that the effects wear off quite fast.

4.2 More on Measuring Monetary Policy Surprises

Kuttner (2001) showed how to measure the surprise components of FOMC announcements concerning the target federal funds rate. But even before the Fed reached the zero lower bound, the announcement about the target federal funds rate was not the only important feature of FOMC announcements—the statement mattered as well. And of course since December 2008, there have been no surprises to the target federal funds rate. Only the statement and quantitative easing announcements have mattered and quantifying these is no easy task.

The literature has gone in two directions in this regard. One is to look at policymakers' statements and speeches to count references to certain key

phrases to help quantify the general tone of communication. This may suffer from some attenuation bias as this method does not distinguish between expected and unexpected communication and therefore provides a lower bound for the effects of communication. Ehrmann and Fratzscher (2003, 2005) have carried out such studies for the ECB and have shown that ECB communication has important bearing on financial markets. Lucca and Trebbi (2009) use tools from computational linguistics to score FOMC statements. Ehrmann and Fratzscher (2009) studied the ECB president's post-meeting press conference. The ECB first releases its policy decision and after a short while its president holds a press conference first reading a prepared set of comments and then having a Q&A session. This allows separating the effects of three different kinds of communication (policy decision, statement, and Q&A).

The second way of quantifying central-bank communication is to resort to using asset prices. Gürkaynak, Sack and Swanson (2005b) use federal funds futures contracts for later months as well as the spot month, and euro dollar futures contracts that cover futures rates over the next year and come up with two policy surprise measures, one for the surprise in the policy action (the target funds rate) and the other for the surprise in the announced policy path using a combination of these different assets. They find that while the policy action surprise influences short maturity interest rates, long maturity interest rates are affected almost exclusively by the path surprise, showing the importance of communication as a policy tool. Work along similar lines has been done by Brand et al. (2010) for the euro area and Aktaş et al. (2009) for Turkey. Wright (2012) does a similar exercise to measure quantitative easing announcement surprises, but uses longer-term Treasury futures.

5 Three New Examples

We now turn to some more illustrative examples of the practical power of event studies. The first is the study of a single event.

On June 7, 2012, the Swedish regulator allowed pension funds to increase the discount rates they use to calculate their liabilities. The very low yield environment brought on by the crisis forced pension funds to increase their holdings of long maturity government securities held for hedging purposes. The regulator, by allowing the use of higher-than-market discount rates, reduced the long maturity government bond demand of these pension funds. The announcements, then, clearly led to an understanding that pension funds would be selling off large quantities of long maturity Swedish government debt. The question is whether this would have a price impact. This is of course closely related to the important question of the effects of large scale asset purchases by central banks, as discussed above, although it is a little different in that it cannot possibly signal anything about future monetary policy.

The expected answer depends on the slope of the demand curve for Swedish government bonds. With downward sloping bond demand there would be a price impact, but with a horizontal demand curve (as would be the case under conventional asset pricing models) there would be no impact. Figure 1, which plots Swedish government bond yields around the time of the announcement, clearly shows a jump in the 10- and 30-year yields of over 30 basis points and also the unresponsive one-year yield underscoring that there was not a simultaneous monetary policy change. The study of this event therefore shows that yields of long maturity Swedish securities—precisely those securities that are affected by the accounting rule change—increased significantly in response to the perceived likely sale of these assets by pension funds. This is only consistent with a downward sloping demand curve, and therefore provides strong evidence

in favor of portfolio constraints or segmented markets. It is striking given that Sweden has highly developed and open capital markets. Findings of this sort are naturally important for issuers of debt, and also for regulators who consider mandating that pension funds or others must hold certain government bonds. Of course, the findings are at least as important for academic research that tries to understand bond pricing.

