Jules H. van Binsbergen, William F. Diamond and M Grotteria
Risk-Free Interest Rates
Article

van Binsbergen, Jules H., Diamond, William F. and Grotteria, M

(2021)

Risk-Free Interest Rates.

Journal of Financial Economics, 143 (1). pp. 1-29. ISSN 0304-405X

DOI: https://doi.org/10.1016/j.jfineco.2021.06.012

Elsevier
https://www.sciencedirect.com/science/article/pii/...

Risk-Free Interest Rates

Jules H. van Binsbergen William F. Diamond Marco Grotteria*

February 9, 2021

Abstract

We estimate risk-free interest rates unaffected by convenience yields on safe assets. We infer them from risky asset prices without relying on any specific model of risk. We obtain interest rates and implied convenience yields with maturities up to 3 years at a minutely frequency. Our estimated convenience yield on Treasuries equals about 40 basis points, is larger below 3 months maturity, and quadruples during the financial crisis. In high-frequency event studies, conventional and unconventional monetary stimulus reduces our rates more than the corresponding Treasury yields, thus broadly impacting rates even outside the narrow confines of the fixed income market.

*We thank Jonathan Berk, Doug Diamond, Greg Duffee, Darrell Duffie, Olivier Darmouni, Peter DeMarzo, Wenxin Du, Vincent Glode, Ben Hébert, Bob Hodrick, Arvind Krishnamurthy, Eben Lazarus, Zhongjin Lu, Ian Martin, Toby Moskowitz, Christian Opp, Michael Roberts, Tano Santos, Krista Schwarz, Ken Singleton, Adrien Verdelhan, Annette Vissing-Jorgensen, Jonathan Wright and two anonymous referees for helpful discussions and/or comments. We thank the Rodney White Center and the Jacobs Levy Equity Management Center for generous financial support. The authors are at the Wharton School of the University of Pennsylvania and London Business School, and can be reached at julesv@wharton.upenn.edu, diamondw@wharton.upenn.edu and mgrotteria@london.edu.

1. Introduction

One of the most important variables in economics is the interest rate on a risk-free investment. It is commonly used as a measure of the time value of money: the required return for receiving a riskless payoff in the future instead of the present. To measure investors' willingness to take risk, the returns on risky assets are compared to the risk-free interest rate, where the difference in average returns is conventionally interpreted as the compensation for bearing an asset's risk, i.e., the asset's risk premium. As a consequence, any attempt to measure either risk premia or the time value money requires a benchmark for the risk-free interest rate.

Empirically, the yield or interest rate on safe assets (such as government bonds) are often used for this benchmark. However, a recent literature has provided evidence that the interest rates on safe assets are driven in part by other forces (Krishnamurthy and Vissing-Jorgensen (2012), Nagel (2016)). Safe assets provide a so-called "convenience yield" that reflects the ease with which they can be traded by uninformed agents, posted as collateral, satisfy regulatory capital requirements, or perform other roles similar to that of money.[1] Thus, the yield on a money-like asset is below the risk-free cost of capital, reflecting the liquidity and collateral value of such assets.

In this paper, we estimate risk-free rates that are unaffected by the convenience yield on safe assets by inferring them from the prices of risky assets. Our empirical measurement is motivated by the fact that in several recent asset pricing models with frictions, risky assets do not earn a convenience yield while safe assets do.[2] As a result, under the assumptions of these models, a risk-free rate inferred from risky asset prices is a pure measure of the risk-free cost of capital. The spread between this inferred rate and the observed yield on a safe asset therefore measures that safe asset's convenience yield for agents able to invest in both rates.[3]

We infer our benchmark rates (which we call Box rates) from the put-call parity relationship for European-style options and show robustness using other implied risk-free rates, such as those inferred from a storage arbitrage in the market for precious metals futures

[1]See Gorton and Pennacchi (1990), Bansal and Coleman (1996), Lagos, Rocheteau, and Wright (2017) and Hazelkorn, Moskowitz, and Vasudevan (2020) among several others. Most generally, a convenience yield is simply the rate of return an investor is willing to forgo in order to hold a more "convenient" asset instead of a less convenient asset with identical cash flows, regardless of the underlying mechanism.

[2]E.g. Frazzini and Pedersen (2014), Stein (2012), Caballero and Farhi (2018), and Diamond (2018).

[3]If all agents trade actively in all markets, this convenience yield will be the same for all investors. On the other hand, if markets are segmented, our rates and convenience yield estimates may only apply to a group of specialized investors such as large financial intermediaries.

and stock futures. We find that the convenience yield on government bonds equals about 40 basis points on average over the sample period 2004-2018, with a relatively flat average term structure across maturities beyond 3 months, and 65 basis points below 3 months. Our estimated convenience yield grows substantialy in a financial crisis, reaching an average of 100 basis points for maturites beyond 3 months and 160 basis points below 3 months in 2008.

Our estimated rates have several advantages over other convenience yield measures. First, our rates are inferred only from risky asset prices, so they have no convenience yield even though their payoff is riskless. Previous work either computes spreads between the yields of two safe assets, which identifies only the *difference* in their convenience yields, or uses the yield of a bond exposed to credit risk, so that a credit spread is included in the convenience yield estimate. Our 40 basis point convenience yield is larger than estimates from this first approach and smaller than estimates from the second, suggesting that we have a reasonable estimate of the true size of the convenience yield on safe assets. Second, we provide minute-by-minute estimates and show that our rate responds faster than Treasuries to new information. Our convience yield measure is therefore the first that is suitable for high frequency event studies, which are the main empirical technique for analyzing how monetary policy impacts asset markets. Third, we obtain an entire term structre of convenience yield estimates with maturities ranging from 1 month to 3 years. This allows us to document that the convenience yield is mostly flat along the term structure and to show the important role convenience yields play in the predictability of bond returns.[4]

One potential concern with our estimated rates is that distortions specific to the option market could impact our option-implied interest rates. We document for our benchmark estimates (using S&P500 index (SPX) options) that observable measures of frictions (such as bid-ask spreads) in the option market are not related to our estimated rates. We repeat our analysis using the substantially less liquid Dow Jones (DJX) index options and find rates that are the same on average but considerably noisier. When there is a spread between the SPX and DJX implied rate, we find that the DJX rate converges towards the SPX rate and not conversely. Further, bid-ask spread measures in both markets (which exhibit substantial independent variation) have minimal predictive power for the difference in the implied rates across the two markets.

While we have precisely estimated candidate convenience-yield-free interest rates, it is possible that our rate measures the risk-free cost of capital (or funding cost) for expert

[4]Our rate has several smaller advantages: 1. There is no margin requirement to lend at our rate. 2. It is collateralized and therefore effectively free of credit risk. 3. Our rate comes from quotes on a centralized exchange, while others are inferred from over-the-counter trades or surveys of market participants which are less suitable for high frequency analysis.

financial intermediaries alone, since our rate is inferred from a sophisticated asset class. That said, our results extend beyond the specific setting of equity option markets, since we also estimate similar average (but noisier) rates using trade data for precious metal futures and index futures. While the extent to which different investor groups have different risk-free rates is an important open question for future research, it is not a question that our data can shed much further light on.

We use our new data set in three applications, which specifically require us to observe a term structure of high frequency interest rates (and convenience yields). First, we do an event study of the effects of quantitative easing and monetary policy on both convenience yields and the risk-free cost of capital.[5] We find that quantitative easing (Q.E.) and accommodating monetary policy reduce convenience yields, particularly during the depths of the financial crisis. This implies that our estimated interest rates are even more sensitive than government bond yields to Q.E and monetary policy. Quantitative easing is the purchase of long term Treasury bonds and agency mortgage-backed securities financed by the issuance of bank reserves, which are a form of overnight debt. An important question about its transmission mechanism is whether its effects spill over beyond the prices of debt securities that are actually purchased.

Because our interest rates are inferred from equity options and not from fixed income securities targeted by quantitative easing, they are ideal for testing whether Q.E. impacts the yields of non-targeted assets. Under the "narrow" view of Q.E.'s transmission mechanism, Quantitative Easing should not spill over broadly into the private sector's discount rates despite the lowering of long-term Treasury or MBS yields (the asset that is bought). A "broad" view of the transmission mechanism of quantitative easing on the other hand is common in much of the theoretical literature (Caballero and Farhi (2018)) and Diamond (2018)) which emphasizes that swapping reserves for risky or long duration assets increases the overall supply of safe/liquid assets and should therefore reduce convenience yields across markets. We find that our risk-free rates are more sensitive to quantitative easing than the associated Treasury yields, implying that quantitative easing reduces the scarcity of safe assets as implied by the theoretical literature. This is particularly true in the depths of the financial crisis, and convenience yields do not seem to respond to conventional monetary policy outside the crisis.

Our second application shows that convenenience yields play an important role in bond re-

[5]Our high frequency data is particularly useful since much modern empirical work (Bernanke and Kuttner (2005), Rigobon and Sack (2004), Nakamura and Steinsson (2018)) on monetary transmission uses high frequency event studies, and we can now study convenience yields using this methodology.

turn predictability.We find that a forecasting factor constructed solely from the cross-section of convenience yields in the spirit of Cochrane and Piazzesi (2005) has forecasting power for both government bond excess returns as well as convenience-yield-free excess returns even when controlling for factors in the literature. In univariate regressions, we outperform the predictive power of conventional predictors in our sample. The results suggest that a full explanation of the predictability in bond excess returns includes time variation in the convenenience yield as well as other sources of predictability (e.g. time-varying risk aversion, volatility, or expectation errors).

The third application studies our option-implied rates across different currencies. The so-called covered interest parity relationship implies that risk-free interest rates in different currencies are related to the ratio of spot to future exchange rates between the currencies. Existing measures of violations of covered interest parity (CIP)[6] use interest rates that feature convenience yields (such as government bond rates) or credit risk (such as LIBOR) which raises the question to what extent such violations persist once convenience-yield-free risk-free interest rates are used. Using option data from Japan we construct an index-option-implied interest rate for that region as well, and find that when using option-implied rates for both countries, previously documented covered interest parity violations are reduced.

Our paper contributes to several related literatures. First, it contributes to the empirical literature mentioned above on safe assets by providing a convenience yield measure that is motivated by and connected to theory. Some existing convenience yield proxies are spreads between yields on two different safe assets, which may be an underestimate if both assets have positive convenience yields. Other proxies in the literature are spreads between a safe asset and a low-risk asset, which may be an overestimate if there is a nontrivial credit risk premium on the risky asset. Our computed spreads are larger than spreads in the first category and smaller than those in the second category. We also find that our spread is almost identical to the LIBOR-Treasury spread before the crisis but substantially smaller after. This is consistent with the view that credit risk in LIBOR was considered negligible before the crisis but significant afterwards. Similarly, we estimate a somewhat smaller convenience yield than the 73 basis points reported in the seminal paper of Krishnamurthy and Vissing-Jorgensen (2012) using a AAA-Treasury spread, perhaps because of some credit risk in AAA bonds. We also contribute to this literature by providing an entire term structure of convenience yields at a minute-level frequency.

Second, we contribute to the literature on monetary policy and quantitative easing event studies. The baseline event study on quantitative easing (Krishnamurthy and Vissing-

[6]See Du, Tepper, and Verdelhan (2018b) and Du, Im, and Schreger (2018a).

Jorgensen (2011)) presents spreads between different yields on safe assets but is constrained to a 2-day event window by slow price discovery, while we are able to use estimates within an hour of all event time stamps. Further, no existing work on monetary transmission studies risk-free rates inferred from assets outside the fixed income market, so our data is ideal for testing how broadly quantitative easing spills over to distant asset classes. In addition, existing high frequency event studies on conventional monetary policy have not examined the response of convenience yields, perhaps due to similar data limitations that we overcome.

Our work also relates to the literature on intermediary asset pricing, particularly the subset of the literature relating arbitrage spreads to financial frictions. He and Krishnamurthy (2013) presents a canonical intermediary asset pricing model, showing theoretically and quantitatively under what assumptions the capitalization of financial intermediaries is a key state variable for the dynamics of asset prices. A related theoretical and empirical paper by Frazzini and Pedersen (2014) presents a model in which the spread between the return on a zero-beta security and the risk-free rate measures the tightness of leverage constraints for levered investors and shows that this zero-beta rate is very high in a large range of asset classes. Measuring a zero-beta rate requires taking a stance on the specific risk factor that the beta is computed against. If the factor does not capture all risks relevant to investors, the zero beta rate includes a risk premium component. In contrast, the risk-free rate we estimate from options markets does not require specifying any particular risk model and implies a smaller spread than the spread estimates in their paper. The spread we estimate therefore measures the tightness of leverage constraints in any multi-factor generalization of their model. Also related to our work is Hébert (2018), who presents a theoretical model in which arbitrage spreads are due to constraints on the trading of financial intermediaries.

Finally, our work advances a literature on interest rates implied by option prices by using a simple, assumption-free estimation procedure on intraday quote data, free of microstructure noise. The majority of the literature compares the price of a put minus a call to the price of the underlying asset on which the options are written, requiring the econometrician to know the present value of dividends paid before the options expire. Recent work by Naranjo (2009), which uses this approach on daily Bloomberg data, finds a considerably more volatile rate than us, which is ascribed to microstructre-related supply/demand imbalances emphasized in Garleanu, Pedersen, and Poteshman (2009). The original paper of Brenner and Galai (1986), as well as Ronn and Ronn (1989), which was the first to study Box spreads, also uses American rather than European options for which put-call parity is only an approximate relationship. Closest to our work is Avino, Stancu, and Wese Simen (2017), which is the only paper to our knowledge that infers rates from European options without assumptions on dividends paid. Their rate series is quantitatively close to ours at low frequencies, although it is estimated with daily OptionMetrics data and a more complex nonlinear estimation procedure. By estimating our rates using minute-by-minute linear regressions with R squared's extremely close to 1, we document that options prices are all consistent with a unique implied interest rate which satisfies put call parity. Moreover, we show that this rate has minimal intraday volatility and rapidly responds to new information, suggesting that it is the most precisely estimated option-implied rate series to date. This is crucial for applications such as our high frequency even studies on monetary transmission.

The paper proceeds as follows. In Section 2, we show how we use the put-call parity relationship on European options to estimate risk-free rates. We perform several robustness analyses with respect to other interest rate measures in Section 3. In Section 4 we document that bid ask spreads do not bias the level of our estimated interest rates. In Section 5 we explore the effects of monetary policy announcements on our estimated rates and compare them to the effects on government bonds and convenience yields (the difference). In Section 6 we explore to what extent the dynamics of the term structure of convenience yields adds to bond return predictability. In Section 7 we explore CIP deviations using option-implied interest rates across markets. In Section 8 we show how the price discovery for Box rates compares favorably to those in Treasury markets. Section 9 concludes. In the Online Appendix we further show that observed trade data lead on average to identical rate estimates, though using trade data would make our rates noisier. In addition, we show in that appendix that (1) the residuals from our rate estimation procedure are extremely small, (2) extreme residuals mean revert within days to zero (though less quickly after the financial crisis), and (3) extreme residuals do not seem related to the level of our estimated rates.

2. Risk-Free Interest Rates without Convenience Yields

In this section we first explain conceptually why it is possible to use risky asset prices to infer an interest that is free from the convenience yield on safe assets. We then propose a novel estimator of a term structure of convenience-yield-free interest rates, which uses European-style options on the S&P500 traded on the Chicago Board Options Exchange (CBOE).

2.1. How to Estimate the Convenience Yield on Safe Assets

One difficulty in estimating the convenience of a safe, money-like asset is finding an appropriate interest rate to compare it to. Generally, the literature cited above approached this problem by using the yield on a less liquid and/or less safe asset for comparison. However, any sufficiently safe asset can itself have a convenience yield, since it potentially has money-like attributes as well. In the IS-LM model (Hicks (1937)), the nominal interest rate measures the return that agents forgo in exchange for holding cash. More recently, Krishnamurthy and Vissing-Jorgensen (2012) show that interest-bearing Treasury bonds earn a convenience yield because they perform a role similar to money in the financial system. As a result, the spread between the nominal rate earned on Treasury bonds and the zero rate earned on cash identifies the difference in convenience yields between these two assets. Similarly, if we estimate the convenience yield on Treasuries by comparing their yield to that of another (nearly) safe asset, such as repos, AAA bonds or commercial paper, we face a similar problem. Any asset that is safe enough to compare its yield to that of a Treasury bond to identify the Treasury's convenience yield may itself perform a money-like role in the financial system and also earn a convenience yield.

A very low risk (but not riskless) asset can therefore have a yield that is either above or below the convenience-yield-free riskless interest rate because the asset may have both a credit spread and a convenience yield. This is illustrated in Figure I which plots expected returns against the covariance of an asset's return with the Stochastic Discount Factor (SDF).[7] The circle labeled T represents the yield on safe assets such as Treasuries. The circle labeled B represents the convenience-yield-free riskless rate. The circles A and C represent two nearly safe assets such as commercial paper and AAA-rated bonds. The ideal measure of the convenience yield is the difference in yields between B and T. If instead of asset B, we would use either asset A or C as the counterfactual, the convenience yield would be mismeasured and could be either overestimated (in case of asset A) or underestimated (in case of asset C).

In this paper we provide a new way of estimating the convenience-yield-free riskless rate (circle B in the figure). To do so, we need to extrapolate from the pricing of risky assets (on the blue line in Figure I) the implied risk-free rate (B). One approach would be to use prices of stocks and construct a portfolio whose return does not covary with the stochastic discount

[7]We present a model in Appendix A in which our convenience yield estimate is valid even in the presence of financial frictions that distort this stochastic discount factor. The model features equity issuance costs, debt overhang costs, and segmentation between the market for financial assets and for loans backed by those assets.

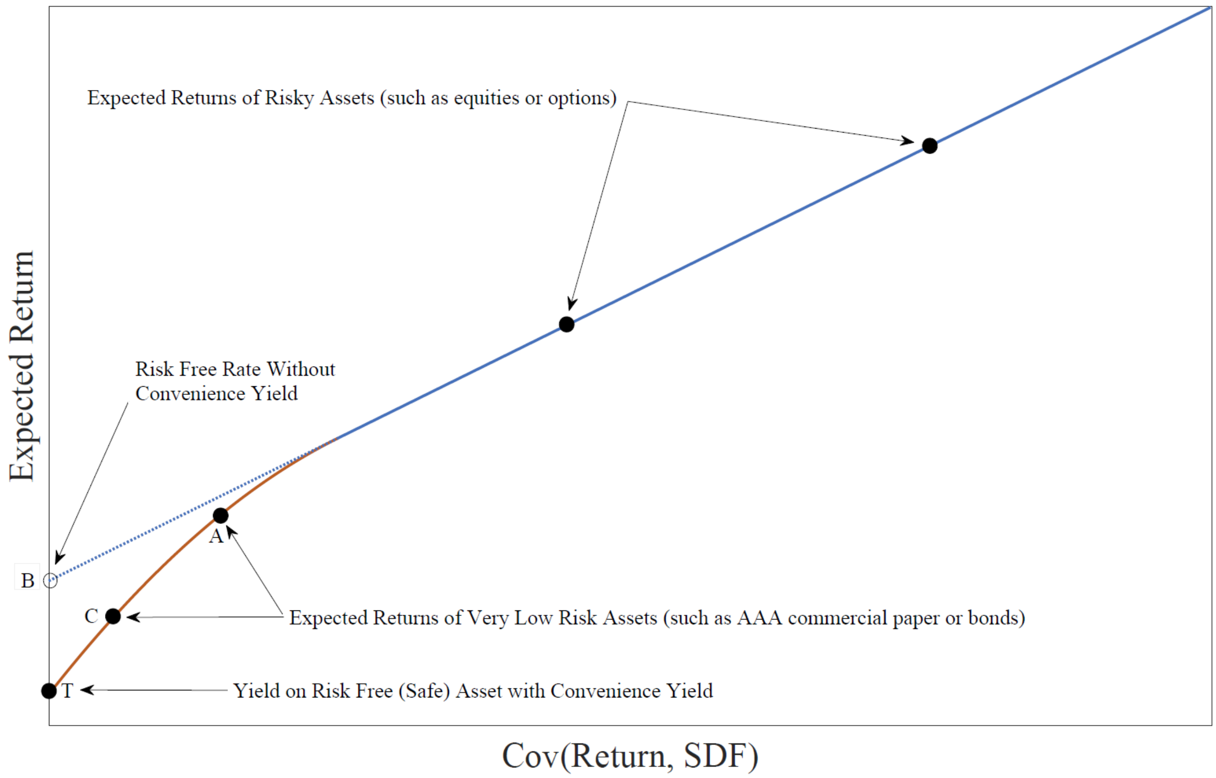

FIGURE I
Risk and Return with a Special Demand for Safe Assets (Convenience).

factor (often called the zero-beta rate (Black (1972) and Frazzini and Pedersen (2014)). One potential downside of using this zero-beta rate as a proxy for the risk-free interest rate is that it depends on a specific model for risk and return. For example, if we use the Capital Asset Pricing Model (Sharpe (1964)) as the risk model, the implied zero-beta rate would be influenced by risk factors not included in that model. Because there is little consensus in the literature as to what the correct risk model is, an accurate estimate of the convenience-yield-free risk-free rate must be invariant to the specific model of risk in asset markets. We propose such an estimate in the next subsection.