The second example relates to a program that the Federal Reserve instituted during the financial crisis, known as the Term Auction Facility (TAF). This was in effect discount window credit that was auctioned (and called by a different name to avoid the stigma associated with discount window borrowing). Taylor and Williams (2012) discuss the impact of this program on the LIBOR-OIS spread—the spread between interbank borrowing rates (LIBOR) and the overnight indexed swap (OIS) rate, which is a bet on future short-term interest rates but, unlike LIBOR, with no principal changing hands. The LIBOR-OIS spread is thought of as a measure of stress in money markets. Taylor and Williams found that the TAF if anything increased the LIBOR-OIS spread. But here there is a huge problem of endogeneity. TAF was not introduced as an exogenous policy decision, but rather in response to turmoil in financial markets, which was driving up spreads. There is thus a two-way causality issue, which the event study methodology can disentangle.

In Table 1, we list announcements pertaining to the size of the TAF program and their effects on the Eurodollar-federal funds futures spread from 15 minutes before the announcement to 1 hour and 45 minutes afterwards. We use futures quotes to have high-frequency data. Eurodollar futures settle to LIBOR, whereas federal funds futures are bets on the future level of the federal funds rate (just like OIS rates). So the eurodollar to federal funds futures spread is a close analog of the LIBOR-OIS spread, but for which we have intradaily data.

Table 1 also lists our judgment about whether this announcement represented an increase or decrease in the size of the TAF program, relative to expectations. Although the effects are small in most cases, it seems striking that in *every single instance*, positive (negative) news leads to a narrowing (widening) of the Eurodollar over federal funds futures spreads. Similar exercises were undertaken by McAndrews, Sarkar and Wang (2008), and Wu (2011) using daily data on LIBOR-OIS spreads. We believe that the results in Table 1 make a compelling case that while TAF was introduced in response to financial market stress, it actually modestly alleviated that stress relative to the counterfactual in which TAF had not been introduced.³

The third exercise we carry out has to do with the determination of which releases financial market participants pay attention to. In particular we are interested in learning whether financial markets pay attention to data releases that they think are the most informative about the underlying state of the economy, or to data releases that policymakers pay attention to. To find out we turn to the systematic study of financial market reactions to the release of the employment report.

Table 2 reports regressions of the intraday change in the fourth Eurodollar futures rate and ten-year Treasury yields onto the employment report’s two main components—the surprises in nonfarm payrolls and in the unemployment rate. Both surprises are divided by their standard deviations so that the financial market responses are easily comparable across the two kinds of surprises. The intraday changes are from 8:25 to 8:45, which brackets the release of the news at 8:30 (Eastern time). The regression is run using all employment reports from January 1985 to September 2012. The fourth Eurodollar futures rate is

³Notice that the measurement of the “surprise” is particularly problematic in this exercise as it depends on our judgment of whether the announced TAF amounts were larger or smaller than what markets expected. We were careful not to look at market responses before making this determination to avoid reverse causality by the econometrician’s hand but our perceptions may nonetheless be different from those of market participants’.

a bet on three-month interest rates in about one years' time, and thus Table 2 considers two points on the term structure. The results suggest that short maturity bond yields are affected by both nonfarm payrolls and unemployment, but that market participants react more to payrolls.

During the sample used in the regressions reported in Table 2, the then Chairman of the Fed, Alan Greenspan, unequivocally stated in a congressional hearing that the Fed views the nonfarm payrolls numbers as far more informative than the unemployment numbers:

I wish I could say the household survey were the more accurate. Everything we've looked at suggests that it's the payroll data which are the series which you have to follow (February 2004).

The two right columns of Table 2 address the question of what difference this statement made, by interacting the surprises with a dummy that takes on the value 1 after the Greenspan statement. The coefficients on the dummy variable interacted with the payrolls surprise are positive and significant for both the Eurodollar futures rate and the Treasury yields. This means that the Greenspan statement made the markets more sensitive to payrolls surprises. For example, the effect of a one standard deviation payrolls surprise on ten-year yields is estimated to have risen from 1.5 to 2.75 basis points.⁴

The evidence in Table 2 suggests that the Greenspan statement led markets to put more weight on the payrolls numbers. There are two interpretations as to why. One is that they were convinced by Greenspan's statement that payrolls data are the less noisy signal. The other is that they want to respond to the series that the Fed is looking at, simply as a means of predicting future policy actions. Either interpretation is possible, though the latter may seem more

⁴The estimated coefficient on the dummy interacted with the payrolls surprise in the Eurodollar futures equation would be even bigger if the sample were ended a couple of years earlier. This is because the zero lower bound has reduced the sensitivity of short- and intermediate-term interest rates to *all* macroeconomic news (Swanson and Williams (2012)).

plausible. The point estimates in Table 2 also indicate that the Greenspan statement led markets to put *less* weight on unemployment numbers, but this difference is not statistically significant.


These three examples give a taste of what can be learned from event studies, be them one time events or systematic ones. Each one hopefully shows the importance of isolating the right event and looking at the right markets to answer the question of interest.

6 Conclusions

Empirical work in macroeconomics and finance faces huge difficulties in disentangling causality. The lumpy manner in which news is released is perhaps the most useful tool that exists for identifying causal effects in the macro-finance literature. The event study methodology helps make use of this data. In this paper, we have reviewed the event study methodology and provided several examples of its power.

Conditioning on the proper event and looking at the right market allows us to gain precious insights about financial market participants' beliefs. This is as close as we get to controlled experiments in macroeconomics.

Figure 1: Swedish Sovereign Yields Around the June 7, 2012 Announcement



Notes: This Figure shows the Swedish government bond yield curve (Bloomberg generic securities) on June 6 and June 7, 2012. These are the daily closing quotes before and after the announcement that Swedish pension funds could increase the discount rates they use to calculate their liabilities, reducing their special demand for long duration Swedish government bonds.

Table 1: Effects of Term Auction Facility (TAF) Announcements on Eurodollar/Federal Funds Futures Spread

Date	Announcement	Sign of News	Intraday Change in Spread (bps)
12/11/2007	No TAF release after FOMC meeting	Negative	+17.5
12/12/2007	Announcement of inception of TAF	Positive	-11.0
12/21/2007	TAF to be retained as long as needed	Positive	-1.5
1/4/2008	TAF to be expanded	Positive	-1.5
3/7/2008	TAF to be expanded	Positive	-6.5
7/30/2008	TAF to be expanded	Positive	-5.0
9/29/2008	TAF to be expanded	Positive	-11.5

Notes: At the December 11, 2007 FOMC meeting, the Federal Reserve Board of Governors decided to introduce the TAF. However, they delayed the announcement until the next morning instead of including it with the FOMC statement released at 2:15. Some market participants had expected TAF to be announced with the FOMC statement, and incorrectly inferred from the absence of any such announcement that the TAF would not go ahead at that time. This announcement is therefore treated as negative news. The other announcements are treated as positive news. The column on the right shows the change in the Eurodollar-federal funds spread (in basis points) from 15 minutes before the event to 1 hour and 45 minutes afterwards. The front Eurodollar futures contract is compared with the average federal funds futures rates for the two months subsequent to the Eurodollar expiration. For example, a December Eurodollar futures quote is compared to the average of the January and February federal funds futures rates. In cases where the front Eurodollar contract is going to settle within the current month, we skip to the next contract instead.

Table 2: Regressions of Changes in Treasury Yields on Payrolls and Unemployment Surprises

	ED4	Ten-Year Yield	ED4	Ten-Year Yield
Payrolls	6.67*** (0.52)	1.75*** (0.15)	6.08*** (0.53)	1.50*** (0.15)
Unemployment	-2.00*** (0.40)	-0.50*** (0.11)	-2.32*** (0.51)	-0.60*** (0.14)
Payrolls*Dummy			3.33** (1.56)	1.25*** (0.41)
Unemployment*Dummy			0.62 (0.80)	0.12 (0.22)
R^2	0.498	0.465	0.513	0.497

Notes: This Table shows coefficient estimates in regressions of intraday changes from 8:25 to 8:45 in the fourth eurodollar futures contract and ten-year Treasury futures yields on the days of employment report announcements onto the surprise components of the payrolls and unemployment releases. The employment report comes out at 8:30. Both surprises are scaled by their standard deviations. Changes in futures yields are obtained by dividing futures returns by the duration of the cheapest-to-deliver security in the futures delivery basket. No constant is included in the regression. White standard errors are shown in parentheses. The sample period is January 1985 to September 2012. The Table also reports the results of regressions with interaction effects with a dummy that is 1 in March 2004 and later (and 0 before). This amounts to allowing the coefficients of the regression to exhibit a break before March 2004, which was the first employment report release following the Greenspan remarks on the relative signal of payrolls and unemployment numbers, discussed in the text.