2.2. Constructing Risk-Free Assets

The starting point of our analysis is the put-call parity relationship for European options. At each time t, for each time to maturity T, option price quotes are available for a large cross-section of different strike prices indexed by $i = 1, ...N$. The put-call parity relationship then states that at time t, for each time to maturity T, and each strike price K_i, the difference between the put price $p_{i,t,T}$ and call price $c_{i,t,T}$ equals the discounted value of the strike K_i

minus the current value of the underlying S_t, where we need to adjust the latter for the present value of the cash flow that the security delivers.[8] Denote this present value of the cash flow by $\mathcal{P}_{t,T}$, then the put-call parity relationship is given by:

$$p_{i,t,T} - c_{i,t,T} = (\mathcal{P}_{t,T} - S_t) + \exp(-r_{t,T}T)K_i. \tag{1}$$

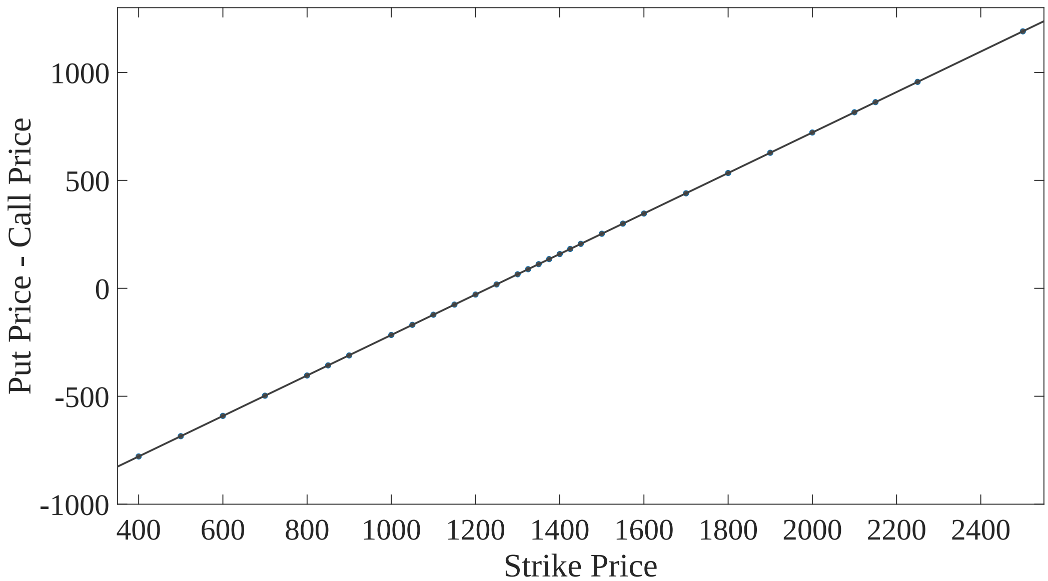

FIGURE II

The figure plots the difference between the put prices and the call prices (using midpoints) on the S&P 500 index against the underlying strike prices of the longest maturity option pairs. The maturity of the options is 824 days. The data is from 3:00pm on the day of Lehman's default (September 15, 2008).

As an illustration, we plot in Figure II option quote data from the S&P 500 index on September 15, 2008 (Lehman's default) at 3:00pm. We plot the difference between the put prices and the call prices (using midpoints) against the underlying strike prices of the longest available maturity option pairs (with a 824 day maturity). The figure also plots the fitted line. The figure illustrates that even on days of financial distress (Lehman's default), the put-call parity relationship holds within visual accuracy (with an R-squared value of 0.9999998),

[8]For dividend paying stock indices this price is the present value of the dividends paid out between time t and T, also called the dividend strip price (van Binsbergen, Brandt, and Koijen (2012)). As an application of our estimated rates, Golez and Jackworth (2020) reestimate dividend strip prices using option-implied interest rates. They find somewhat lower estimates of the dividend risk premium compared to van Binsbergen et al. (2012) and conclude that the moments of the dividend assets are in line with the index. As such there is still a substantial short-duration dividend risk premium in excess of the short-duration maturity matched bond. Binsbergen (2020) constructs a duration-matched government bond portfolio for the stock index and shows that the compensation investors have received for long duration dividend risk is zero (and potentially even negative) between 1970 and 2020.

leading to highly precise estimates of the implied interest rate. The estimated slope equals 0.93778, which corresponds to an annualized interest rate of 2.85%. The OLS standard error of the estimated slope equals 8×10^{-5}.

The put-call parity relationship provides two ways of obtaining a time series and term structure of the risk-free interest rate $r_{t,T}$ implied by the option market.

Estimator 1: At each time t and for each maturity T, we run the following cross-sectional regression:

$$p_i - c_i = \alpha + \beta K_i + \varepsilon_i \tag{2}$$

where the slope of the line is equal to:

$$\beta = \exp(-r_{t,T} T), \tag{3}$$

and where the intercept is equal to:

$$\alpha = \mathcal{P}_{t,T} - S_t. \tag{4}$$

The continuously compounded risk-free interest rate at time t for maturity T therefore equals:

$$r_{t,T} = -\frac{1}{T}\ln(\beta). \tag{5}$$

The estimated β of this regression can also be interpreted as the realized risk-free return that is earned on a particular trading strategy. To see this, consider the Ordinary Least Squares (OLS) estimator of the slope:

$$\beta_{OLS} = \frac{\sum_i \left((p_i - c_i - \overline{p-c})(K_i - \overline{K})\right)}{\sum_i (K_i - \overline{K})^2} \tag{6}$$

where

$$\overline{p-c} = \frac{\sum_i (p_i - c_i)}{N} \tag{7}$$

and

$$\overline{K} = \frac{\sum_i K_i}{N}. \tag{8}$$

So the strategy (also sometimes called the "Box" trade) involves buying (writing) a total of $K_i - \bar{K}$ put options for which the strike is above (below) average and writing (buying) a total of $K_i - \bar{K}$ calls for which the strike is above (below) average for each $i \in 1, ..., N$. This strategy will deliver the continuously compounded risk-free rate equal to $-\frac{1}{T}\ln(\beta_{OLS})$.

Estimator 2: At each time t and for each maturity T take all possible combinations of strikes, indexed by $1, ..., A$ where $A = \frac{N(N-1)}{2}$ and compute an implied risk-free rate for that strike pair. That is, $\forall i \in i = 1, ..., N$ and $\forall j \in i = 1, ..., -i, ..., N$ for which $K_i > K_j$, we compute:

$$r_{t,T,a} = -\frac{1}{T}\ln\left(\frac{(p_{i,t,T} - c_{i,t,T}) - (p_{j,t,T} - c_{j,t,T})}{K_i - K_j}\right), \tag{9}$$

with $a \in 1, ..., A$. We then compute the estimate for the risk-free as the median over all these implied rates:

$$r_{t,T} = \text{median}_{a \in A}\left(r_{t,T,a}\right). \tag{10}$$

This estimator, known as the Theil–Sen estimator, allows for robust estimation of the slope of the regression line even when there are large outliers in the underlying data. It also corresponds to a trading strategy, which is to invest in the strike pair i and j that deliver the median risk-free rate observation. Buying the put of strike K_i and the call of strike K_j while writing the call of strike K_i and the put of strike K_j yields a riskless payoff of $K_i - K_j > 0$ if all options are held to maturity. Because buying and writing these puts and calls costs a total of $(p_{i,t,T} - c_{i,t,T}) - (p_{j,t,T} - c_{j,t,T})$, this trading strategy earns exactly the risk-free rate corresponding to the Theil-Sen estimator.

2.3. Data

Our options data contains all option trades and quotes from the Chicago Board Options Exchange (CBOE) on two underlying assets: the S&P 500 index (SPX) and the Dow Jones Index (DJX), between 2004 and 2018. The traded options on these underlying assets are European, implying that the put-call parity relationship should hold exactly (for American options it only holds with an inequality).[9] The data set contains the bid price, the ask price, the strike and the maturity date for a large range of strike prices for each minute. We compute risk-free rate estimates at the minute level using the mid prices of puts and calls of all strike prices with a particular maturity. To compute daily estimates, we then take a median over the minute-level estimates in the day.

[9] Brenner and Galai (1986) estimate implied interest rates using American options and use both an early exercise correction as well as a correction for the dividend strip price. Our approach requires neither.

2.4. Counterparty Risk and Margin Requirements

In order for our interest rates to be risk free, there can be no meaningful credit risk in equity options traded on the CBOE. We believe this is true for several reasons. First, the CBOE requires investors to post margin and make daily variation margin payments, which ensure the safety of options except in the event of very extreme movements in the S&P 500. Second, the CBOE itself posts additional margin with the Options Clearing Corporation (OCC), which clears the options traded on the CBOE. In the event of a bankruptcy, derivatives are "super-senior" and exempt from the automatic stay, so derivatives traders can immediately seize the collateral backing their trades. The CBOE has an A credit rating and the OCC a AAA rating, so the odds of both institutions being unable to make payments is remote. The OCC is also considered a "Systemically Important Financial Market Utility" and has access to emergency liquidity support from the Federal Reserve. During the extreme stock market crash of 1987, the Federal Reserve intervened to ensure that all derivatives contracts were successfully paid off (see Bernanke (1990)). The default of Lehman (active in the derivatives market) did not even expose the OCC to losses since Lehman's collateral was sufficient to pay off all derivatives counterparties (Faruqui, Huang, and Takts, 2018). In addition, if an investor is to lend at our estimated rate, they do not have to post any additional margin beyond simply providing the cash up front that is to be lent. This is because the CBOE recognizes that the Box trade is riskless and adjusts its margin requirements appropriately.[10]

2.5. Results: Estimated Interest Rates

We now describe the results. We estimate our benchmark rate using S&P 500 (SPX) index options, which gives us precisely estimated interest rates. We study risk-free rates implied by the Dow Jones (DJX) options for robustness in Section 3. In Table 1 we provide summary statistics for SPX-implied yields for three maturities: 6 months, 12 months and 18 months, and we compare them with the corresponding yields on government bonds as implied by the Nelson-Siegel-Svensson (NSS) parameters estimated by Gürkaynak, Sack, and Wright,[11] the continuously compounded LIBOR rates, and the fixed rate of an interest rate swap contract written on the Federal Funds rate (OIS).[12]

[10]See the CBOE margin manual at https://www.cboe.com/learncenter/pdf/margin2-00.pdf .

[11]For a description of the NSS procedure see Section E.2. Note further that almost identical results are obtained by using the Fama Bliss (1987) Treasury data.

[12]There are two ways to compute SPX-implied yields for fixed maturities. For simplicity, we linearly interpolate the two closest SPX-implied yields around each fixed maturity. Alternatively, we could fit a NSS

Table 1
Summary Statistics of SPX Option Implied Interest Rates 2004-2018

Zero Coupon Yields: 6 month maturity		
	Mean	St. Dev.
Option Implied SPX	0.0178	0.0174
LIBOR Implied	0.0185	0.0173
Government Bond	0.0142	0.0167
OIS	0.0143	0.0178
LIBOR Implied - Option Implied SPX	0.0007	0.0021
Option Implied SPX - Government Bond	0.0035	0.0022
Option Implied SPX - OIS	0.0035	0.0023

Zero Coupon Yields: 12 month maturity		
	Mean	St. Dev.
Option Implied SPX	0.0185	0.0171
LIBOR Implied	0.0210	0.0160
Government Bond	0.0148	0.0164
OIS	0.0148	0.0177
LIBOR Implied - Option Implied SPX	0.0024	0.0026
Option Implied SPX - Government Bond	0.0037	0.0021
Option Implied SPX - OIS	0.0036	0.0020

Zero Coupon Yields: 18 month maturity		
	Mean	St. Dev.
Option Implied SPX	0.0194	0.0167
Government Bond	0.0157	0.0159
Option Implied SPX - Government Bond	0.0037	0.0021
Option Implied SPX - OIS	0.0037	0.0017

The table shows that for all maturities the average yields on the SPX-implied interest rates are above those of the corresponding government bonds and interest rate swaps, and below those of the LIBOR rate. The average difference between the SPX-implied rate and the government bond rate (i.e., the convenience yield), is 35-37 basis points per year, with very little variation across maturities. The average difference between the SPX-implied rate and the interest-rate swap fixed rate is also 35-37 basis points, and also essentially constant across maturities. The average difference between the LIBOR rate and the SPX-implied rate is positive for both the 6-month and 12-month maturities, and equal to 7 basis points and 24 basis points respectively. For the 18-month maturity, a LIBOR rate is not available. Furthermore, the LIBOR rate has the lowest volatility, and the interest-rate swap fixed rate has the highest.

To better understand the variation and comovement in the rates, we plot in Figures III,

yield curve. However, our shortest maturity yields are less precisely estimated, because a small amount of price mismeasurement implies a large error in yield estimates due to the scaling by maturity in the calculation of yields. Elsewhere in the paper, we handle this problem by dropping maturities less than 30 days and weighting our loss function by the inverse of duration.

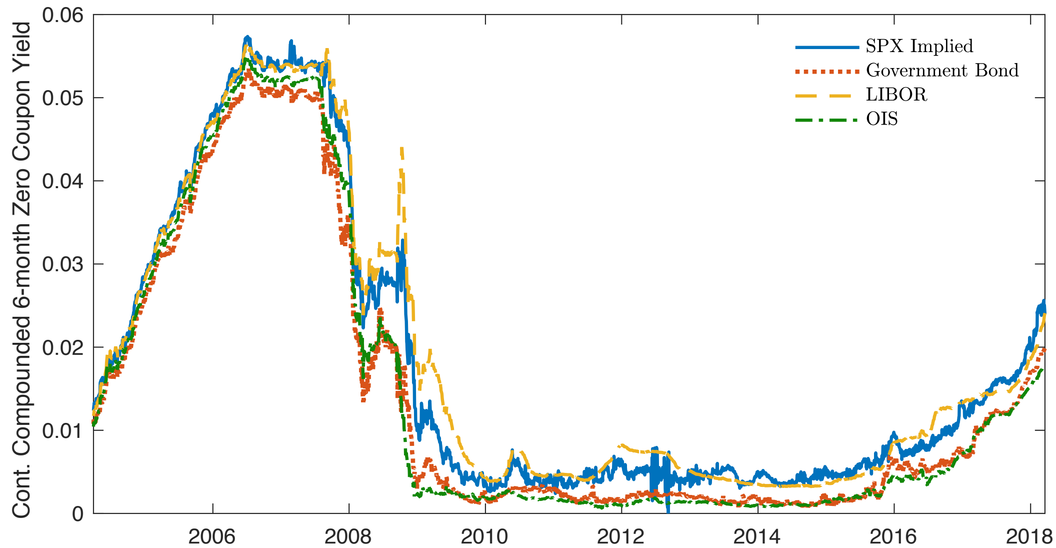

FIGURE III

Comparison of 6-month zero coupon interest rates implied from SPX options with government bond, LIBOR, and OIS rates. All rates are continuously compounded.

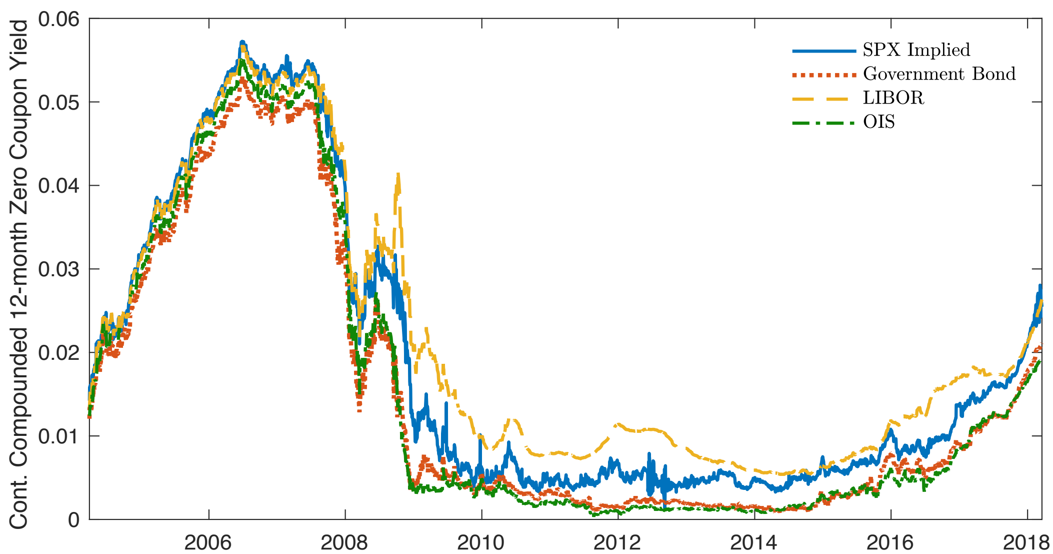

FIGURE IV

Comparison of 12-month zero coupon interest rates implied from SPX options with government bond, LIBOR, and OIS rates. All rates are continuously compounded.

IV and V the four interest rates for all three maturities. The three graphs show a consistent pattern. Before 2008 the SPX-implied yields are above the corresponding government bond yield, and closely follow LIBOR. Between 2008 and 2017 a substantial deviation from LIBOR occurs and the SPX-implied yields are in between the LIBOR rate and the government bond yield. This suggests that between 2008 and 2017 banks faced substantial credit risk, as measured by the spread between LIBOR and the SPX-implied zero coupon yield.

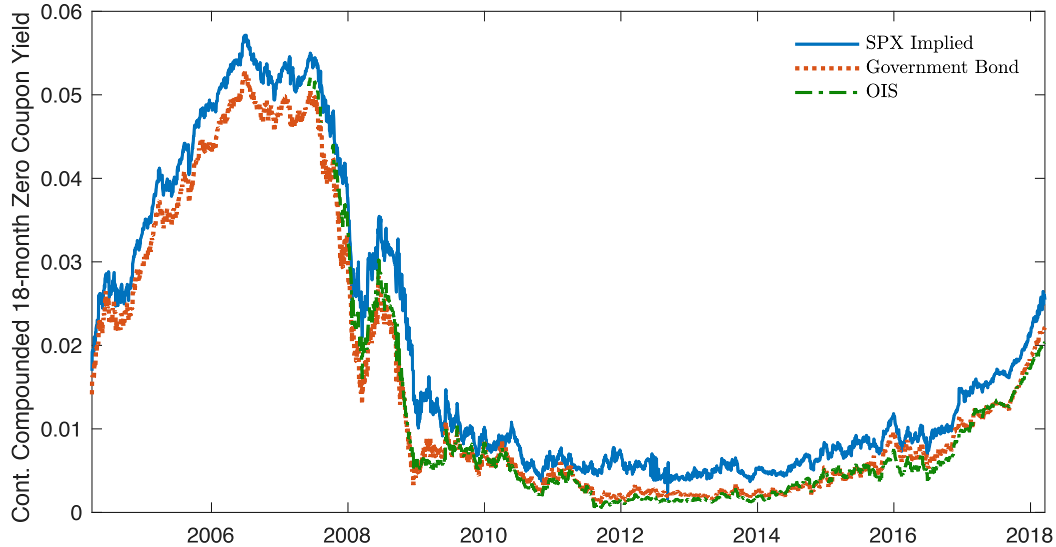

FIGURE V

Comparison of 18-month zero coupon interest rates implied from SPX options with government bond rates and OIS rates. All rates are continuously compounded.

Next, we present in Figures VI and VII a time series average of daily Nelson-Svensson-Siegel (NSS) yield curves fit to our SPX-implied rates and compare it to the benchmark Treasury yield curve of Gürkaynak et al. (2007) for both the full sample, as well as the year 2008. We add to these pictures a curve that is fit to constant maturity Treasury bill rates. The average spread between our yield curve and the Treasury curve is remarkably flat. The Treasury yield curve that is fit to long maturity notes and bonds implies higher short-term yields than bills themselves. This spread between actual T-bill yields and the short-term yields implied from the curve that is fit to long-term notes and bonds identifies the additional convenience yield on Treasury bills compared to notes and bonds. This additional convenience yields equals roughly 25 basis points. This is consistent with the idea common in the banking literature that short-term safe assets are somehow special, and financial institutions therefore have an incentive to finance themselves with large amounts of short-term safe debt to exploit the additional convenience yield it earns. Further, our entire term structure of convenience yields shifts outward but remains relatively flat if we restrict our data to only 2008, when the financial crisis was severe. This suggests that the scarcity of safe assets during the financial crisis was not restricted only to short-term debt, and investors were willing to pay a large premium for the safety of even 2.5 year Treasury bonds. Our data does not allow us to compute convenience yields beyond this maturity without extrapolating, so it is an open question whether the convenience yields on 10 or 30 year bonds behave similarly.