References

- [1] Aktaş, Zelal, Harun Alp, Refet Gürkaynak, Mehtap Kesriyeli, and Musa Orak, (2009): Türkiye’de Para Politikasının Aktarımı: Para Politikasının Mali Piyasalara Etkisi, *Iktisat İşletme ve Finans*, 24, pp. 9-24.
- [2] Andersen, Torben, Tim Bollerslev, Francis X. Diebold and Clara Vega (2007): Real-Time Price Discovery in Stock, Bond and Foreign Exchange Markets, *Journal of International Economics*, 73, pp.251-277.
- [3] Balduzzi, P., E.J. Elton and T.C. Green (1999): Economic News and Bond Prices: Evidence from the U.S. Treasury Market, *Journal of Financial and Quantitative Analysis*, 36, pp.523-543
- [4] Bauer, Michael. (2012): Inflation Expectations and the News. Federal Reserve Bank of San Francisco working paper.
- [5] Beechey, Meredith J., Benjamin K. Johansson and Andrew T. Levin (2011): Are Long-Run Inflation Expectations Anchored More Firmly in the Euro Area Than in the United States?, *American Economic Journal: Macroeconomics*, 3, pp.104-129.
- [6] Bernanke Ben S. and Kenneth N. Kuttner (2005): What Explains the Stock Market’s Reaction to Federal Reserve Policy?, *Journal of Finance*, 60, pp.1221-1257.
- [7] Boyd, John H., Jian Hu and Ravi Jagannathan (2005): The Stock Market’s Reaction to Unemployment News: Why Bad News Is Usually Good for Stocks, *Journal of Finance*, 60, pp.649-672.
- [8] Brand, Claus, Daniel Buncic, and Jarkko Turunen. (2010): The Impact of ECB Monetary Policy Decisions and Communication on the Yield Curve, *Journal of the European Economic Association*, 8, pp. 1266-1298.
- [9] Cutler, David M., James M. Poterba and Lawrence H. Summers (1989): What Moves Stock Prices?, *Journal of Portfolio Management*, 15, pp.4-12.
- [10] D’Amico, Stefania, William English, David López-Salido and Edward Nelson (2012): The Federal Reserve’s Large-Scale Asset Purchase Programmes: Rationale and Effects, *Economic Journal*, 122, pp.F415–F446.
- [11] D’Amico, Stefania and Thomas B. King (2012): Flow and Stock Effects of Large-Scale Treasury Purchases: Evidence on the Importance of Local Supply, *Journal of Financial Economics*, forthcoming.
- [12] Doh, Taeyoung (2010): The Efficacy of Large-Scale Asset Purchases at the Zero Lower Bound, *Federal Reserve Bank of Kansas City Economic Review*, 2, pp.5-34.

- [13] Duran, Murat, Gülserim Özcan, Pınar Özlü and Deren Ünalmış (2012): Measuring the Impact of Monetary Policy on Asset Prices in Turkey, *Economics Letters*, 114, pp. 29-31.
- [14] Duffie, Darrell (2010): Asset Price Dynamics with Slow-Moving Capital, *Journal of Finance*, 65, pp.1238-1268.
- [15] Ederington, Louis H. and Jae H. Lee (1993): How Markets Process Information: News Releases and Volatility, *Journal of Finance*, 48, pp.1161-1191.
- [16] Ehrmann, Michael and Marcel Fratzscher (2003): Monetary Policy Announcements and Money Markets: A Transatlantic Perspective, *International Finance*, 6, pp.1-20.
- [17] Ehrmann, Michael and Marcel Fratzscher (2005): Equal Size, Equal Role? Interest Rate Interdependence Between the Euro Area and the United States, *Economic Journal*, 115, pp.930-950.
- [18] Ehrmann, Michael and Marcel Fratzscher (2009): Explaining Monetary Policy Decisions in a Press Conference, *International Journal of Central Banking*, 5, pp.41-84.
- [19] Ehrmann, Michael, Marcel Fratzscher and Roberto Rigobon (2011): Stocks, Bonds, Money Markets and Exchange Rates: Measuring International Financial Transmission, *Journal of Applied Econometrics*, 26, pp.948-974.
- [20] Faust, Jon, John H. Rogers, Shing-Yi B. Wang and Jonathan H. Wright (2007): The High-Frequency Response of Exchange Rates and Interest Rates to Macroeconomic Announcements, *Journal of Monetary Economics*, 54, pp.1051-1068.
- [21] Fleming, Michael J. and Eli M. Remolona (1997): What Moves the Bond Market?, *Federal Reserve Bank of New York Economic Policy Review*, 3, pp.31-50.
- [22] Fleming, Michael J. and Eli M. Remolona (1999): Price Formation and Liquidity in the U.S. Treasury Market: The Response to Public Information, *Journal of Finance*, 54, pp.1901-1915.
- [23] French, Kenneth and Richard Roll (1986): Stock Return Variances: The Arrival of Information and the Reaction of Traders, *Journal of Financial Economics*, 17, pp.5-26.
- [24] Gagnon, Joseph, Matthew Raskin, Julie Remache and Brian Sack (2011): Large-Scale Asset Purchases by the Federal Reserve: Did They Work?, *International Journal of Central Banking*, 7, pp.3-44.
- [25] Galati, Gabriele, John Lewis, Steven Poelhekke and Chen Zhou (2011): Have Market Views on the Sustainability of Fiscal Burdens Influenced Monetary Authorities' credibility?, DNB working paper 304.

- [26] Gürkaynak, Refet S., Burcin Kisacikoglu and Jonathan H. Wright (2012): Identifying Effects of Partially-Measured News Surprises, working paper.
- [27] Gürkaynak, Refet S., Andrew Levin and Eric T. Swanson (2010): Does Inflation Targeting Anchor Long-Run Inflation Expectations? Evidence from Long-Term Bond Yields in the U.S., U.K., and Sweden, *Journal of the European Economic Association*, 8, pp.1208–1242.
- [28] Gürkaynak, Refet S., Brian Sack and Eric T. Swanson (2005a): The Excess Sensitivity of Long Term Interest Rates: Evidence and Implications for Macroeconomic Models, *American Economic Review*, 95, pp.425-436.
- [29] Gürkaynak, Refet S., Brian Sack and Eric T. Swanson (2005b): Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements, *International Journal of Central Banking*, 1, pp.55-93.
- [30] Gürkaynak, Refet S., Brian Sack and Jonathan H. Wright (2007): The U.S. Treasury Yield Curve: 1961 to the Present, *Journal of Monetary Economics*, 54, pp.2291-2304.
- [31] Gürkaynak, Refet S., Brian Sack and Jonathan H. Wright (2010): The TIPS Yield Curve and Inflation Compensation, *American Economics Journal: Macroeconomics*, 2, pp.70-92.
- [32] Gürkaynak, Refet S. and Justin Wolfers (2006): Macroeconomic Derivatives: An Initial Analysis of Market-Based Macro Forecasts, Uncertainty and Risk, NBER International Seminar on Macroeconomics.
- [33] Joyce, Michael A.S., Ana Lasaosa, Ibrahim Stevens and Matthew Tong (2010): The Financial Market Impact of Quantitative Easing, Bank of England working paper 393.
- [34] Joyce, Michael A.S. and Matthew Tong (2012): QE and the Gilt Market: a Disaggregated Analysis, *Economic Journal*, 122, pp.F348–F384.
- [35] Krishnamurthy, Arvind and Annette Vissing-Jorgenson (2011): The Effects of Quantitative Easing on Long-term Interest Rates, *Brookings Papers on Economic Activity*, 2, pp.215-265.
- [36] Kuttner, Kenneth N. (2001): Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market, *Journal of Monetary Economics*, 47, pp.523-544.
- [37] La Ferrara, Eliana and Massimo Guidolin (2007): Diamonds Are Forever, Wars Are Not. Is Conflict Bad for Private Firms? *American Economic Review*, 97(5), pp.1978-93.
- [38] Lucca, David O. and Francesco Trebbi (2009): Measuring Central Bank Communication: An Automated Approach with Application to FOMC Statements, NBER working paper 15367.