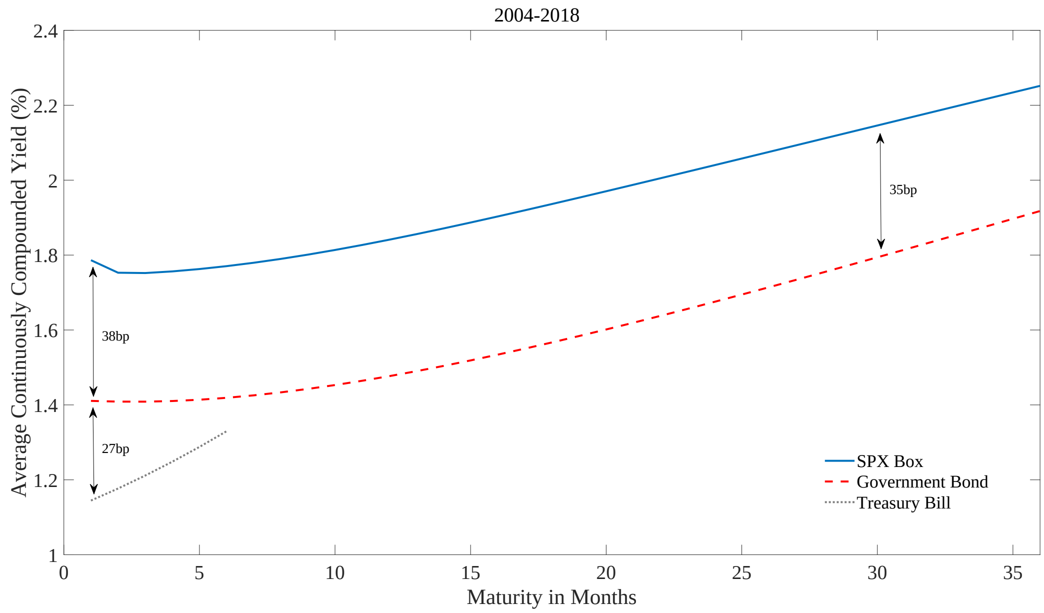

FIGURE VI

Average NSS yields curves fit to SPX Box rates and Treasury bond rates together with Treasury bill rates, 2004-2018. All rates are continuously compounded

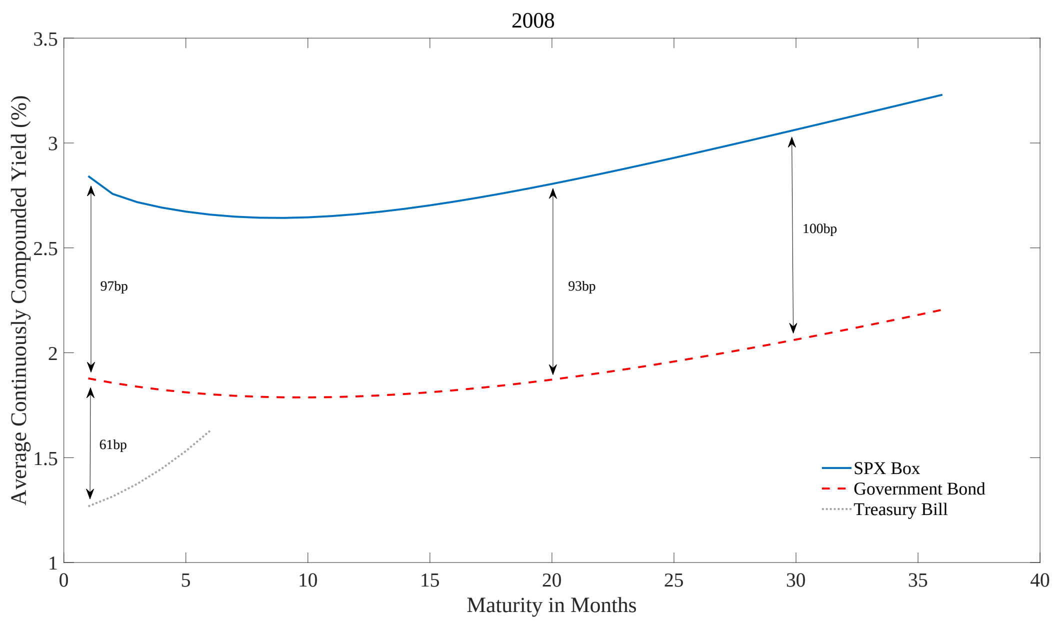

FIGURE VII

Average NSS yields curves fit to SPX Box rates and Treasury bond rates together with Treasury bill rates, 2008. All rates are continuously compounded

To further study the differences between the various available interest rates, we plot in Figures VIII, IX and X the spreads between the SPX-implied yield and the government bond yield, as well as the spread between LIBOR and the SPX-implied yield with maturities of 6 months, 12 months and 18 months. As LIBOR rates only have maturities up to 12 months, we only plot the spread between the SPX-implied yield and the government bond yield for that maturity. For all three maturities, both spreads exhibit large variation, and they both go up during the crisis and have since been reduced to levels closer to zero.

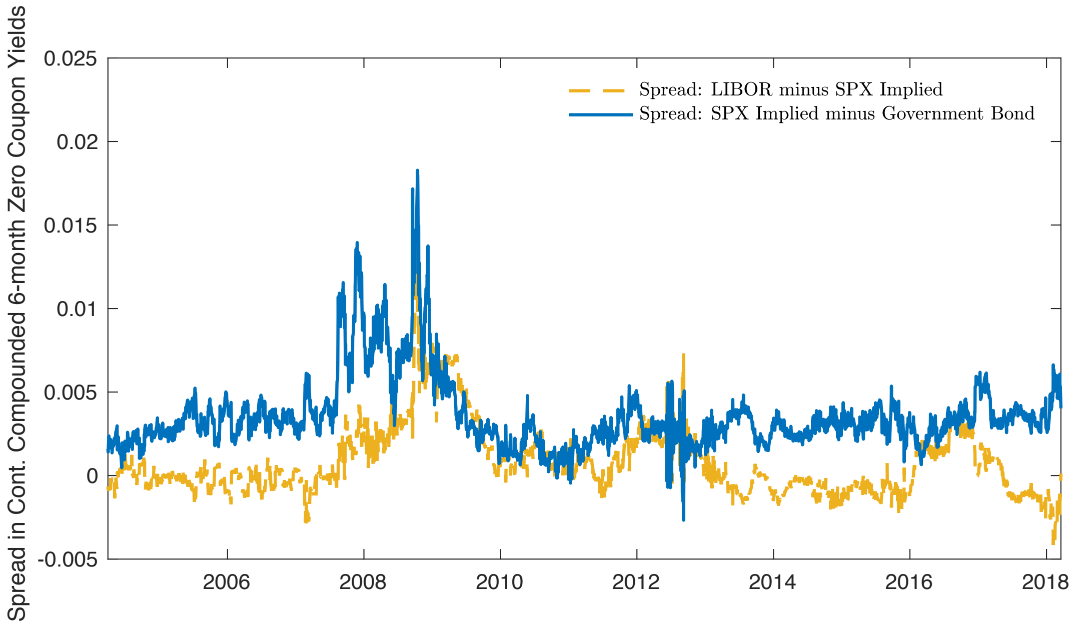

FIGURE VIII

Spreads of 6-month zero coupon interest rates implied from SPX options with government bond rates (the convenience yield) and LIBOR rates. All rates are continuously compounded.

2.6. Results: Precision of Estimated Rates

In this subsection we evaluate the precision with which our rates are estimated. If no-arbitrage conditions held perfectly, the R-squared of the regression in Equation 2 would equal 1 and there would be no estimation error in our rates. As such, the R-squared of the regression can be interpreted as a measure of efficiency within the market for options for this particular underlying asset. Because the slope of the regression is so close to 1, we can also easily map this measure of market efficiency to variation in estimated (non-annualized) rates ($Tr_{t,T}$) across the strikes. To see this, note that the population R-squared of the regression in Equation 2 is given by:

$$R^2 = \frac{\text{var}(\beta K)}{\text{var}(\beta K) + \text{var}(\varepsilon)} = \frac{1}{1 + \frac{\text{var}(\varepsilon)}{\beta^2 \text{var}(K)}} \tag{11}$$

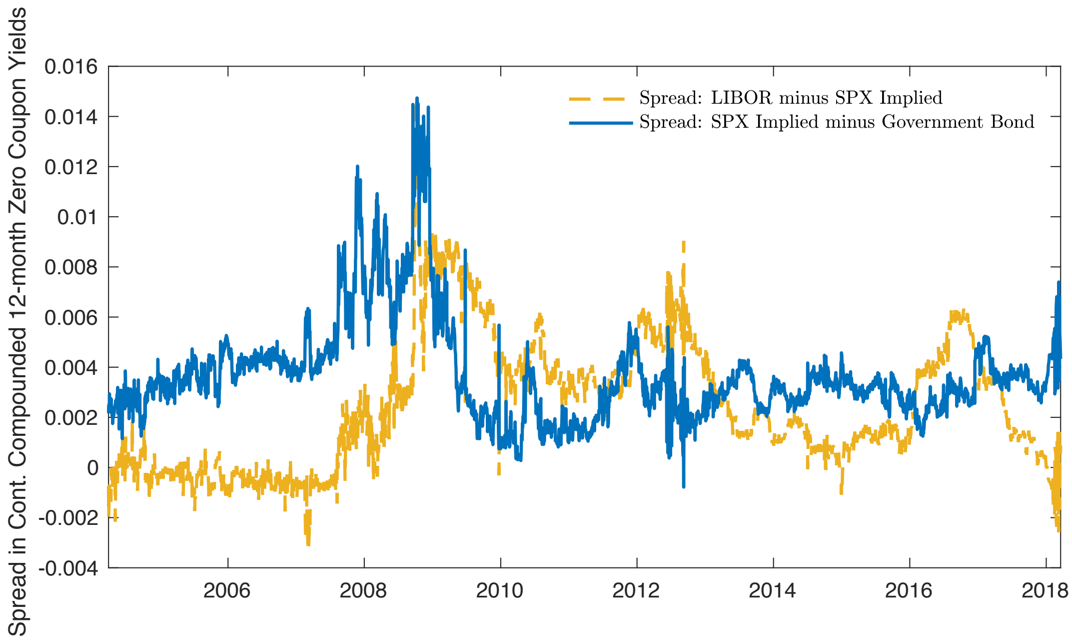

FIGURE IX

Spreads of 12-month zero coupon interest rates implied from SPX options with government bond rates (the convenience yield) and LIBOR rates. All rates are continuously compounded.

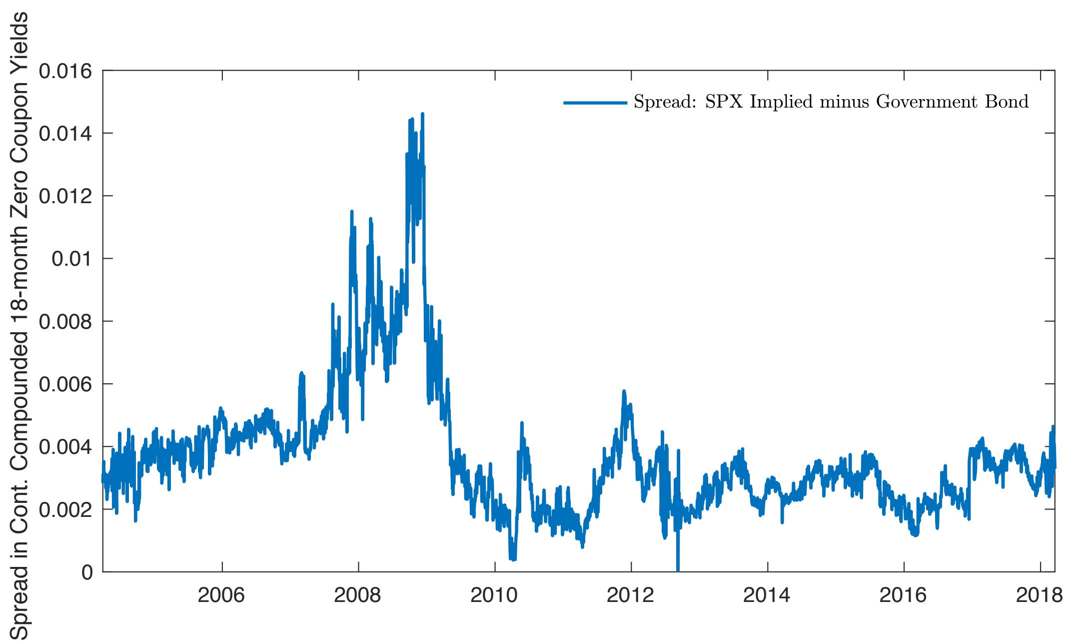

FIGURE X

Spreads of 18-month zero coupon interest rates implied from SPX options with government bond rates (the convenience yield). All rates are continuously compounded.

Rewriting this equation, we find:

$$\frac{1}{R^2} - 1 = \frac{\mathrm{var}(\varepsilon)}{\beta^2 \mathrm{var}(K_i)} \approx \frac{\mathrm{var}(\varepsilon)}{\mathrm{var}(K_i)}. \tag{12}$$

Assuming uncorrelated error terms, the asymptotic variance of the univariate OLS estimator equals the variance of the error term scaled by N times the variance of the right-hand side variable, that is, the variance across the strike prices. This then implies that the variance of the OLS estimated interest rates can be approximated by (using the approximation that β is close to 1 and that the log-linearized regression coefficient uncovers the interest rate):

$$\sigma(\hat{r}_{t,T}) \equiv \sigma\left(\frac{1}{T}\ln(\beta_{OLS})\right) \approx \sqrt{\frac{1}{NT^2}(\frac{1}{R^2}-1)}. \tag{13}$$

As a consequence, for a regression for maturity $T = 1$, with $N = 20$ strike prices and an R-squared of 0.999999, the standard error of the estimate at each time t (i.e. each minute) is in the order of magnitude of 2 basis points. For 100 strikes, this number is 1 basis point. Given that our daily estimates are computed by taking a median over the minutely observations, the standard error of the daily estimate is even smaller than that. As an illustration, we plot in Figure XI a daily series of the standard error of the minute-level risk-free zero coupon yield estimate for the 18 month maturity. We use the actual standard error implied by the regression, which is approximately equal to the non-linear transform of the R-squared as explained in Equation 13, and as such can be interpreted as a measure of market efficiency. To arrive at a daily series for this minute level standard deviation, we take the median standard error across all minutes within a day. The graph shows that the

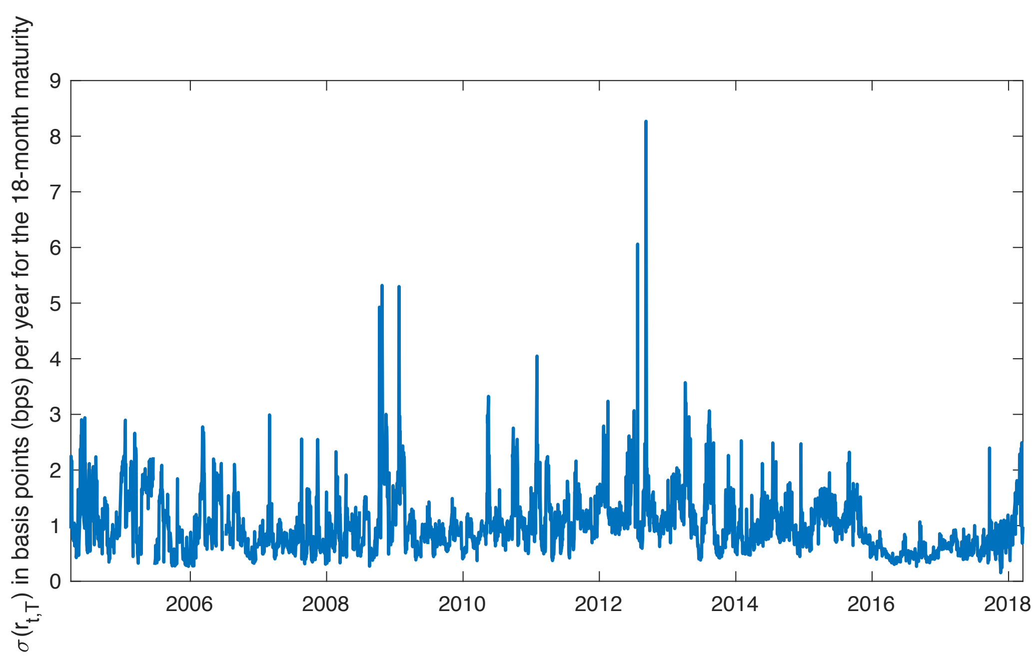

FIGURE XI

Efficiency in the SPX option market expressed as the standard error around the implied risk-free rate.

typical standard error is in the order of magnitude of 1 basis point, but it can occasionally spike. The maximum over our sample period is 8 basis points.

One implication of our very small standard errors (or equivalently the very high values of our cross-sectional R-squared measures) is that put-call parity holds almost exactly (with the appropriate interest rate) in our SPX option data. If our R-squared values were not so high, the deviation of our rates from Treasury rates could arguably be interpreted as a measure of put-call parity violations unrelated to the convenience yield on safe assets. However, the data shows that the implied interest rates across various different strike prices seem to be highly internally consistent. This suggests the absence of put-call parity violations and the presence of a convenience yield for safe assets. Under this interpretation, our option-implied rates precisely measure the risk-free cost of capital for investors in this option market.

An additional implication of our very small standard errors is that it is unlikely that our rate estimates are being distorted by our use of midpoint price quotes. If the midpoint of a bid and an ask did not accurately measure the value of an option, then there is no reason to believe that such midpoints would satisfy no-arbitrage relationships such as put-call parity with such precision. The only way for midpoints to satisfy put-call parity is either if (1) mid points do accurately reflect the value of an option or (2) a knife-edge case holds where the difference between the midpoint price and true value of an option happens to be an exactly linear function of the strike price. The simplest interpretation of this evidence is that midpoints do in fact provide an accurate measure of the value of an option and that the rates we infer from put-call parity measure the risk-free cost of capital, at least for market makers and other intermediaries that can execute trades at midpoint prices.[13]

2.7. Interpretation of Rate Estimates

The Box rates estimated in this paper are a candidate measure for riskless interest rates free of the convenience yield on safe assets. Other candidate rates in the literature are inferred directly from the fixed income market rather than indirectly from the prices of risky assets. As a consequence, such rates may themselves be affected by varying degrees of convenience. Many low risk assets (such as Treasuries, Repos, and Refcorp bonds) can help a financial institution satisfy its regulatory requirement to hold so-called high quality liquid assets. Other assets (such as Treasury bills and high grade Commercial Paper) can be held by money market funds and therefore be used to back money-like assets held by

[13]In unreported results, we also find empirically that the R-squared of our regressions is higher using midpoints than using either bids or asks as our option price measure.

consumers. The rates on these securities should therefore include a premium for meeting regulatory requirements unlike our rate. Consistent with this, Treasuries and Repos have yields below our rate. The volume of trade in some safe assets (such as Refcorp bonds and Commercial Paper) is low relative to SPX options, and the illiquidity of these assets may raise their yields for reasons distinct from the convenience yield we want to estimate.

In Appendix A, we provide a theoretical model that provides conditions under which the spread between Box and Treasury yields can be interpreted as a measure of convenience. The spread between the Box rate and the yield on a safe asset is the convenience yield to an investor that 1. is able to lend at the Box rate we estimate using midpoints of bid and ask prices and 2. holds a strictly positive quantity of the safe asset. While we do need to assume that the investor is active in both markets and that midpoints reflect its valuations of options, the investor can be subject to a range of financial frictions that we describe in Appendix A.

One class of agents who are active in option markets and can trade with minimal transaction costs are option market makers. The most restrictive interpretation of our convenience yield estimate is that it is specific to these option market makers. For some of the equity options traded on the CBOE, the CBOE lists a "Designated Primary Market Maker" who is obligated to provide price quotes on those options to other traders. The identities of these Designated Primary Market Makers are given in the table below.

Table 2
Identities of Designated Primary Market Makers for CBOE Equity Options

Name of Institution	Number of Underlyings	Percent of Underlyings
Susquehanna Securities, LLC	1,588	35.27
Citadel Securities, LLC	824	18.30
Wolverine Trading, LLC	731	16.24
Barclays Capital Inc.	408	9.06
Group One Trading, L.P.	309	6.86
Sumo Capital, LLC	252	5.60
Morgan Stanley & Co. LLC	199	4.42
Two Sigma Securities, LLC	105	2.33
Citigroup Derivatives Markets Inc.	86	1.91

Three of these market makers (Barclays, Morgan Stanley, Citigroup) are regulated dealer banks accounting for 15.4 percent of underlying assets, while the remaining 84.6 percent of market making is performed by hedge funds who are less regulated. For our dealer banks, regulatory requirements are one likely reason why they would be willing to invest in Treasuries with lower yields than the option-implied rates. Treasury securities count as High Quality Liquid Assets that help the banks satisfy their Liquidity Coverage Ratio. Hedge funds, however, do not face such requirements. Given that hedge funds are the dominant participants in this market, our model in Appendix A is motivated by the frictions faced by hedge funds. The model features frictions in raising new equity capital that impact asset prices as in He and Krishnamurthy (2013), debt overhang frictions in the relationship between a hedge fund and its prime brokers (Andersen, Duffie, and Song (2019)), differences in the hedge fund's and prime broker's pricing kernels (as in Diamond (2018)), as well as a convenience yield for the hedge fund to hold certain special safe assets (e.g. Krishnamurthy and Vissing-Jorgensen (2012)). We show in this model, that even in the presence of all these other frictions, the spread between our rate and the rate on Treasuries identifies the convenience yield of Treasuries for a hedge fund that is active both in the Treasury market as well as in option market making.

In order for our theoretical model to be practically relevant, it must be the case that hedge funds are active in multiple asset classes. This is indeed true in the position data available. Citadel Securities, who provides the most detailed annual reports, states that "The Company primarily engages in market making and liquidity provision in U.S. options, equities, government securities, and foreign exchange products, as well as trade execution." They have \$8.9 billion in government securities and \$11.1 billion of options on the asset side of their balance sheet. Susquehanna Securities also holds a modest quantity of government securities (and likely more in their parent company Susquehanna International Group). Wolverine securities is also active in the corporate bond markets, and all hedge funds are active in equity markets. If these hedge funds use a single risk free rate to price assets across asset classes, then the option-implied rate we estimate would be the risk free rate implicit in stock and bond prices as well. In summary, given that the primary market makers in option markets are also active in many other asset markets, the risk free rate that we estimate could identify the risk free rate implicit in a pricing kernel for a larger group of risky assets beyond the narrow confines of the option market. It could therefore be used as a benchmark for the mean of an intermediary stochastic discount factor.

2.8. Convenience Yields and the Demand for Safe Assets and Derivatives

While the theoretical model in Appendix A provides general conditions under which the spread between our Box rate and Treasury yields have an interpretation as a convenience yield to intermediaries, final investors' demand for either options or Treasuries can impact this spread. In the "Demand-Based Option Pricing" literature starting with Garleanu et al.

(2009), an increase in investors' demand to buy options impacts option prices but not the underlying stock on which the options are written. This is because intermediaries cannot risklessly exploit arbitrage opportunities between stock prices and option prices. However, proposition 2.i of Garleanu et al. (2009) shows that the put-call parity relationship we exploit still holds in such a model, so our Box rates correctly identify the risk-free cost of capital for derivatives dealers.[14]

The more recent work of Hazelkorn et al. (2020) constructs an interest rate implied by spot and futures prices on the S&P 500 index and documents that this rate is impacted by shocks to the demand for futures trading. In their model, they say "the basis between futures and spot prices emerges because of the balance sheet cost that futures dealers face to meet customer demand for futures and hedge their exposure in the spot market." In the interpretation of the term "convenience yield" justified by our model in Appendix A, their model shows that the futures basis they compute has an interpretation as a convenience yield. The basis shows the rate of return that the dealer is willing to forgo in order to hold a riskless asset instead of a collection of risky assets that replicates the same riskless payoff. Empirically, Hazelkorn et al. (2020) document that their futures basis has a correlation of .8 with our convenience yield measure, and that shocks to the demand for futures are correlated with the size of both our spread and their futures basis. Together with the evidence in our paper, this suggests that the convenience yield (which can be interpreted as the "price" of holding a safe asset) we estimate is impacted both by demand for safe assets, as well as shocks to the demand for futures (and perhaps other risky securities) that impact the balance sheets of financial institutions. Even so, the evidence we present below shows that our Box rate is consistent with discount rates in a broad range of asset markets and therefore provides a useful candidate measure of the cost of capital for financial institutions active in many asset classes.

2.9. Relation to Previous Option-Implied Rate Estimates

Our paper is related to two other papers that use option prices to construct an interest rate series, Naranjo (2009) and Avino et al. (2017). Both papers use daily data from OptionMetrics. Naranjo (2009)'s estimation procedure compares a synthetic version of the underlying asset to the price of the underlying asset and requires an additional proxy for the

[14] In Garleanu et al. (2009), the fact that derivatives are not perfectly spanned by a trading strategy in the underlying asset implies that the risk preferences of derivatives dealers are reflected in the equilibrium pricing kernel. Our theoretical framework in Appendix A shows that a convenience yield for safe assets can be estimated even in the presence of a pricing kernel impacted by such financial frictions.

present value of dividends paid by the underlying asset before the options mature. He finds that "implied interest rates do not resemble benchmark interest rates such as the three-month T-bill rate or LIBOR, but instead are much more volatile. [He] argue[s] that the volatility in the implied short-term rate in futures and option markets is due to frictions arising from borrowing and short-selling costs." Our implied rates do not feature this high degree of volatility, due to both our different estimation procedure and by taking a median each day over our minute-by-minute rate estimates. Like us, Avino et al. (2017) estimates interest rates only from option prices without requiring additional assumptions about dividend payments. They estimate a spread over Treasury yields that, on average, is very close to ours (36 vs 40 basis points). Unlike our rates, their term structure inverts during the 2008 financial crisis, while our term structure of convenience yields stays (on average throughout the year) almost completely flat. Our series therefore seems relatively immune to short-term fluctuations in frictions specific to the options market, as we document in our robustness analysis. One other advantage of our series is that it is estimated minute-by-minute, allowing us to perform high frequency event studies and document a rapid response to news. Finally, we are the first to document that the the linear relationship implied by put call parity holds so precisely in the data, confirming that there is indeed a unique discount rate consistent with options of all strike prices rather than a large degree of heterogeneity and noise.

3. Robustness to Other Interest Rates

In this section we perform several robustness analyses related to our implied interest rates. First, we compare our SPX-implied Box rates to four other interest rates: the GC repo rate, the DJX option-implied Box rate, the risk-free rate implied by precious metal futures, and the risk-free rate implied by combining stock futures and options.

3.1. Relation to Other Interest Rates

In this subsection, we compare our Box rate to three other risk-free interest rate proxies. First, we show that the average level of the General Collateral (GC) repo rate (which is only available for short maturities) equals that of the government bond yields implied by the NSS curve. As a result, GC repo also seems to earn a convenience yield close to that of government bonds. More importantly, we can confidently conclude that our rate is distinct from other common benchmarks in the literature, including government bond yields, OIS rates, LIBOR and GC repo rates.

Second, we estimate risk-free rates from DJX (Dow Jones) options instead of SPX options. We find that the implied rates and associated convenience yields are highly similar to those implied by the SPX, though the DJX estimated rates are substantially noisier, with lower R-squared values (and associated higher standard errors) in our estimated cross-sectional regressions. This demonstrates that our data on SPX options yields uniquely precise rate estimates and that the nearly perfect fit of the put-call parity relationship is due to the quality of the SPX option market, rather than a mechanical feature of how option quotes are generated.

Third, we combine the option data on the S&P 500 with futures data on the same underlying and maturity. We compute the risk-free rate implied from an arbitrage trade that takes a position in both assets. We document that the rate implied from the option-futures trade is nearly identical to our benchmark rate, but the rate estimates are substantially noisier than our Box rate.

Fourth, in the Online Appendix, we use the cost-of-carry formula for precious metal futures to infer implied interest rates from that derivatives market. Generally, the time-varying cost of storage of a commodity can complicate the estimation of a risk-free rate from the cost-of-carry formula. We resolve this issue by focusing on precious metals for which the storage cost as a fraction of the value of the underlying asset is minimal. Using this risk-free interest rate proxy, we once again find similar convenience yields to those implied by option markets (equal to about 40bp over our sample period).

3.1.1. GC Repo

In this section we study the GC repo rate and compare it to several other interest rates. The GC repo rate is the interest rate earned on a loan collateralized by a safe financial asset such as Treasuries, agency securities, or other members of the so-called "General Collateral" basket of safe assets. It is commonly used in the literature Nagel (2016) to measure a riskless rate of return that is higher than that earned by special liquid assets such as Treasury Bills. It is generally available for shorter maturities than the ones we study in this paper.

In Table 3 we compare the summary statistics of our 3-month Box rate to those of the 3-month government bond yield (implied by the NSS curve), the 3-month OIS rate and the 3-month GC repo rate. We find that the average rate across government bonds, OIS and GC repo are all very similar, whereas our implied Box rate is substantially above all three. We can therefore conclude that our rate is distinct from other common benchmarks in the literature, which all seem to feature some form of convenience yield.

Table 3
Summary Statistics of Various (Implied) Interest Rates 2004-2018

Zero Coupon Yields: 3-month maturity		
	Mean	St. Dev.
Option Implied SPX	0.0174	0.0176
Government Bond	0.0142	0.0169
OIS	0.0139	0.0178
GC Repo	0.0141	0.0173

3.1.2. Comparison with Other Convenience Yield Measures

We next compare the time series of our convenience yield estimate to the two most prominent previous estimates in the literature. We take a monthly average of our 6-month SPX Box-Treasury spread and plot it against the monthly GC Repo-Treasury Bill spread from Nagel (2016) and the annual AAA Bond-Treasury spread from Krishnamurthy and Vissing-Jorgensen (2012). All three spreads have large increases during the financial crisis. After the crisis, our spread and the AAA-Treasury spread remain somewhat elevated while the GC Repo-Treasury Bill spread drops even below its pre-crisis level. Unlike the other two, our spread increases in late 2011 and early 2012 before falling after Draghi's July 2012 "Whatever it Takes" speech, suggesting that it alone is exposed to the European financial crisis. Unlike our spread, the AAA-Treasury spread increases significantly in 2016 and 2017, while the GC Repo-Treasury Bill spread increases moderately. Because AAA bonds tend to be longer maturity than the SPX box yields we construct, this may reflect an increase in the convenience yield specific to long maturity Treasuries. If instead the convenience yield of Treasuries has a flat term structure (as we find within the maturities we consider), the increased AAA-Treasury spread in 2016-2017 must reflect other forces, such as a growing credit risk or illiquidity premium specific to fixed income markets.[15] At a monthly frequency, our series has a correlation of 0.826 with Nagel's spread in levels and a correlation of 0.517 in first differences. At an annual frequency, our rate has a correlation 0.858 with the GC Repo-Treasury Spread and 0.726 with the AAA-Treasury spread in levels and 0.901 and 0.803 in first differences. The GC Repo-Treasury Bill spread and AAA-Treasury spread only have a correlation of 0.413 in levels and 0.634 in first differences. Because our spread is more correlated with the other two spreads than they are with each other, ours is a good candidate for a single summary measure of the convenience yield of Treasuries, in the sense that it seems to capture the common variation. In addition, because at a monthly frequency our series

[15] We constructed a GC Repo-Treasury Bill spread until 2018 that had a correlation of .999 with that of Nagel (2016). This spread did not increase after the financial crisis like the AAA-Treasury spread.

is more correlated with Nagel's in levels than in differences, high frequency fluctuations in our spread (which are important for event studies) may not be reflected in other candidate measures of the convenience yield.

The level of our convenience yield estimate, above that of Nagel (2016) and below that of Krishnamurthy and Vissing-Jorgensen (2012), is also consistent with it being a good measure of the convenience yield of Treasuries. Because Repos themselves are relatively safe, money-like assets (often held, for example, by money market mutual funds to satisfy their regulatory requirements), the GC Repo-Treasury Bill spread should be interpreted as the difference in the convenience yields of the two assets. Conversely, AAA bonds have low trading volumes and some degree of credit risk. This suggests that the AAA-Treasury spread reflects both the illiquidity and (small) risk of AAA bonds unlike our spread. One caveat in interpreting our spread is that it may be exposed to shocks specific to the demand for leveraged trade through options or other derivatives as documented in Hazelkorn et al. (2020). While any asset class will face some such idiosyncratic shocks, we show in a multivariate analysis of many arbitrage spreads in the Online Appendix that ours has the least idiosyncratic variation. While some fluctuations in our spread are likely due to the demand imbalances studied in Hazelkorn et al. (2020), it seems to be the least exposed to such forces in comparison to other alternatives.

3.1.3. Box Rates from the Dow Jones Industrial Index (DJX)

Next, we repeat the Box rate estimation that we performed for the S&P 500 index for options on the Dow Jones industrial index (DJX). In Table 4 we summarize the results for the median estimator (estimator 2 in Equation 10).[16] We find highly comparable results to those of the SPX: the implied interest rate is on average higher than the government bond yield by about 40 basis points, which is invariant to maturity.

Given how comparable the results for the DJX are to the SPX, we only plot the implied continuously compounded interest rate for the 1-year maturity as an illustration in Figure XIII. The graphs exhibit very much the same pattern, though the DJX implied rates are somewhat noisier than the ones implied by the SPX. Next, we repeat the efficiency analysis of Figure XI but now for DJX. The results are summarized in Figure XIV, where we plot the standard error of the OLS estimate of Equation 2. The results are comparable to the SPX though the average level of efficiency is substantially lower, with an average standard error of the minutely level estimated rate equal to 3.4 basis points, and spikes that occasionally go as high as 38 basis points. Because our daily estimates are computed by taking a median

[16]The regression-based estimator gives highly comparable results.

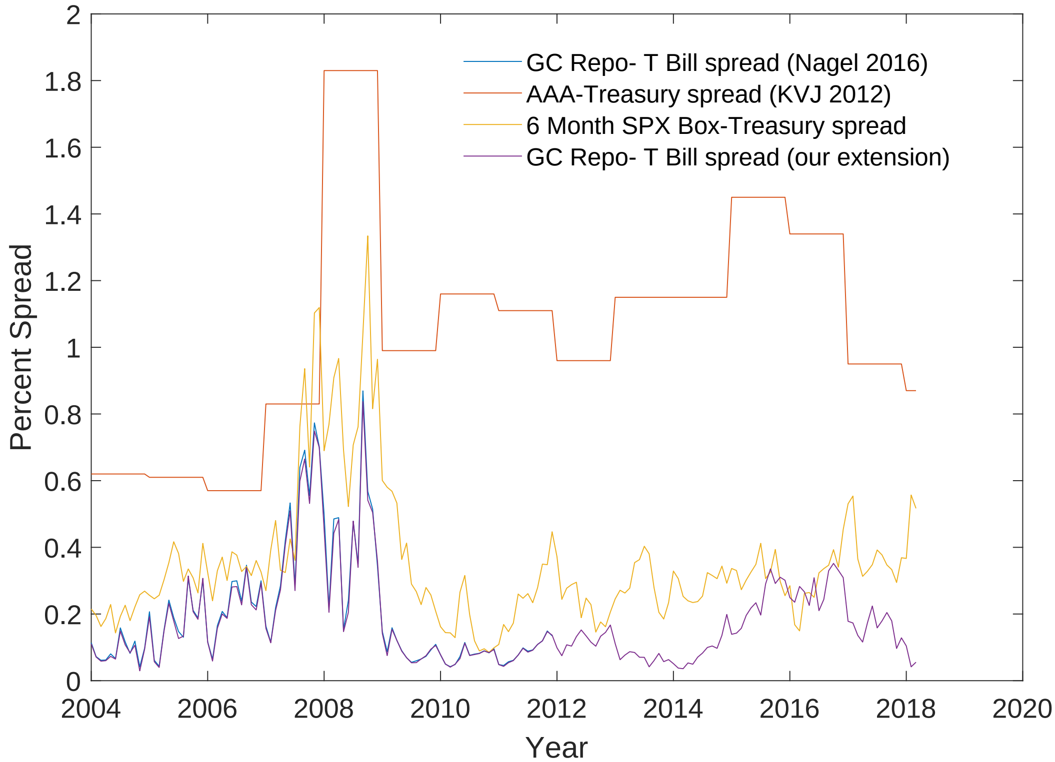

FIGURE XII

Comparison of 6 month SPX Box-Treasury spread to previous convenience yield estimates

Table 4

Summary Statistics of DJX Option Implied Interest Rates 2004-2018

Zero Coupon Yields: 6 month maturity		
	Mean	St. Dev.
Option Implied DJX	0.0184	0.0171
Government Bond	0.0144	0.0166
Option Implied DJX - Government Bond	0.0040	0.0029

Zero Coupon Yields: 12 month maturity		
	Mean	St. Dev.
Option Implied DJX	0.0190	0.0168
Government Bond	0.0150	0.0163
Option Implied DJX - Government Bond	0.0040	0.0023

Zero Coupon Yields: 18 month maturity		
	Mean	St. Dev.
Option Implied DJX	0.0197	0.0164
Government Bond	0.0159	0.0158
Option Implied DJX - Government Bond	0.0039	0.0021

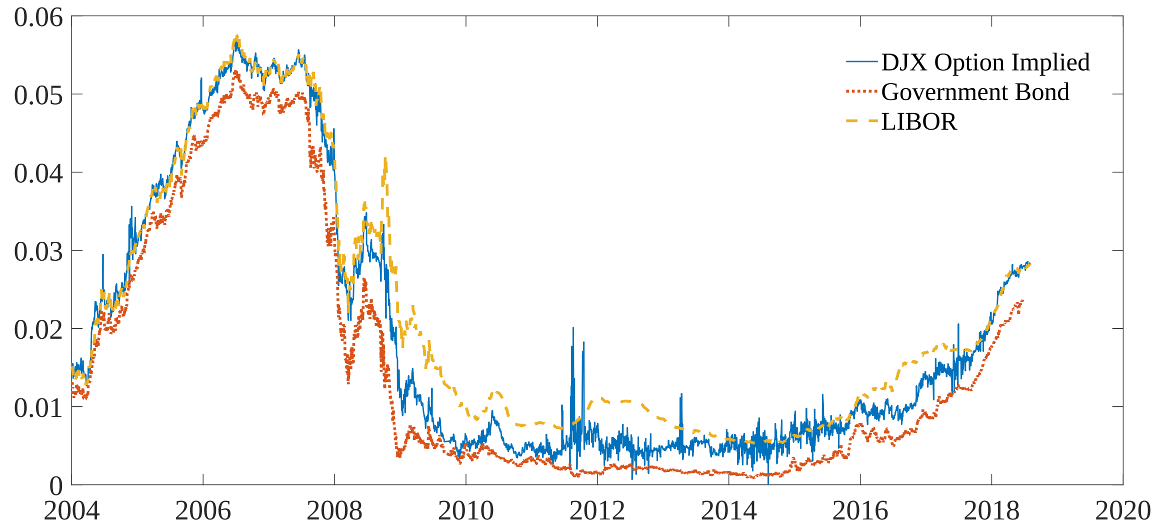

FIGURE XIII
Comparison of 1-year zero coupon interest rates implied from DJX options with government bond rates and LIBOR rates. All rates are continuously compounded.

over all the minute-level observations, those estimates will have smaller standard errors.

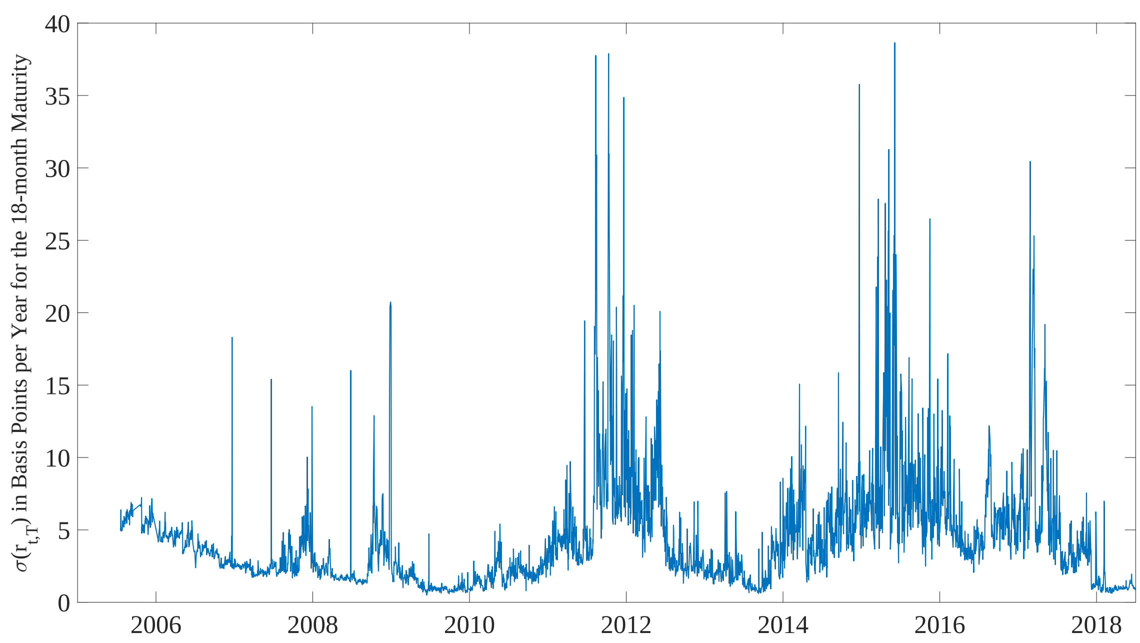

FIGURE XIV
Efficiency in the DJX option market expressed as the standard error around the implied risk-free rate.

Finally, we study how the interest rates implied by the DJX differ from those implied by the SPX. For each maturity, we compute a difference between the DJX and the SPX rate and we report the characteristics of that series in table 5.

Table 5
Difference between DJX and SPX Option Implied Interest Rates 2004-2018

Maturity	6-month	12-month	18-month
Mean	0.00046	0.00023	0.00021
Stdev	0.00224	0.00121	0.00103
AR(1) (daily)	0.4302	0.49587	0.5710

The table shows that while on average the rates are very close, substantial persistent daily deviations occur. As an illustration, Figure XV plots the differences between the two yields for the 12-month maturity.

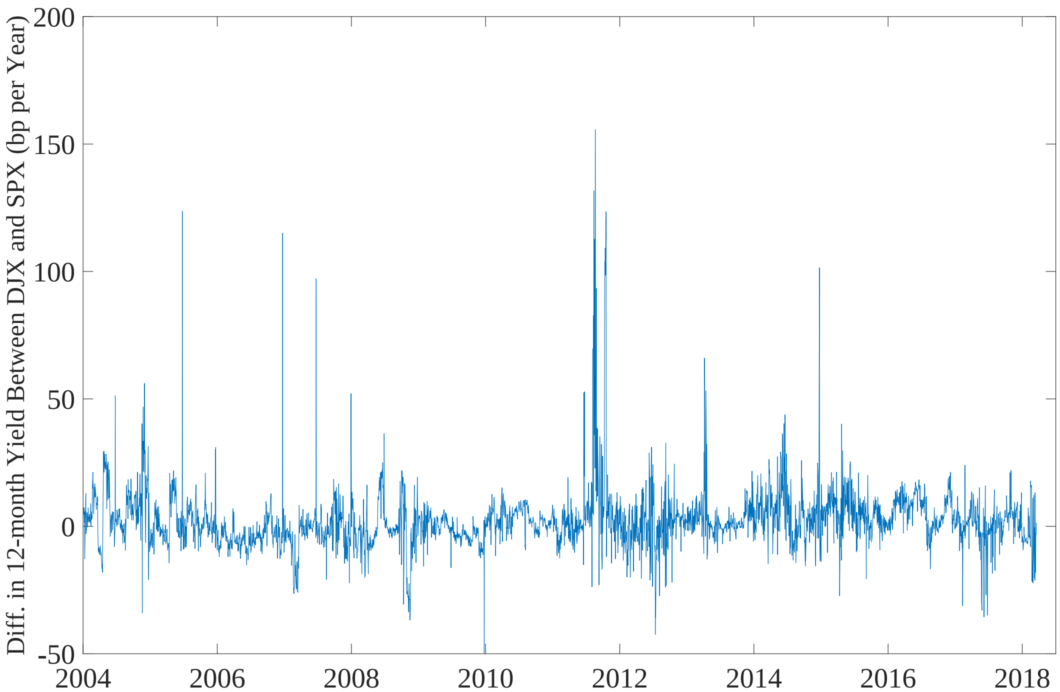

FIGURE XV
Difference in Daily Continuously Compounded Implied 12-month Zero Coupon Yield Between the DJX and the SPX in Basis Points per Year.

3.2. Combining Futures and Option Prices

In this subsection we combine data from SPX option and futures prices to assess the extent to which the rates implied from such an arbitrage trade would lead to similar rate estimates compared to those inferred from option prices alone. We find that the average rate inferred from this option-future arbitrage trade is nearly identical to our benchmark estimates. However, given the quality of the futures data, the estimates are substantially nosier making these implied rates less suitable for high-frequency studies.

The arbitrage trade that we explore uses the fact that the price level of the index net of the dividend strip price equals the discount value of the futures price with the same maturity T (Golez, Jackwerth, and Slavutskaya (2018)):

$$\begin{aligned} p_{i,t,T} - c_{i,t,T} &= (\mathcal{P}_{t,T} - S_t) + \exp(-r_{t,T}T)K_i && (14) \\ p_{i,t,T} - c_{i,t,T} &= -\exp(-r_{t,T}T)F_{t,T} + \exp(-r_{t,T}T)K_i. && (15) \end{aligned}$$

Rewriting this equation, we find:

$$r_{t,T} = -\frac{1}{T} ln \left(\frac{p_{i,t,T} - c_{i,t,T}}{K_i - F_{t,T}} \right) \tag{16}$$

For our SPX futures data we use the best-of-book trade and quote data (BBO) on E-mini futures from the CME. We first compute for each available minute and for each maturity the median futures price. We then match these futures prices with our minute-level option quotes and take for each minute the median of the estimated interest rates. To arrive at daily estimates, we then take the median across all minute-level estimates during the day.