- [39] McAndrews, James, Asani Sarkar and Zhenyu Wang (2008): The Effect of the Term Auction Facility on the London Inter-Bank Offered Rate, Federal Reserve Bank of New York Staff Report 335.
- [40] Meese, Richard A. and Kenneth Rogoff (1983): Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample?, *Journal of International Economics*, 14, pp.3-24.
- [41] Mitchell, Mark, Lasse H. Pedersen and Todd Pulvino (2007): Slow Moving Capital, *American Economic Review*, 97, pp.215-20.
- [42] Neely, Christopher (2010): The Large-Scale Asset Purchases Had Large International Effects, Federal Reserve Bank of Saint Louis working paper 2010-18A.
- [43] Rigobon, Roberto (2003): Identification Through Heteroskedasticity, *Review of Economics and Statistics*, 85, pp.777-92.
- [44] Rigobon, Roberto and Brian Sack (2003): Measuring the Response of Monetary Policy to the Stock Market, *Quarterly Journal of Economics*, 118, pp.639-669.
- [45] Rigobon, Roberto and Brian Sack (2004): The Impact of Monetary Policy on Asset Prices, *Journal of Monetary Economics*, 51, pp.1553-1575.
- [46] Rigobon, Roberto and Brian Sack (2005): The Effects of War Risk on US Financial Markets, *Journal of Banking and Finance*, 29, pp.1769-1789.
- [47] Roley, V. Vance (1982): The Effect of Federal Debt-Management Policy On Corporate Bond and Equity Yields, *Quarterly Journal of Economics*, 97, pp.645-668.
- [48] Snowberg, Erik, Justin Wolfers, and Eric Zitzewitz (2012): Prediction Markets for Economic Forecasting, NBER working paper 18222.
- [49] Stein, Jeremy C. (2012): Evaluating Large-Scale Asset Purchases, speech delivered at the Brookings Institution, Washington, October 11 2012.
- [50] Swanson, Eric T. (2011): Let's Twist Again: A High-Frequency event study Analysis of Operation Twist and Its Implications for QE2, *Brookings Papers on Economic Activity*, 1, pp.151-188.
- [51] Swanson, Eric T. and John C. Williams (2012): Measuring the Effect of the Zero Lower Bound on Medium- and Longer-Term Interest Rates, Federal Reserve Bank of San Francisco working paper 2012-12.
- [52] Taylor, John B. and Volker Wieland (2012): Surprising Comparative Properties of Monetary Models: Results from a New Monetary Model Base, *Review of Economics and Statistics*, 94, pp.800-816.

- [53] Taylor, John B. and John C. Williams (2009): A Black Swan in the Money Market, *American Economic Journal: Macroeconomics*, 1, pp.58-83.
- [54] White, Halbert (1980): A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity, *Econometrica*, 48, pp.817–838.
- [55] Wright, Jonathan H. (2012): What Does Monetary Policy do to Long-Term Interest Rates at the Zero Lower Bound?, *Economic Journal*, 122, pp.F447-F466.
- [56] Wu, Tao. (2011): The U.S. Money Market and the Term Auction Facility in the Financial Crisis of 2007—2009, *The Review of Economics and Statistics*, 93, pp. 617-631.