The results are presented in Figure XVI using data between 2008 and 2018 (the sample for which we have high quality futures data). The graph shows that the option-future implied rate is a noisy version of the option-implied rate. The average rates are nearly identical (with an insignificant difference in means of about 3bp). In Table 6 we present the results of regressing the option-future implied rate (left-hand side variable) on the options-implied rate. As the table shows, the slope is 1 and the intercept not statistically different from 0. The R-squared of the regression is only 66.8% confirming that the option-future implied rate is a noisy version of the option-implied rate.

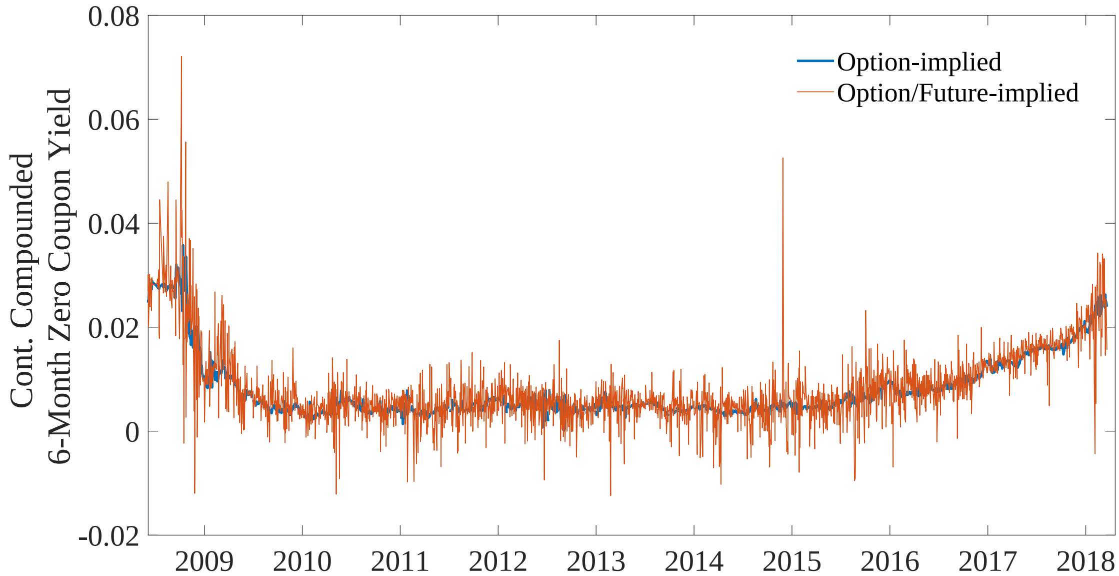

FIGURE XVI

The figure compares the estimated 6-month continuously compounded zero coupon yield from options alone to those inferred from combining options and futures.

Table 6

Regressing the Future-Option Implied Rate on the Option-Implied Rate

	Estimate	Standard Error
Intercept	0.0003	(0.0002)
Slope	1.0034	(0.0304)
R-squared	66.8%	

4. Robustness to Bid Ask Spreads

One potential concern with our option-implied rates is that frictions distinct from the convenience yield on safe assets could bias our estimates. For example, it is possible that our rates are systematically above or below true measures of option traders' risk-free cost of capital when the option market is particularly illiquid. While our SPX options do seem to be highly liquid and imply a very precise rate estimate (as measured by the cross-sectional R-squared values when estimating Equation 2 and the standard error measures in Figure XI), we demonstrate in this section that option market illiquidity does not induce bias in our estimated rates.

To compute our rates, we have followed the common practice of using midpoints of bids and asks (also used in the computation of the VIX index by the CBOE). If bid-ask spreads are symmetric around a true measure of an option's value, then an expanding bid ask spread would not change the value of the midpoint. However, if bid ask spreads expand and contract asymmetrically, they could induce bias in an option's midpoint value. In this section, we explore this issue, and document that observable measures of frictions (bid-ask spreads) in the option market do not seem related to our estimated rates. A first indication that such frictions are not influencing the level of our rates is that, on average, the DJX option market is substantially less liquid then the SPX option market (with wider bid ask spreads and substantially lower R-squared values) yet it produces essentially the same average level of interest rates (as reported in Table 4). Furthermore, and as argued before, we find similar convenience yields using the futures market for precious metals, suggesting that the only frictions of concern must be common across several markets.

We believe our usage of midpoint prices produces a more accurate measure of the risk-free cost of capital compared to other approaches that use bids or asks directly, or that use data on observed trades.[17] First, while there are non-trivial bid-ask spreads on individual options in our data, anecdotal evidence suggests that there are substantially smaller bid-ask spreads specifically for doing the Box trade. After all, a Box trader is unexposed to the underlying asset, so that market makers need not worry about asymmetric information regarding the underlying. Second, a trader posting limit orders or a trader who is well-connected in the trading network, can execute at substantially better prices than the reported bid and ask (Li, Wang, and Ye (2019)). Third, using trade data only would drastically reduce the cross-sectional R-squareds that we obtain and therefore yield much less precise rate estimates (yet similar unconditional means). We therefore take the approach of using midpoints and showing empirically that the resulting rate estimates are not biased by microstructure effects.

To refute the possibility of bias induced by option-market frictions, we study the time variation in the spread between the DJX-implied rate and the SPX-implied rate. Without such frictions, both rates would accurately estimate the risk-free cost of capital using only risky asset prices. As a result, fluctuations in the DJX-SPX spread reflect potential biases in both rates due to market-specific frictions. We demonstrate the lack of bias in the SPX-implied rate using two methods. First, we document that a positive (negative) DJX-SPX spread predicts negative (positive) changes in the DJX rate towards the SPX rate but the converse effect is an order of magnitude smaller. Second, we show that proxies for option microstructure frictions such as bid-ask spreads in each market predict almost no variation in the DJX-SPX spread. In our data, measures of frictions in the DJX and SPX option markets are only moderately correlated with each other, so we can observe the effects of

[17] We explore trade data in the Online Appendix and find that while using trades introduces noise in our estimates, it does not change the mean (and median) level of our rate estimates.

changing liquidity conditions separately in each market.

We proceed with the analysis as follows. First, we compute daily changes in the 12-month SPX-implied Box rate as well as daily changes in the 12-month DJX-implied rate and evaluate to what extent the spread between the two rates on the previous day predicts these changes. The results are reported in Table 7. The table shows that the DJX-SPX box spread negatively predicts changes in the level of the 12-month DJX-implied rate, indicating a convergence of the DJX-implied rate towards the SPX-implied rate. The R-squared of this predictive regression is 25%. While the DJX-SPX spread also positively predicts changes in the 12-month SPX rate, the coefficient is 14 times smaller in absolute magnitude (0.04 vs -0.54), and the explained variation (R-squared) is only 1.3%. It thus appears that the SPX-implied rate more accurately measures the risk-free cost of capital and that deviations of the DJX-implied rate from this more accurate estimate dissipate over time. Given that the DJX-SPX spread has a daily autocorrelation of 0.5 over our sample period, the DJX-implied rate seems to converge towards to SPX-implied rate not in a matter of hours or minutes, but rather in a matter of days.

Table 7
Predicting Daily Yield Changes with the DJX-SPX Spread

	$\Delta r^{DJX}_{t,1y}$	$\Delta r^{SPX}_{t,1y}$
Constant	0.0001	0.0000
t-stat	6.22	-0.84
$r^{DJX}_{t-1,1y} - r^{SPX}_{t-1,1y}$	-0.5444	0.0429
t-stat	-34.66	6.78
R^2	0.252	0.013

Next, we compute a daily bid-ask spread measure. We do this separately for the SPX and the DJX. For each minute, each maturity, and each strike, we compute the bid ask spread by taking the difference between the bid and the ask and dividing it by the mid price. We then construct a daily measure by taking the median over all those bid-ask spreads within the day. We do this separately for puts and calls, so that we can evaluate the influence of each option type in the regression. We then regress the daily spread between the 12-month DJX and SPX-implied interest rates on the natural logarithm of these bid-ask spread measures, so that the coefficients can be interpreted as the effect of roughly a doubling of the bid-ask spreads. The results are reported in Table 8.

Table 8
The Influence of Bid Ask Spreads on Estimated Rates

	$r^{DJX}_{t-1,1y} - r^{SPX}_{t-1,1y}$
Constant	0.0019
t-stat	6.51
ln(bid-ask-djx-call)	0.00026
t-stat	5.81
ln(bid-ask-djx-put)	0.00022
t-stat	4.90
ln(bid-ask-spx-call)	-0.00004
t-stat	-0.59
ln(bid-ask-spx-put)	0.00020
t-stat	3.37
R^2	0.042

The table shows that while there is some influence of the bid-ask spreads on the spread between the two rates, the effect is generally small. An increase of one unit in the log-bid-ask spread (roughly a doubling of the spread) leads to a change of at most 2.6 basis points in the spread. More importantly, the R-squared of the regression is only 4.2%. Because the standard deviation of the DJX-SPX spread (the dependent variable) equals 12 basis points, this R-squared value translates to a standard deviation of the fitted value of the regression of only 2.4 basis points. In Figure XVII we plot this fitted value together with the raw spread. The graph once again illustrates that very little of the DJX-SPX spread is related to movements in bid-ask spreads across the two markets.

Finally, we repeat the analysis above, but now using the LIBOR rate as the counterfactual interest rate. After all, in the period before 2007, when our rate is almost identical to LIBOR (and banks were not thought to be risky borrowers), the Box-LIBOR spread is another potential measure of the accuracy of our rate. After 2007, this spread is driven primarily by the credit risk of banks and is no longer an accurate proxy for the distortions we wish to measure. We therefore run regressions of the Box-LIBOR spread on our bid-ask spread

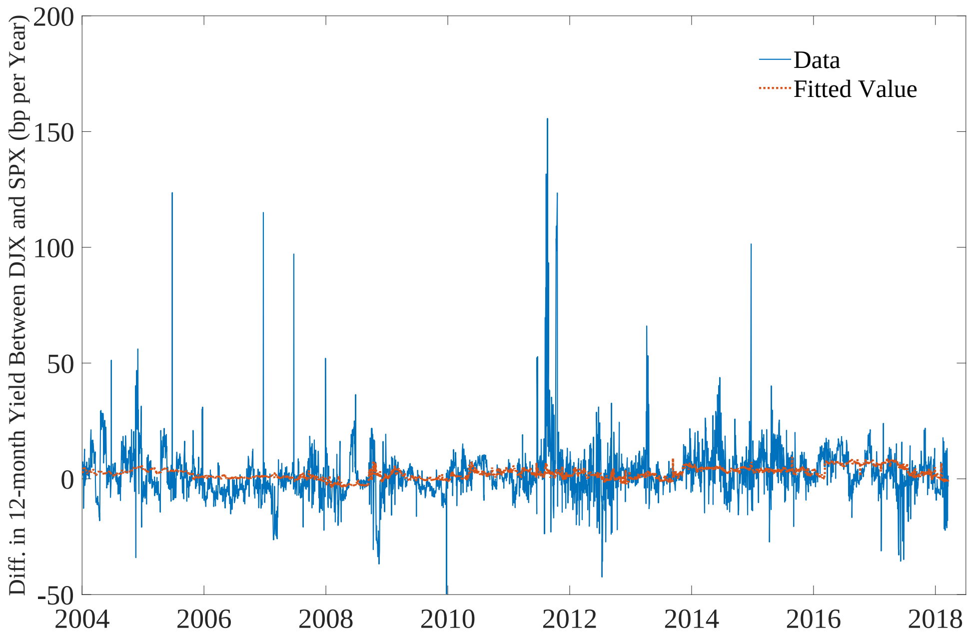

FIGURE XVII

Difference in Daily Continuously Compounded Implied 12-month Zero Coupon Yield Between the DJX and the SPX in Basis Points per Year. The graph also plots the fitted value that uses bid-ask spreads on puts and calls in both markets as the explanatory variables.

from 2004 to 2006 to test the hypothesis that option microstructure frictions can bias our Box rate estimates away from the risk-free cost of capital.

We summarize our results in Table 9. The left-hand side of the regression is the daily Box-LIBOR spread (1-year maturity) between January 1st 2004 and December 31st 2006, i.e. 755 daily observations, which we regress on the bid ask spread measures mentioned above. The regression results show once again that the option market frictions have a negligible explanatory power for the spread with a low R-squared value (2.3%). The coefficients on the log bid ask spread measures indicate that a doubling of bid ask spreads lowers our Box rate estimate by a few basis points. Given that put and call bid-ask spreads predict the Box-LIBOR spread with opposite signs, it is not even clear whether an overall decrease in option market liquidity predicts an increase or a decrease in the Box-LIBOR spread.

Table 9

The Influence of Bid Ask Spreads on Estimated Rates, Cont.

	$r_{t,1y}^{SPX} - r_{t,1y}^{libor}$
Constant	-0.00048
log (bid-ask-spx-calls)	-0.00029
t-stat	-3.19
log (bid-ask-spx-puts)	0.00010
t-stat	1.26
R-squared	0.023
Number of daily obs	755

5. Convenience Yields, Monetary Policy, and Q.E.

We now use our data to perform high-frequency event studies of the effects of monetary policy and quantitative easing on the term structure of convenience yields and Box rates. Our minute-level data is ideal for this purpose and broadens the set of questions that can be examined using high frequency event studies. Existing event studies on the effects of quantitative easing (Krishnamurthy and Vissing-Jorgensen (2011)) use two-day event windows because of issues related to slow price discovery. While price discovery in Treasury bonds themselves is quite fast, more illiquid bonds such as agency debt, corporate debt, or mortgage-backed securities require longer event windows. Because our Box rate estimates seem to have price discovery roughly as fast as Treasuries (see Section 3), we are able to measure spreads between different risk-free rates using a considerably shorter event window. There is a large literature on the high frequency effects of monetary policy on asset prices (Bernanke and Kuttner (2005), Rigobon and Sack (2004), Nakamura and Steinsson (2018)), but before our paper this literature has not presented results on convenience yields, arguably for similar reasons related to slow price discovery.

If liquidity is an important channel in the monetary transmission mechanism, we should find that monetary stimulus (either conventional and/or unconventional) reduces convenience yields. An idea going back to the LM curve of the IS-LM model (Hicks (1937)) that has been justified with recent empirical support (Nagel (2016)) is that the nominal interest rate measures the liquidity premium on assets such as cash and checking accounts (that pay no interest). As a result, an interest rate increase should make liquidity more scarce and increase the convenience yield on safe assets. Because our Box rate is inferred from risky assets which should have little to no convenience yield, we are able to decompose the effects of central bank policy on bond yields into changes in the risk-free cost of capital and changes in convenience yields. Our contribution is to present more direct evidence that monetary stimulus reduces convenience yields. Existing evidence mainly shows that nominal interest rates are correlated with spreads between different rates.

Our results on quantitative easing provide evidence for a broad transmission mechanism during the crisis, making progress on the state of knowledge in which Ben Bernanke said it "works in practice but not in theory." Because quantitative easing is the purchase of long term Treasury bonds and agency mortgage-backed securities financed by the issuance of bank reserves (which are a form of overnight debt), it is not clear whether its effects are determined by what is bought or what is sold. The seminal paper by Krishnamurthy and Vissing-Jorgensen (2011) presents empirical evidence on the transmission mechanisms of quantitative easing and discusses several possible channels. One view is that reducing the supply of Treasuries should make long duration safe assets more scarce and therefore increase their convenience yield. Under this "narrow" view of Q.E.'s transmission mechanism, the yield on the specifically purchased asset is disproportionately affected by the policy. A "broad" view of the transmission mechanism of quantitative easing is common in much of the theoretical literature (Caballero and Farhi (2018) and Diamond (2018)), which emphasizes that swapping reserves for long duration assets increases the overall supply of safe assets and should therefore impact the prices of untargeted assets. Because our Box rates are inferred from equity option prices, an asset class quite distinct from the fixed income market, our data is well-suited to measure the spillover of Q.E. to untargeted asset classes. We find substantial broad spillover effects from Q.E. 1 but not from Q.E. 2 or Q.E. 3. This is complimentary to the evidence in Krishnamurthy and Vissing-Jorgensen (2013) that Q.E. 3 had important "narrow" effects on agency MBS yields.

5.1. Effects of Quantitative Easing

Our results on the effects of quantitative easing follow a literature which has identified specific dates and times at which policymakers conveyed news about their intention to increase or decrease the size of the program. For the first two rounds of the program, which occurred respectively in 2008/2009 and in 2010, we use the same dates as Krishnamurthy and Vissing-Jorgensen (2011). For Q.E. 3, which happened after the aforementioned paper, we follow the dates in Di Maggio, Kermani, and Palmer (2019). The five event dates for Q.E. 1 are 11/25/2008/, 12/1/2008, 12/16/2008, 1/28/2009 and 3/18/2009. For Q.E. 2 we consider the event dates 8/10/2010, 9/21/2010, and 11/3/2010, and for Q.E. 3 we consider the event dates 9/13/2012, 5/22/2013, 6/19/2013,7/10/2013, and 9/18/2013. For each date we have precise time stamps of the event.

To illustrate the quality of our minute-level data, we plot in Figure XVIII the minute-level Box rates on March 18th 2009, for three different maturities. The time of the Q.E. announcement was at 2.15pm. The picture clearly shows the effect of the announcement for all three maturities and particularly for the shortest maturity, it seems that rates started moving before the actual announcement.[18]

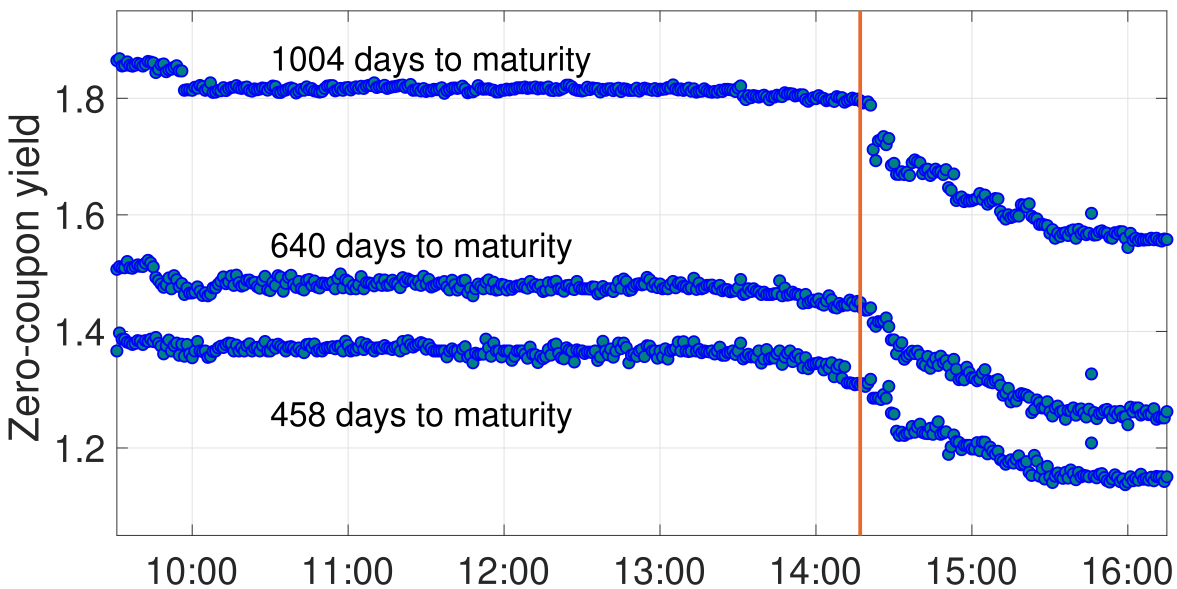

FIGURE XVIII

S&P500 Box Rates on March 18, 2009: The figure plots the minute-level box rates for three different maturities on March 18, 2009. The vertical line represents the time of the release (14:17).

To analyze the effect across all Q.E. dates, we take the median yield on every asset in a window 30 to 60 minutes before the time stamp and 30 to 60 minutes after the time stamp.[19] We then fit Nelson-Siegel-Svensson yield curves to these median yields before and after each event. Specifically, we fit one yield curve to our intraday SPX Box rates and a second yield curve to intraday indicative quotes on Treasury yields from GovPX. For all quotes we use a midpoint of bid and ask.

To summarize our results, we find that both monetary policy and quantitative easing

[18] See also Cieslak, Morse and Vissing-Jorgensen (2019).

[19] This is not possible for 2 of our time stamps, since they occur too early in the day. On 11/25/2008, the time stamp is before the start of trading, so we use the median yield in the last 30 minutes of the previous day and the median yield in the first 30 minutes of the day. On 5/22/2013, we use the median yield between 0 and 30 minutes before the 10am time stamp and 60 to 90 minutes after.

have quite strong effects on convenience yields during the worst of the financial crisis (the second half of 2008 and first half of 2009) but considerably more modest effects otherwise. We report in the figures below the effects of the central bank policies on 3-month, 12-month, and 30-month yields. We report results on Treasury yields, Box yields, and the convenience yield (which equals their difference). The maturities of 3 and 30 months are the most extreme durations for which we can present results without extrapolating beyond where our data lies.

We find that for Q.E. 1 (i.e., the first round of quantitative easing) which occurred between November 2008 and March 2009, Box yields fell considerably more than Treasury yields. In Figure XIX below, we show that 12 and 30 month Box yields fell by 88 and 86 basis points respectively, while Treasury yields of the same maturity only fell by 46 and 61 basis points. This results in a reduction in 12- and 30-month convenience yields of 42 and 25 basis points. At the shorter 3-month maturity, government yields fell by only 2 basis points while Box yields fell by 37 basis points rounding to a 36 basis point reduction in the convenience yield. The lack of response in short-term Treasury rates is likely because those rates were already at the zero lower bound, while all Box rates were considerably higher than Treasury yields at this time. The greater drop in box than Treasury yields provides evidence in favor of a broad transmission of quantitative easing, in which asset prices outside of narrowly defined fixed income markets also respond. Because risky assets are priced without the convenience yield (that is, consistently with our Box rate rather than Treasury rates), this implies that quantitative easing reduced investors' cost of capital by even more than is suggested by the drop in Treasury yields. It also implies that Q.E. 1 can be thought of as an increase in the supply of safe assets, by swapping more scarce reserves for less scarce Treasuries or agency mortgage-backed securities.[20] In particular, Q.E. 1 reduced the convenience-yield-free short-term interest rate. If this rate solves the consumption Euler equation, Q.E. 1 therefore stimulated aggregate demand by the same mechanism as conventional monetary policy.[21] This relative scarcity is consistent with our finding that the convenience yield is largest at the shortest maturities.

For Q.E. 2 and 3, we find considerably smaller effects on Treasury yields and effects with ambiguous signs on the convenience yield on safe assets as reported in Figure XX. Summing up across all 8 event dates in this period, we find that 3, 12, and 30-month Treasury yields fell by 3, 4 and 11 basis points respectively. At the same time, 3, 12 and 30-month Box yields fell by 6, increased by 3, and fell by 8 basis points respectively. This leads to a 3 basis point

[20]Because our longest maturity is only 30 months, we are unable to measure any specific scarcity of long-term safe assets, as documented by Krishnamurthy and Vissing-Jorgensen (2011).

[21]This transmission mechanism is analyzed in Caballero and Farhi (2018); Diamond (2018) and our empirical findings confirm predictions of these models.

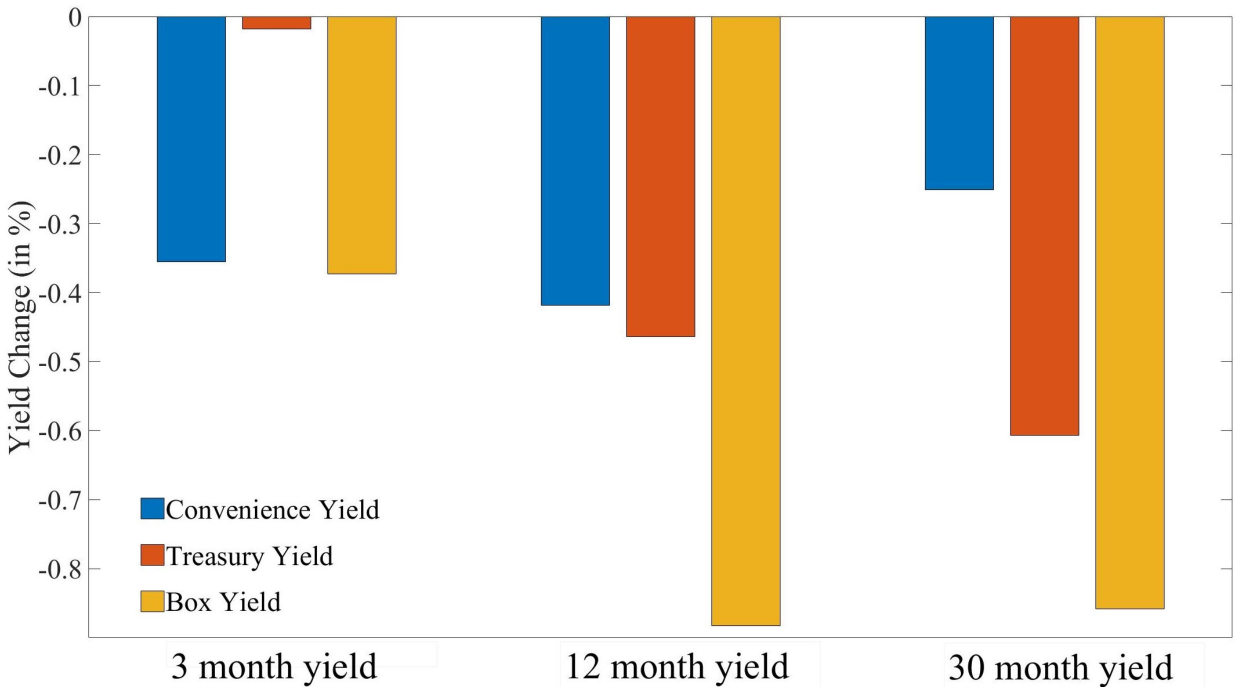

FIGURE XIX

Cumulative effect of Q.E. 1 on government bond yields, Box yields, and convenience yields (i.e. the difference between the two) across various different maturities.

decrease in the 3-month convenience yield and a 7 and 3 basis point increase in the 12- and 30-month convenience yields. The aggregate effect of all Q.E. 2 and 3 announcements is of considerably smaller magnitude than the effect of Q.E. 1. In particular, if anything, it seems to increase the convenience yield on Treasuries, though the effect is small and of ambiguous sign across the yield curve.

One possible explanation of our results is that quantitative easing after 2009 was performed outside of the depths of the financial crisis, at which point convenience yields had already converged back to normal levels. It may be that quantitative easing is a weaker policy tool when the financial system is not in distress. Another possible explanation is that the news in this sample on average did not surprise investors as much, with event days including both news that increased and decreased investors' expectations about the size of the program. Regardless of the explanation, it is immediately clear that the large effects found in Q.E. 1 do not seem to generalize to this extension of the program after the depths of the crisis.

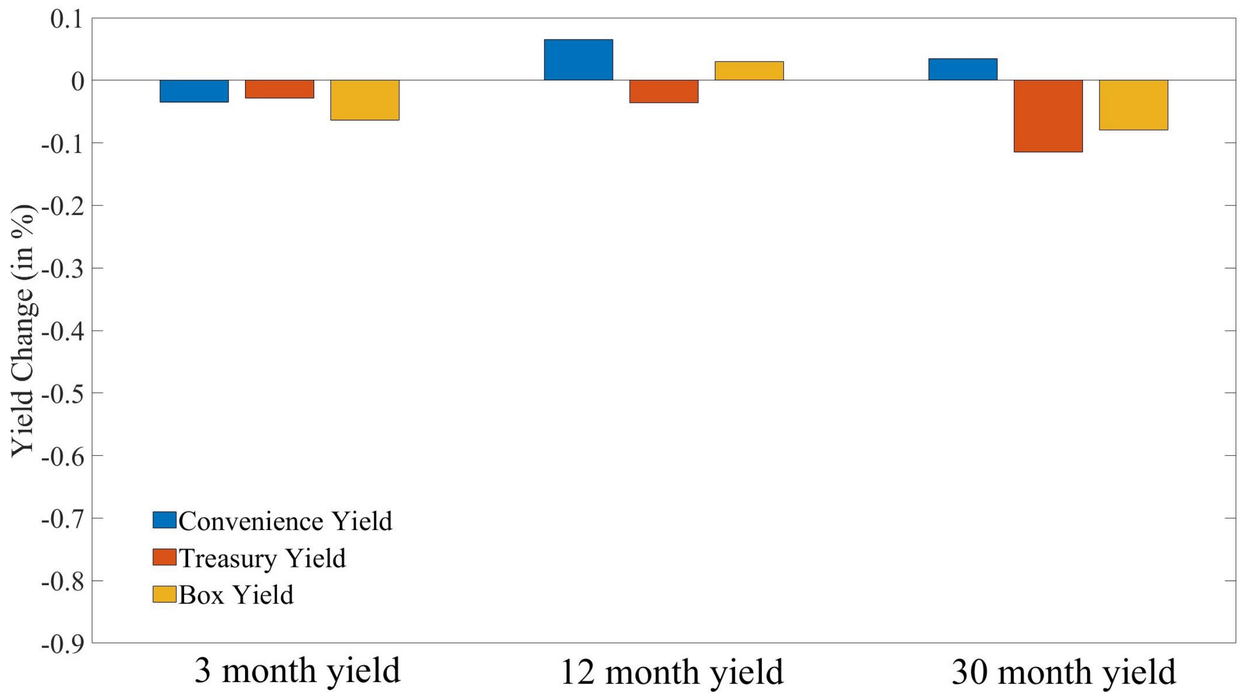

FIGURE XX
Cumulative effect of Q.E. 2 and 3 on government bond yields, Box yields, and convenience yields (i.e. the difference between the two) across various different maturities.

5.2. Monetary Policy Event Studies on FOMC announcement dates

To study the effect of conventional monetary policy on convenience yields, we perform a high frequency event study using all FOMC announcements from 2004 to 2018, the time period in which we have Box rate data. We measure unanticipated shocks to monetary policy using innovations in Federal Funds futures from the Chicago Mercantile Exchange (CME) around each FOMC announcement. Our measure of a monetary policy shock is analogous to that of (Bernanke and Kuttner (2005) and Nakamura and Steinsson (2018)). We use the first trade more than 10 minutes before and the first trade more than 20 minutes after each announcement to compute our monetary surprise. Given this monetary policy shock, we fit yield curves to GovPX Treasury quotes and our Box yields in windows 30 to 60 minutes around each announcement. We then regress the change in each yield around an announcement on our measure of the monetary shock associated with that announcement and use our estimated regression coefficient to predict the effects of a 100 basis point surprise increase in the Federal Funds rate when reporting our results below.

Similar to our quantitative easing results, we find that monetary policy has considerably stronger effects on convenience yields in the depths of the crisis than at other times. In Figure XX below we show the results for the whole sample. A 100 basis point rate increase leads to a 54, 88, and 74 basis point increase in Box yields for the 3, 12 and 30-month maturities respectively. It leads only to 52, 63, and 45 basis point increases in the 3, 12, and 30-month Treasury yields. This results in an increase of the convenience yield of 2, 26, and 28 basis points respectively. This implies that particularly at longer maturities, an increase in the Federal Funds rate leads to increases in the convenience yield that are more than a third the size in the increase in Treasury yields. Similar to quantitative easing, the effect of monetary policy spills over to unrelated asset classes like equity index options.

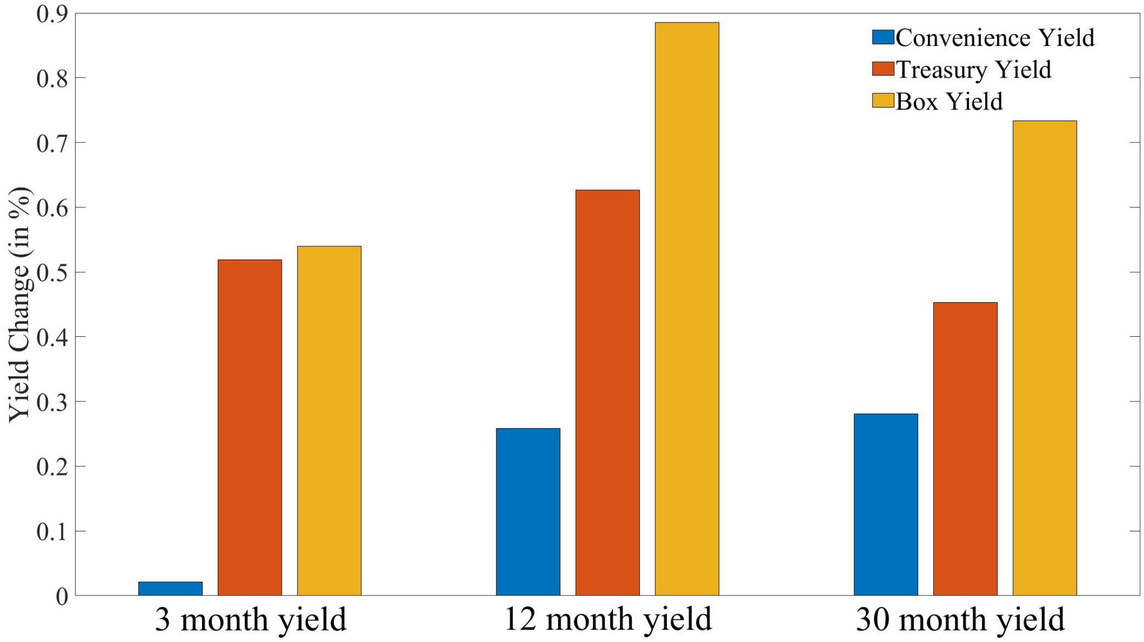

FIGURE XXI

Effect of FOMC Announcements on government bond yields, Box yields, and convenience yields (i.e. the difference between the two) across various maturities over the sample 2004-2018.

Next, we present results in Figure XXII from an identical event study but ignoring data from the second half of 2008 and first half of 2009. The results change considerably. Like before, we find that Treasury yields respond quite strongly to monetary policy. A 100 basis point increase in the fed funds rate leads to 70, 55, and 34 basis point increases in the 3, 12, and 20 month Treasury yields. However, there is only a 32, 66, and 46 basis point increase in the 3, 12, and 20 month Box rate. This leads to a 38 basis point decrease in the 3 month convenience yield and a 11 and 13 basis point increase in the 12 and 30 month convenience yields. It therefore seems that by simply removing one year of the worst of the financial crisis from the data, the results imply a considerably weaker (and ambiguously signed) effect of monetary policy on convenience yields. While this evidence does clearly show that monetary policy shocks pass through to interest rates implicit in option prices, there is no clear evidence that monetary policy shocks impact our convenience yield measure outside of a severe financial crisis. This is consistent with our evidence on Q.E., where only the first round of Q.E. in 2008-2009 had important effects on our convenience yield measures.

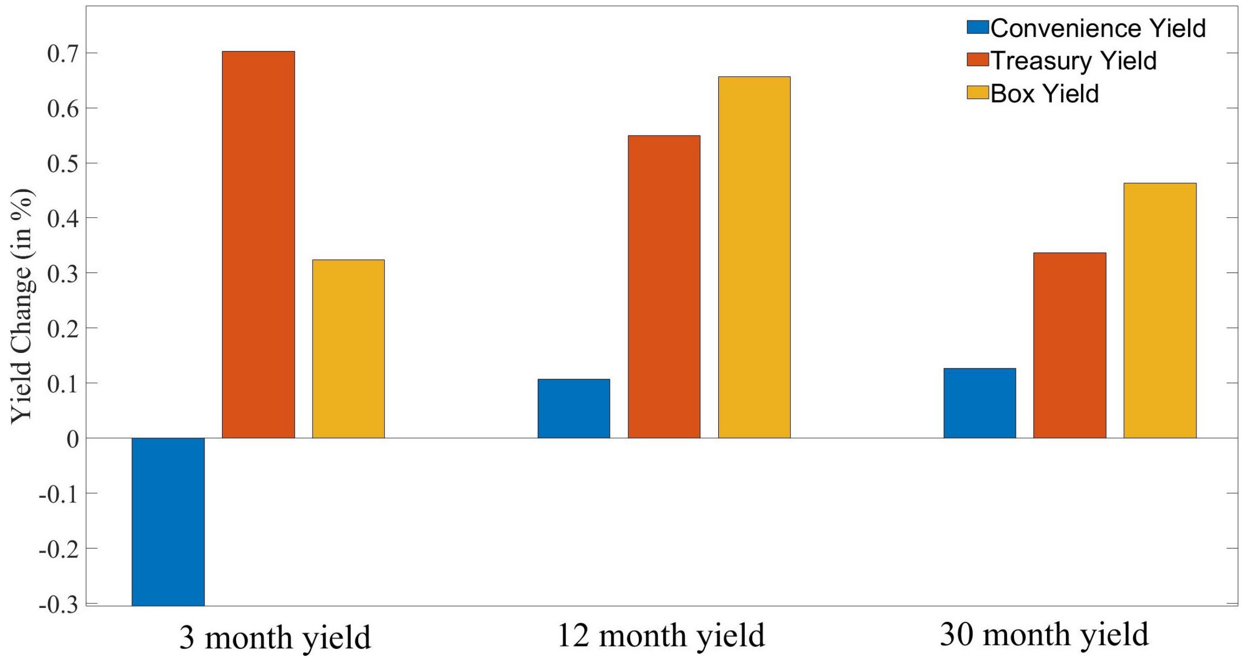

FIGURE XXII

Effect of FOMC Announcements on government bond yields, Box yields, and convenience yields (i.e. the difference between the two) across various maturities for the sample 2004-2018 but excluding the crisis period.

6. Bond Return Predictability

A literature as early as Fama and Bliss (1987) has focused on the predictability of government bond excess returns using information contained in the term structure of bond returns. This predictability has been one of the more difficult empirical findings to square with asset pricing theory, particularly given the seeming disconnect of these time varying expected excess returns from the documented variation in expected excess return in stock markets.

Because of our unique term structure of convenience yields, we can decompose government bond returns into movements in the option-implied risk-free cost of capital and convenience yields. In particular, we have defined the convenience yield $cy_{t,n}$ as the difference between the implied (continuously compounded) yield inferred from S&P500 options and the yield on government bonds:

$$cy_{t,n} = y_{t,n}^{box} - y_{t,n}^{gov}, \tag{17}$$

where n is the time until maturity. Rewriting this equation we find:

$$y_{t,n}^{gov} = y_{t,n}^{box} - cy_{t,n}. \tag{18}$$

The excess return on government bonds, which is given by:

$$rx_{t+1,n}^{gov} = ny_{t,n}^{gov} - (n-1)y_{t+1,n-1}^{gov} - y_{t,1}^{gov} \tag{19}$$

can then be written as the difference between two return components, the one related to changes in the Box rate and the one related to changes in the convenience yield:

$$rx_{t+1,n}^{gov} = ny_{t,n}^{box} - (n-1)y_{t+1,n-1}^{box} - y_{t,1}^{box} - ncy_{t,n} + (n-1)cy_{t+1,n-1} + cy_{t,1}. \tag{20}$$

This then naturally raises two questions. First, to what extent is the predictability in government bond returns related to each of these two components? Is it driven by predictable variation in excess Box returns, or predictable variation related to the convenience yield component of returns? Second, is the predictive power in current yields due to the component due to convenience yields or the component due to the option-implied risk-free cost of capital?

To provide a first answer to these two questions, we use the approach proposed by Cochrane and Piazzesi (2005) and use the cross-section of yields to construct a return forecasting factor. We construct two such factors. The first replicates the one of Cochrane and Piazzesi (2005) for the 2004-2018 sample. For the construction of the second forecasting factor, we follow the exact same procedure (using the same left-hand side variables), but instead of using as the forecasting variables the cross-section of government bond yields, we use the cross-section of convenience yields. We then evaluate the forecasting power of these two factors individually and jointly using excess returns on government bonds and excess returns on Box rates focusing on the 2-year maturity only (we do not have longer maturity claims available for the Box rate).

The results are summarized in the table below. The results show that the factor constructed from the convenience yield substantially predicts both government bond returns as well as Box rate returns. In the joint regression, both the convenience yield factor and the Cochrane Piazzesi factor show up significantly. In traditional asset pricing models where the consumption Euler equation prices all assets, there is no convenience yield and thus no predictability resulting from it. Overall, the results therefore suggest that a complete explanation of bond return predictability requires a model that features both time varying risk aversion (or volatility) as well as a time varying convenience yield. Decomposing fluctuations in the convenience yield into time varying demand for safe assets, shocks to intermediary balance sheets, and demand fluctuations for derivatives such as options is a natural question for future work.

	$rx^{gov}_{t+1,2}$	$rx^{gov}_{t+1,2}$	$rx^{gov}_{t+1,2}$	$rx^{box}_{t+1,2}$	$rx^{box}_{t+1,2}$	$rx^{box}_{t+1,2}$
β^{CP}	0.299***		0.196***	0.357***		0.204***
β^{CY}		0.415***	0.329***		0.580***	0.489***
Adj. R^2	0.231	0.319	0.403	0.258	0.488	0.560

7. The Box-Rate-Implied CIP Deviation

One important no-arbitrage relation that has received increasing attention in recent years, is the so-called Covered Interest Parity (CIP) relationship. CIP states that the ratio of an exchange rate's forward and spot rate equals the ratio between the nominal gross interest rates in the two countries. Recently, however, Du et al. (2018a) and Du et al. (2018b) have shown large, and persistent violations of CIP using government bond and LIBOR yields as the measure for the nominal interest rate. As argued previously, government bond yields feature a convenience yield and LIBOR contains credit risk. As a consequence, it is informative to compute a measure of CIP deviations using our convenience-yield-free risk-free interest rates.

Take a US-based agent at time t facing two alternative strategies. He can either invest in a riskless asset denominated in dollars with n years to maturity (in our case the U.S. Box rate), or exchange money into a foreign currency, invest it into the riskless asset denominated in that currency for n years (the Box rate constructed for the foreign country) and buy a promise to exchange the money back into dollars at a predetermined rate at time $t+n$.

More formally, denote with $r_{t,n}$ and $r^*_{t,n}$ the continuously-compounded Box rates at time t with n-year maturity for the domestic and foreign country. The CIP relation that we are exploring in this section is then given by:

$$e^{nr_{t,n}} = e^{nr^*_{t,n}} \frac{S_t}{F_{t,n}}, \qquad (21)$$

where S_t is the time-t spot exchange rate between dollars and foreign currency and $F_{t,n}$ the forward rate of exchange, set at time t with a n-year maturity.

We construct the so-called cross-currency basis, in logs, as

$$x_{t,n} = r^*_{t,n} - r_{t,n} - \frac{1}{n}\ln(S_t/F_{t,n}) \tag{22}$$

The last term $(\frac{1}{n}\ln(S_t/F_{t,n}))$ is the annualized continuously-compounded "forward premium".

We analyze the CIP deviation between the U.S. dollar and the Japanese yen. The data on futures trades are obtained from the Chicago Mercantile Exchange (CME). Spot exchange-rate quotes are from the TrueFX dataset, which offers historical tick-by-tick market data for interbank foreign exchange rates at the millisecond frequency. For spot exchange rate quotes, we take the mid-point between bid and ask rates. We compute the median spot and the median forward rate every minute, and match the spot and forward rates in the same minute. We construct the forward premium, and compute the daily median.

To compute the Japanese Box rate, we use European options on the Nikkei 225 index provided by the Japan Exchange Group. Each observation corresponds to a quote in response to a new order. We then construct the mid quote at the minute level. In particular, we consider the best bid and the best ask each minute, and use only those minutes in which both an order to sell and to buy are submitted. Finally, we restrict our attention to those minutes, maturities, and strikes for which the minimum ask price is not larger than 1.5 times the maximum bid price. The rest of the procedure for constructing the Box rates from mid quotes follows the one outlined in Section 2. Per each maturity and date, we compute a daily median.

We then linearly interpolate Box rates on both currencies to match the maturity of the forward contract. Figure XXIII depicts the cross-currency basis in bps implied by Box rates, and the ones computed by Du et al. (2018a,b). These authors use U.S. and Japanese government bonds and LIBOR respectively to construct their cross-currency bases. To make our results comparable to existing work, we use the forward contract with the maturity closest to 90 days. The average maturity for our series is 97 days. The figure shows that the box-rate-implied CIP deviation is almost always smaller (closer to zero) than previous estimates. The average value of the cross-currency basis we calculate is 41 bps relative to the 79 bps for the Du et al. (2018a) series in the same period. The difference of 38 bps is driven by two effects: (1) a more precise estimate of the forward premium (daily median of minute-

FIGURE XXIII

CIP Deviation: Box rate, government bond, and LIBOR

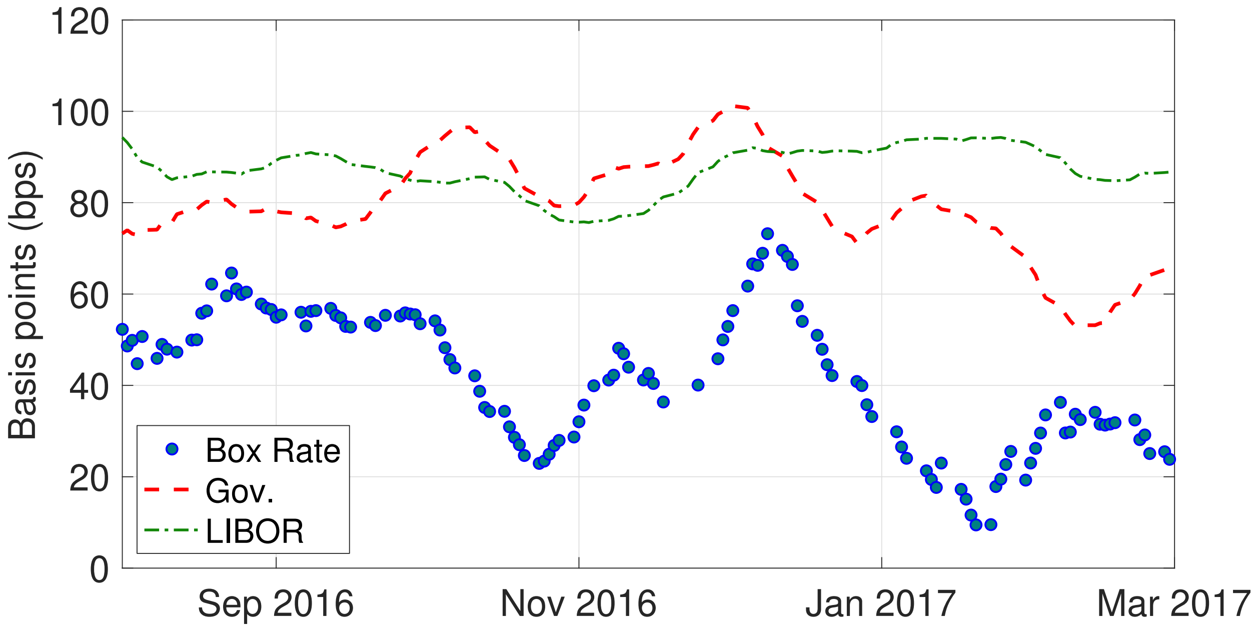

FIGURE XXIV

The figure shows a ten-day rolling moving average of the Dollar-Yen cross-currency basis implied by Box rates, as well as those implied by government bonds and LIBOR presented in Du et al. (2018a) and Du et al. (2018b). The maturity equals 3 months.

level forward premia instead of end-of-the-day values), and the usage of Box rates rather than government bonds. In the same period, the CIP violation computed using LIBOR rates is 87.5 bps. This exercise shows that previous estimates of CIP deviations reflected the greater convenienec of Dollar over Yen-denominated safe assets. The remaining spread may reflect demand imbalances (as in Hazelkorn et al. (2020)) for borrow and lending in different currencies driven by a desire to exploit the carry trade.

8. Price Discovery

As a last robustness analysis, we examine the speed of price discovery for Box rates and compare it to the speed in Treasury markets. In particular, we investigate how fast rates converge to a new stable level that incorporates news in Federal Open Market Committee (FOMC) announcements. Previous research (Fleming and Piazzesi (2005)) has used Treasury yield data from GovPX to show that government bond yields respond quickly to such announcements. We demonstrate that the speed of convergence in our minute-level Box rate data is as fast, if not faster, than the convergence in Treasury yields. The advantage of our Box rates is that we have quotes at every minute for all maturities while GovPX is irregularly spaced with frequent gaps. For this reason, to appropriately compare our high-frequency rates with the Treasury security tick data from GovPX, we first select the government bonds with the highest market activity in a given FOMC meeting day. We then match these bonds with the closest Box rates in terms of maturity with a maximum difference in maturities of 30 days.

Our price-discovery exercise closely mimics Fleming and Piazzesi (2005). As in Section 5, we derive policy surprises from the prices of fed funds futures traded on the Chicago Board of Trade (CBOT). We compute policy surprises as the difference between the yield implied by the last trade executed at most 10 minutes before the announcement release and the first trade made at least 20 minutes after the announcement. Following Kuttner (2001) and Nakamura and Steinsson (2018) we define the policy surprise to be the innovation in the current-month futures rate scaled up to reflect the number of days left in that month. For announcements made in the last seven days of the month we use the next month's futures contract.

We compute changes in both Box and Treasury rates in 5-minute intervals around the announcement release. Figure XXV reports the average 5-minute absolute yield changes (volatility) for all matched securities and all FOMC announcement days. Both Treasury yields and Box yields appear to adjust immediately at the time of announcement. Furthermore, a slightly higher volatility seems to persist in the hours after the announcement. These volatility patterns provide evidence that both rates respond with similar speed to news.[22]

To calculate the speed of convergence, we then proceed by regressing rate changes over various time intervals around announcements on the Fed Funds target rate surprises. For both the government and the Box rates, the largest responses occur in the interval including the announcement release. However, unlike the Box rate, the government bond yields exhibit a sluggish response to target rate surprises consistent with the evidence documented by Fleming and Piazzesi (2005). The regression estimates yield additional support for the usage of the Box rate when evaluating market behavior at higher frequencies.

9. Conclusion

We have constructed and analyzed a novel panel of risk-free interest rates that are free of the convenience yield on safe assets. We have presented three important applications of this

[22] In Online Appendix C we present several interesting individual cases.

Table 10
Effects of Fed funds rate surprises on yields around FOMC announcements

	Interval of Analysis				
Security	$(-45,-25)$	$(-25,-5)$	$(-5,25)$	$(25,55)$	$(55,80)$
Box rate	0.034	-0.035	1.196***	0.354**	0.080
	(0.065)	(0.081)	(0.209)	(0.170)	(0.087)
Government	0.010	-0.020	0.461**	0.187	0.165***
	(0.039)	(0.034)	(0.205)	(0.122)	(0.053)

Notes: The table shows the results from regressing the innovations in Box rates or government bond yields on Fed Funds rate surprises for various intervals around the announcement release. The Fed Funds surprise is the variation in the current-month futures rate from the last trade executed at most 10 minutes before the announcement to the first trade made at least 20 minutes after the announcement.

FIGURE XXV
Yield volatility on FOMC announcements

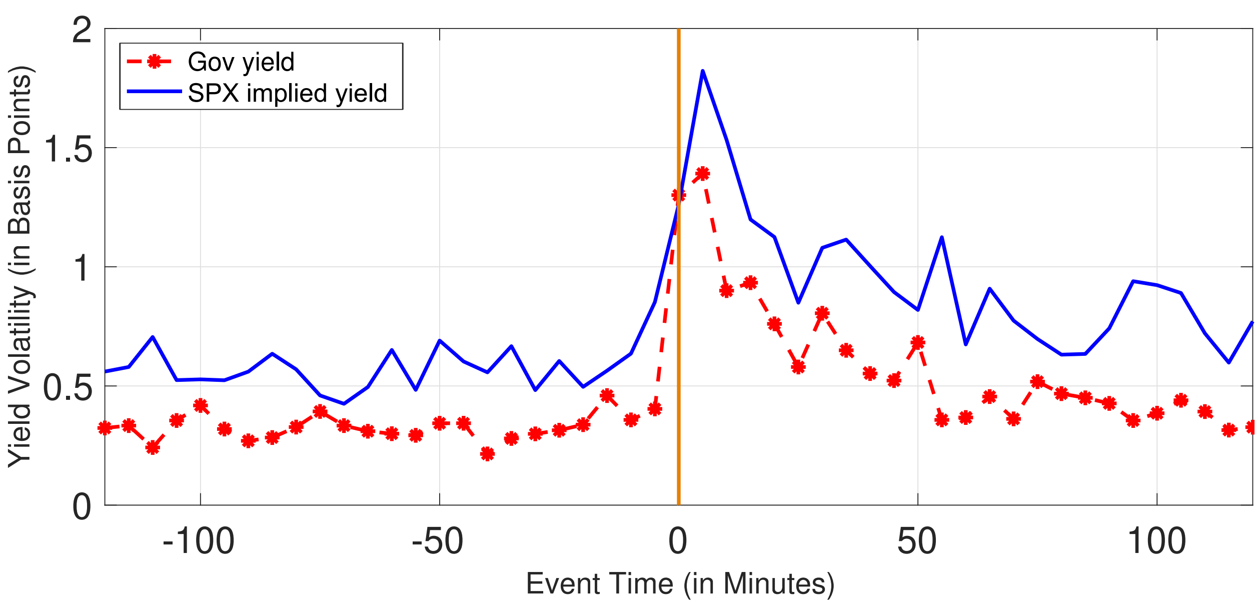

The figure shows the average absolute rate changes around FOMC announcement releases.

novel data set: (1) high frequency event studies of the effects of central bank policy, (2) the role of convenience yields in bond return predictability, and (3) constructing a convenience-yield free measure of Covered Interest Parity (CIP) deviations in foreign exchange markets. More generally, we wish to advocate for our rates' widespread use in the macro, monetary, international, and financial economics literatures. For example, our data is important for the accurate measurement of risk premia on stocks and credit instruments in asset pricing, as it prevents researchers from inadvertently confusing the convenience yield on safe assets with compensation for risk. In addition, our rates allow monetary economists to isolate monetary transmission mechanisms that flow specifically through the scarcity of safe/liquid assets.

One issue we have not yet explored is if the risk-free cost of capital in options markets differs from the time value of money for households, which would imply that the convenience yield on a safe asset depends on which investor buys it. Understanding how the costs of capital and convenience yields provided by safe assets varies across investors is a promising direction for future research. Another interesting avenue is to further identify the mechanisms underlying our documented spread (convenience yield). Given that our spread is a difference between option-implied rates and Treasury yields, it would be interesting to further decompose the variation into fluctuations in option (leverage) demand (Garleanu et al. (2009) Hazelkorn et al. (2020)) and Treasury demand (Krishnamurthy and Vissing-Jorgensen (2012) and Nagel (2016)).

References

Amihud, Y., Mendelson, H., 1991. Liquidity, maturity, and the yields on u.s. treasury securities. Journal of Finance 46, 1411–1425.

Andersen, L., Duffie, D., Song, Y., 2019. Funding Value Adjustments. Journal of Finance 74, 145–192.

Avino, D., Stancu, A., Wese Simen, C., 2017. Option Implied Discount Rates. Tech. rep., working Paper.

Bansal, R., Coleman, W. J., 1996. A Monetary Explanation of the Equity Premium, Term Premium, and Risk-free Rate Puzzles. Journal of Political Economy 104, 1135–1171.

Bernanke, B., 1990. Clearing and settlement during the crash. Review of Financial Studies 3, 133–151.

Bernanke, B., Kuttner, K., 2005. What Explains the Stock Market's Reaction to Federal Reserve Policy? Journal of Finance 60, 1221–1257.

Binsbergen, J. H. v., 2020. Duration-based stock valuation: Reassessing stock market performance and volatility, unpublished paper, The Wharton School, University of Pennsylvania.

Black, F., 1972. Capital market equilibrium with restricted borrowing. Journal of Business 45, 444–454.

Brenner, M., Galai, D., 1986. Implied interest rates. The Journal of Business 59, 493–507.

Brolley, M., Malinova, K., 2018. Informed trading in a low-latency limit order market. Working paper, Wilfrid Laurier University.

Caballero, R., Farhi, E., 2018. The Safety Trap. Review of Economic Studies 85, 223–274.

Cochrane, J. H., Piazzesi, M., 2005. Bond risk premia. American Economic Review 95, 138–160.

Di Maggio, M., Kermani, A., Palmer, C., 2019. How Quantitative Easing Works: Evidence on the Refinancing Channel. Tech. rep., working Paper, Harvard Business School.

Diamond, W., 2018. Safety Transformation and the Structure of the Financial System. working paper .

Du, W., Im, J., Schreger, J., 2018a. The U.S. Treasury Premium. Journal of International Economics 112, 167 –181.

Du, W., Tepper, A., Verdelhan, A., 2018b. Deviations from covered interest rate parity. The Journal of Finance 73, 915–957.

Fama, E. F., Bliss, R. H., 1987. The information in long-maturity forward rates. American Economic Review 77, 680–692.

Faruqui, U., Huang, W., Takts, E., 2018. Clearing risks in otc derivatives markets: the ccp-bank nexus. Bank of International Settlements Quarterly Review .

Fleming, M. J., Piazzesi, M., 2005. Monetary policy tick-by-tick. Working paper, Stanford University.

Frazzini, A., Pedersen, L., 2014. Betting against beta. Journal of Financial Economics 111, 1–25.

Garleanu, N., Pedersen, L. H., Poteshman, A. M., 2009. Demand-Based Option Pricing. Review of Financial Studies 22, 4259–4299.

Golez, B., Jackwerth, J., 2020. Quarter-of-a-century of dividend pricing, unpublished paper, University of Notre Dame.

Golez, B., Jackwerth, J., Slavutskaya, A., 2018. Funding illiquidity implied by s&p 500 derivatives. Working paper.

Gorton, G., Pennacchi, G., 1990. Financial Intermediaries and Liquidity Creation. Journal of Finance 45, 49–71.

Gürkaynak, R. S., Sack, B., Wright, J. H., 2007. The U.S. Treasury yield curve: 1961 to the present. Journal of Monetary Economics 54, 2291 – 2304.

Hazelkorn, T. M., Moskowitz, T. J., Vasudevan, K., 2020. Beyond basis basics: Liquidity demand and deviations from the law of one price. Working paper.

He, Z., Krishnamurthy, A., 2013. Intermediary asset pricing. American Economic Review 103, 732–770.

Hébert, B., 2018. Externalities as arbitrage. Working paper, Stanford University.

Hicks, J., 1937. Mr. keynes and the 'classics': A suggested interpretation. Econometrica 5, 147–159.

Krishnamurthy, A., Vissing-Jorgensen, A., 2011. The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy (Brookings Papers on Economic Activity, 2011, No. 2). Brookings Papers on Economic Activity pp. 215–287.

Krishnamurthy, A., Vissing-Jorgensen, A., 2012. The Aggregate Demand for Treasury Debt. Journal of Political Economy 120, 233–267.

Krishnamurthy, A., Vissing-Jorgensen, A., 2013. The ins and outs of large scale asset purchases. Kansas City Federal Reserve Symposium on Global Dimensions of Unconventional Monetary Policy .

Kuttner, K. N., 2001. Monetary policy surprises and interest rates: Evidence from the fed funds futures market. Journal of Monetary Economics 47, 523 – 544.

Lagos, R., Rocheteau, G., Wright, R., 2017. Liquidity: A New Monetarist Perspective. Journal of Economic Literature 55, 371–440.

Li, S., Wang, X., Ye, M., 2019. Who provides liquidity and when? Working paper.

Musto, D., Nini, G., Schwarz, K., 2018. Notes on bonds: Illiquidity feedback during the financial crisis. The Review of Financial Studies 31, 2983–3018.

Nagel, S., 2016. The liquidity premium of near-money assets. Quarterly Journal of Economics 131, 1927–1971.

Nakamura, E., Steinsson, J., 2018. High Frequency Identification of Monetary Non-Neutrality: The Information Effect. Quarterly Journal of Economics 133, 1283–1330.

Naranjo, L., 2009. Implied Interest Rates in a Market with Frictions. Tech. rep., working Paper.

Rigobon, R., Sack, B., 2004. The Impact Of Monetary Policy On Asset Prices. Journal of Monetary Economics 51, 1553–1575.

Ronn, A., Ronn, E., 1989. The box spread arbitrage conditions: Theory, tests, and investment strategies. Review of Financial Studies 2, 91–108.

Sharpe, W., 1964. Capital asset prices: A theory of market equilibrium under conditions of risk. Journal of Finance 19, 425–444.

Stein, J., 2012. Monetary Policy as Financial-Stability Regulation. Quarterly Journal of Economics 127, 57–95.

Svensson, L. E., 1994. Estimating and interpreting forward interest rates: Sweden 1992 - 1994. Working Paper 4871, National Bureau of Economic Research.

van Binsbergen, J., Brandt, M., Koijen, R., 2012. On the timing and pricing of dividends. American Economic Review 102, 1596–1618.

Appendix: Interpreting Rates Through a Simple Model

This section presents a theoretical model that provides conditions under which our approach provides a valid estimate of the convenience yield of a safe asset. Motivated by evidence that hedge funds are the primary market makers for the options we study, we analyze the investment decisions of a hedge fund facing a variety of potential financial frictions that distort its pricing of assets. While all asset prices are distorted by financial frictions distinct from the convenience yield on safe assets, we show that the difference in the prices of two riskless assets is driven purely by the convenience yield these safe assets provide. The hedge fund at time $t = 0$ chooses a portfolio of assets that pay off at $t = 1$ in order to maximize the present value of cash flows it pays to its clients. The hedge fund has an existing portfolio holding quantities $q_{i,old}$ of each asset i, where asset i pays δ_i at time 1. At time 1, these assets in place (*old* assets) pay:

$$\delta_{old} = \sum_i \delta_i q_{i,old}$$

The hedge fund can also buy an additional quantitiy $q_{i,new}$ of each asset i, paying:

$$\delta_{new} = \sum_i \delta_i q_{i,new},$$

so that its total portfolio pays:

$$\delta = \delta_{old} + \delta_{new}.$$

Further, the hedge fund has a promised payoff γ_{old} on its outstanding debt to its prime broker. The fund can also borrow more by promising an additional payoff γ_{new}. Its total promised debt payoff is therefore $\gamma = \gamma_{old} + \gamma_{new}$. If the hedge fund defaults on its debt, its prime broker seizes all of its assets which pay δ. The payoff of the hedge fund's debt is therefore

$$\delta_{debt} = min(\delta, \gamma)$$

leaving a residual payoff to hedge fund of

$$\delta - min(\delta, \gamma) = max(\delta - \gamma, 0).$$

Finally, it can issue equity e, for which it pays a convex transactions cost $C(e)$. The hedge fund remains in business as long as $\delta > \gamma$ and defaults otherwise.

The prime broker uses a pricing kernel m_b to price the cash flows $min(\delta, \gamma)$ paid by the hedge fund's debt.[23] The value of the hedge fund's debt to the prime broker is therefore $E[m_b min(\delta, \gamma)]$. Importantly, because the hedge fund has already issued debt promising γ_{old} a debt overhang problem emerges.[24] Given the payoff δ of the hedge fund's portfolio, this old debt entitles the prime broker to a payoff of $min(\delta, \gamma_{old})$ regardless of how much it lends today. As a result, the hedge fund cannot borrow against any cash flow added to its portfolio in states of the world where it would default on its old debt. After all, such a cash flow is already promised in an existing loan. Instead of lending $E(m_b[min(\delta, \gamma)])$ to the hedge fund, debt overhang induces the prime broker to only lend $E(m_b[min(\delta, \gamma)]) - E(m_b[min(\delta, \gamma_{old})])$.

The model also features certain types of "special" collateral posted by the hedge fund that the prime broker values. This is motivated by evidence that government debt and other money-like assets can be pledged as collateral to borrow at particularly low so-called "special repo rates." In the model s_i is our measure of the convenience yield of asset i. For each such asset i the prime broker is willing to lend an additional s_i (compared to a non-special asset with identical cash flows) if the hedge fund posts such an asset as collateral. As a result, the total amount lent by the prime broker is $\sum_i q_i s_i + E(m_b[min(\sum_i \delta_i q_i, \gamma_{new} + \gamma_{old})]) - E(m_b[min(\sum_i \delta_i q_i, \gamma_{old})])$. This amount plus the equity e raised by the hedge fund equals the value of assets it is able to purchase.

The hedge fund maximizes the present value of its payoffs to investors, with the pricing kernel m_h:

$$\max_{\{q_i\}, e, \gamma_{new}} E(m_h max(\sum_i \delta_i q_i - \gamma_{new} - \gamma_{old}, 0)) - e - C(e)$$

subject to the budget constraint:

$$\sum_i (q_i - q_{i,old}) p_i \quad = \quad e + \sum_i q_i s_i + E(m_b[min(\sum_i \delta_i q_i, \gamma_{new} + \gamma_{old})]) - E(m_b[min(\sum_i \delta_i q_i, \gamma_{old})])$$

Plugging in the budget constraint yields:

$$\max_{\{q_i\}, e, \gamma_{new}} E(m_h[max(\sum_i \delta_i q_i - [\gamma_{old} + \gamma_{new}], 0)]) -$$
$$G(\sum_i (q_i - q_{i,old}) p_i - \sum_i q_i s_i - [E(m_b[min(\sum_i \delta_i q_i, \gamma_{new} + \gamma_{old})]) - E(m_b[min(\sum_i \delta_i q_i, \gamma_{old})])])$$

[23]This assumption reflects the fact that prime brokers tend to use sophisticated portfolio margin models. See https://www.theocc.com/risk-management/cpm/.

[24]A hedge fund that borrows only at short maturities, which anecdotally is common, need not have a debt overhang problem.

where G is the function $G(x) = x - C(x)$. Differentiating with respect to q_i yields the first order condition:

$$\begin{aligned}
p_i &= s_i + E\left[\frac{m_h}{1+C'(e)}\delta_i I_{[\delta>\gamma]}\right] + E\left[m_b \delta_i I_{[\gamma_{old}<\delta<\gamma]}\right] \\
&= s_i + E\left[\left(\frac{m_h}{1+C'(e)} I_{[\delta>\gamma]} + m_b I_{[\gamma_{old}<\delta<\gamma]}\right)\delta_i\right] \\
&= s_i + E\left[M\delta_i\right],
\end{aligned}$$

where the hedge fund prices assets with the stochastic discount factor

$$M \equiv \left(\frac{m_h}{1+C'(e)} I_{[\delta>\gamma]} + m_b I_{[\gamma_{old}<\delta<\gamma]}\right).$$

The indicator function $I_{[\delta>\gamma]}$ equals 1 when $\delta > \gamma$ (i.e., the no default region) and 0 otherwise. The second indicator function $I_{[\gamma_{old}<\delta<\gamma]}$ equals 1 when the hedge fund defaults on the new debt, but would not have defaulted on the old debt. Because the hedge fund faces a debt overhang problem, any marginal cash flow paid when the old debt would have been defaulted on does not impact the hedge fund's willingness to pay for an asset. Finally, the equation shows that there is a unique stochastic discount factor M that prices all assets held by the hedge fund, even though frictions can distort it away from the hedge fund's preferences (m_h).

The expression above shows that the price of an asset is equal to its present value (discounted using the stochastic discount factor M) plus its convenience yield s_i. The stochastic discount factor M reflects both the hedge fund's preferences m_h as well as several frictions. First, the hedge fund's pricing kernel is distorted by its marginal cost of equity financing. If the hedge fund is financially constrained, this will distort its pricing kernel away from that m_h that prices its payouts to equity holders. Second, the hedge fund and its prime broker may have different pricing kernels, reflecting the potential for segmentation between markets dominated by banks and those dominated by hedge funds. The prime broker may be more risk averse than the hedge fund (perhaps due to regulatory constraints or due to its desire to borrow from less sophisticated investors), or may have investment opportunities that differ from the hedge fund. Finally, there may be a debt overhang problem. The cash flow $\delta_i I_{[\delta<\gamma]}$ is going to the prime broker, but the portion $\delta_i I_{[\delta<\gamma_{old}]}$ would have been paid to its existing debt claim. This debt overhang problem also distorts the hedge fund's willingness to pay for financial assets, reducing the hedge fund's willingness to pay by: $E\left[m_b \delta_i I_{[\delta<\gamma_{old}]}\right]$.

Despite all these potential frictions, for any two assets with identical cash flows, the difference in the prices of the two assets is due entirely to the assets' convenience yields.

Suppose that there exists a set of risky assets paying δ_j that can be used to replicate a riskless payoff, so that $\sum_j^J w_j \delta_j = 1$, where the w_j's are the portfolio weights required to create the riskless payoff. It follows that:

$$\sum_{j=1}^{J} w_j \left(p_j - s_j\right) = E\left[\frac{m_h}{1+C'(e)} I_{[\delta>\gamma]}\right] + E\left[m_b I_{[\gamma_{old}<\delta<\gamma]}\right]$$

If there is another asset with $\delta_i = 1$, we have that

$$p_i - s_i \quad = \quad E\left[\frac{m_h}{1+C'(e)} I_{[\delta>\gamma]}\right] + E\left[m_b I_{[\gamma_{old}<\delta<\gamma]}\right]$$

and thus

$$p_i - s_i = \sum_{j=1}^{J} w_j (p_j - s_j).$$

If each of the risky assets $(j = 1, .., J)$ is itself not special (e.g. options), the specialness of asset i is given by

$$s_i = p_i - \sum_{j=1}^{J} w_j p_j.$$

The important insight from this analysis is that even in the presence of a range of frictions distinct from the specialness of safe assets, such as debt overhang, frictions in equity issuance and differing pricing kernels for hedge funds and prime brokers, the specialness of government debt is still identified by our spread between the Box rate and Treasuries. However, this specialness that we have identified is specific to a hedge fund that is active in both option and bond markets and whose transaction costs for trading in the option market are sufficiently low so that midpoint prices are representative.[25] While the fund must be active in both markets to validate our convenience measure, we do not need to assume that the hedge fund uses the option market for large quantities of risk-free borrowing or lending. In addition, our option-implied rate can be interpreted as the correct risk free rate for a pricing kernel for risky asset markets where our hedge is active, even though this pricing kernel is distorted

[25] Section 4 provides empirical evidence that the bid-ask spread contracts and expands symmetrically around the midpoint price and do not cause distortions in the level of our option-implied rate, consistent with the midpoint price of an option representing the market maker's true valuation of the underlying asset. The midpoint of bid and ask also reflects the fair value of an asset to a market maker in some theoretical models of market making (e.g. Brolley and Malinova (2018)).

by financial frictions.

Online Appendix

A. Trade Data

In this appendix, we check whether options prices from actual trades are consistent with the rates we estimate from option quotes. We repeat our calculations using transaction prices for all options with maturities over 59 days. In each day-minute, and for each option maturity and strike, we check whether or not a transaction happened. If so, we substitute the quote value in our dataset with the volume weighted average mean price for that option in that minute, and recompute the difference between put and call price with this new observation. We then repeat our risk free rate calculations as before. For each maturity-day-time, we calculate an "error" as the difference between our benchmark estimate of the risk-free rate, and the new estimate that includes transaction prices. We then take the median of this difference (or error) at the daily level, and report the summary statistics of its distribution in Table 11.

The Table has four lines because to check the robustness of our estimates, we compute risk-free rates in the new dataset using three different weighting schemes. In all cases, results are nearly identical. First, we run a simple OLS regression. As the first line in the Table shows,the median error is zero, and the 1st and 99th percentile are in the range of 0.2 bps. We further run three weighted linear regressions.

In the first of them, for each given option maturity, and in each given day-minute, we assign a weight of 25% split equally among all observations including trades, and a weight of 75% split among the only-quote observations. In the second, we compute risk free rates by equally weighting trade data and quote data. That is, we assign a weight of 50% split among the observations including a trade and quote-only observations. Finally, we assign a weight of 75% split equally among all observations including trades, and a weight of 25% split among the only-quote observations. The average and median errors are larger than the ones from a simple linear regression, but the overall inference is the same. The median error is less than 1 b.p. in all cases, while the 1st and 99th percentile are in the range of 1-2 bps for the weighted linear regression with weights 25/75 (majority of weight on quotes), 2-3 bps for the weighted linear regression with weights 50/50, and 2-4 bps for the weighted linear regression with weights 75/25 (majority of weights on trades). These results suggest that our estimates are robust to the use of trade instead of quote data when it comes to the average level of our convenience yield, though the standard errors around our estimated rates do increase significantly when including such trade data.

Table 11
Rates Using Trade and Quote Data

Variable	Mean	St. Dev.	Min	1%	25%	Median	75%	99%	Max
OLS	0.00000	0.00001	-0.00004	-0.00002	-0.00000	0.00000	0.00000	0.00002	0.00005
WLS – 25%	0.00001	0.00004	-0.00024	-0.00007	-0.00000	0.00000	0.00003	0.00016	0.00034
WLS – 50%	0.00003	0.00008	-0.00086	-0.00014	-0.00000	0.00002	0.00005	0.00032	0.00069
WLS – 75%	0.00004	0.00011	-0.00123	-0.00019	-0.00001	0.00003	0.00007	0.00044	0.00091

Notes: The Table shows summary statistics for the difference between our benchmark estimate of the risk-free rate, and a new estimate including transaction prices (trades). The first column lists the three regression weighting schemes as explained in the text: equally weighting all observations (OLS), assigning a weight of 25% to trade data (WLS – 25%), assigning a weight of 50% to trade data (WLS – 50%), assigning a weight of 75% to trade data (WLS – 75%). The remaining columns report the mean, standard deviation and various percentiles of the distribution of such a difference computed under each of the three weighting schemes.

B. Analysis of Error Terms

This appendix provides evidence that all option prices are consistent with an effectively unique option-implied discount rate at any point in time. We do so by analyzing the cross-sectional distribution of residuals ε_i from the regression in Equation 2 whose slope β provides our option-implied discount factor. We have three main results. First, the residuals in Equation 2 are quite small in every year of the data. We scale the residual ε_i by dividing by the corresponding strike K_i, so that the residuals can be interpreted as an percentage interest rate deviation. We report summary statistics of the distribution of residuals in our minute-by-minute regressions for each year in Table 12, and even the largest deviations are single digit basis points or up to an extreme of 15 basis points. Before the financial crisis the 1st and 99th percentiles are roughly 3-5 basis points in absolute value, growing to the slightly larger 9-11 basis points in 2008-2015, and then falling to 7-8 basis points after.

Second, we find that our residual errors revert back to 0 rapidly, though less so after the financial crisis. We demonstrate this in two ways. First, we report minute-level autocorrelations of the persistence of our residuals in Table 12. All residuals coefficients are considerably less than 1 at a minute-level frequency, suggesting that deviations decay in a manner of hours. Second, we individually analyze and plot the extreme outliers that may potentially impact single rate estimates without showing up in an aggregate analysis. We compute the median residual across our minute-by-minute regressions in each day, and find the largest (below 1st and above 99th percentile) daily residuals and then plot this residual in an 11 day event window. There are sufficiently few such extreme observations that we can plot every such series in Figure B.1 below, and the few most extreme observations are visually salient. We find before the financial crisis that each extreme residual visually mean reverts in the day before and after it appears. During and after the financial crisis, there is some additional persistence in our most extreme residuals. This suggests that supply and demand imbalances specific to individual strike prices as in Garleanu et al. (2009) are more important for options prices after the financial crisis, even though our convenience yield measure returns back to its pre-crisis level.

Third, we explore how our Theil-Sen interest rates relate to the most extreme residuals that we see. In Figure B.2, we plot the correlation of our rate estimate with the magnitude of the residuals plotted in Figure B.1 across the 11-day event windows considered in the plot. If these outliers were an important driver of the estimated rates that we use in our analysis, we should expect to find a correlation between these outliers and our rates. We find relatively small correlations, both in levels and in absolute values whose sign varies year by year, suggesting no clear pattern relating individual extreme observations to potential contamination of our rates away from a risk-free rate estimate, consistent with prices across multiple asset classes. The only exception is that the correlation in absolute values is somewhat elevated in both the 2009 and 2012 crisis years. Note however that the correlation in absolute values would also be elevated if the overall variance of rates is positively correlated with the occurrence of outliers. In addition, the results at the bottom of table 12 show that our extreme residuals are not predicted by the strike price of the option considered, suggesting that the prices of far in- and out-of-the-money options are consistent with overall interest rates we estimate.

Table 12

Regression Error Analysis by Calendar Year

ε_i/K_i	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017
Minute-level														
1%	-0.00040	-0.00037	-0.00034	-0.00031	-0.00091	-0.00100	-0.00121	-0.00113	-0.00144	-0.00105	-0.00092	-0.00102	-0.00077	-0.00071
5%	-0.00017	-0.00018	-0.00014	-0.00013	-0.00029	-0.00041	-0.00051	-0.00051	-0.00058	-0.00041	-0.00038	-0.00046	-0.00032	-0.00028
10%	-0.00011	-0.00012	-0.00009	-0.00008	-0.00016	-0.00025	-0.00033	-0.00033	-0.00037	-0.00025	-0.00024	-0.00029	-0.00019	-0.00017
25%	-0.00005	-0.00005	-0.00004	-0.00004	-0.00006	-0.00010	-0.00013	-0.00013	-0.00015	-0.00011	-0.00011	-0.00012	-0.00007	-0.00007
50%	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000	0.00000
75%	0.00004	0.00005	0.00004	0.00003	0.00005	0.00009	0.00012	0.00012	0.00014	0.00010	0.00010	0.00011	0.00007	0.00006
90%	0.00011	0.00011	0.00009	0.00008	0.00016	0.00025	0.00033	0.00032	0.00036	0.00025	0.00025	0.00029	0.00019	0.00017
95%	0.00019	0.00017	0.00015	0.00015	0.00029	0.00041	0.00055	0.00052	0.00058	0.00043	0.00040	0.00047	0.00032	0.00029
99%	0.00050	0.00037	0.00040	0.00045	0.00094	0.00111	0.00139	0.00125	0.00157	0.00108	0.00095	0.00102	0.00079	0.00074
AC(1)	0.9064	0.9775	0.9959	0.9644	0.8702	0.9177	0.7306	0.8252	0.9145	0.9297	0.8533	0.8092	0.7763	0.9384
S.E. AC(1)	0.0003	0.0001	0.0000	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001	0.0000	0.0001	0.0001	0.0001	0.0000
Coefficients on strikes														
Strike	-6.60E-09	-1.00E-08	2.08E-10	2.74E-09	-1.09E-08	-1.26E-08	-9.13E-09	-9.17E-09	-3.45E-09	-4.87E-09	-2.6E-09	-9.71E-09	-1.84E-09	-1.30E-10
Strike^2	1.18E-11	1.83E-11	-3.85E-13	-3.93E-12	1.54E-11	2.17E-11	1.51E-11	1.31E-11	4.66E-12	6.20E-12	2.97E-12	1.08E-11	2.03E-12	1.23E-13
Strike^3	-4.92E-15	-7.83E-15	1.68E-16	1.39E-15	-5.16E-15	-8.56E-15	-5.82E-15	-4.32E-15	-1.46E-15	-1.86E-15	-8.20E-16	-2.84E-15	-5.23E-16	-2.80E-17
Rate change due to a change in strike from 1500 to 2500														
Strike	-0.00001	-0.00001	0.00000	0.00000	-0.00001	-0.00001	-0.00001	-0.00001	0.00000	0.00000	0.00000	-0.00001	0.00000	0.00000
Strike^2	0.00005	0.00007	0.00000	-0.00002	0.00006	0.00009	0.00006	0.00005	0.00002	0.00002	0.00001	0.00004	0.00001	0.00000
Strike^3	-0.00006	-0.00010	0.00000	0.00002	-0.00006	-0.00010	-0.00007	-0.00005	-0.00002	-0.00002	-0.00001	-0.00003	-0.00001	0.00000

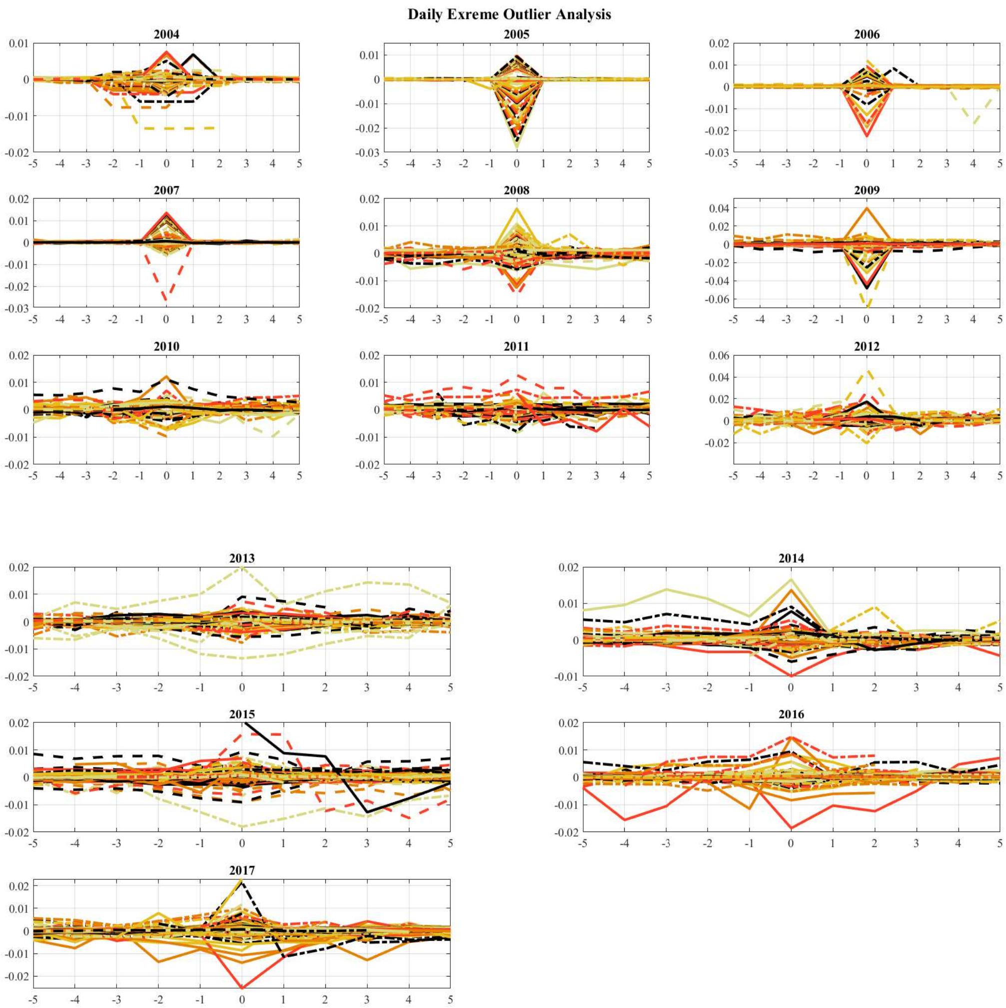

FIGURE B.1

Eleven-day Event Studies Around Extreme Outliers. We plot for each calendar year 11-day event windows around the most extreme outliers (below 1st percentile and above 99th percentile) of ε_i/K. The plots follows the same option contract (i.e., the same strike and the same maturity) 5 days before and 5 days after the event.

C. Rates from Commodity Markets

To construct a risk-free asset in commodity markets we use the cost-of-carry relationship between the futures price ($F_{t,T}$ and the spot price S_t) which states that:

$$F_{t,T} = S_t \exp\left((r_{t,T} + c_{t,T})T\right). \tag{C.1}$$

where $r_{t,T}$ is the implied continuously compounded risk-free interest rate and $c_{t,T}$ is the net storage cost of the commodity. To derive estimates of the risk-free interest rate, we focus on

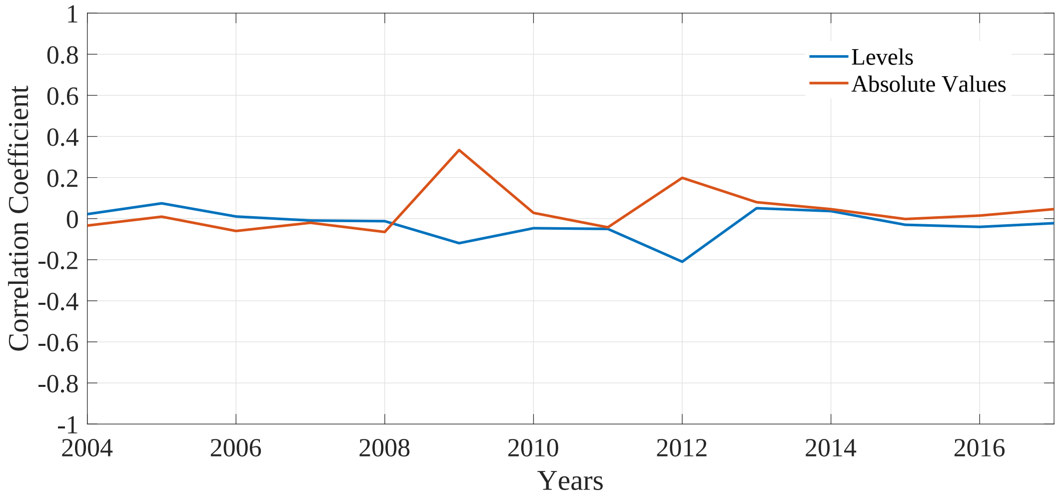

FIGURE B.2
Correlation Between Extreme Outliers And Estimated Rates. We plot the correlation of our rate estimate with the magnitude of the residuals plotted in figure B.1 across the 11-day event windows considered in the plot.

futures contracts on underlying assets that are very cheap to store relative to their underlying value, implying that the term $c_{t,T}$ is essentially zero. As such we focus on precious metals: gold and silver. The risk-free rate is then computed as:

$$r_{t,T} = \frac{1}{T} \ln \left(\frac{F_t}{S_t} \right). \tag{C.2}$$

Our data set contains all futures trades made between May 2007 and January 2018 on the Chicago Mercantile Exchange (CME) regarding two precious metals: gold, silver. Unlike the CBOE data, which also contains quotes, the database we purchased only contains trades.

In Table 13 we summarize the key statistics for gold and silver implied interest rates. We compare these rates to the government bond yields as implied by the NSS parameters (Gürkaynak et al. (2007)). The table shows that the estimated convenience yield for government bonds relative to metal-implied interest rates is the same as for our previous estimates and equal to about 40 basis points for gold with no apparent relation to maturity. For silver the order of magnitude is the same, but there now seems to be a maturity dependence of the estimate, with the convenience yield decreasing with maturity.[26]

[26] In unreported results, for platinum, the data is not sufficiently rich to obtain (interpolated) term structure data. However, the average convenience yield across all available maturities is 50 basis points. The volatility of the daily estimates is large, partly due to the fact that we only have trade data and not quote data.

Table 13
Risk-free rates and convenience yields implied by precious metal prices

Zero Coupon Yield Curve				
	Gold		Silver	
	Mean	St. Dev	Mean	St. Dev
Metal implied 6m	0.0118	0.0123	0.0133	0.0174
Metal implied 12m	0.0120	0.0117	0.0116	0.0126
Metal implied 18m	0.0127	0.0112	0.0116	0.0124
Metal Implied - Gov. Bond 6m	0.0043	0.0035	0.0054	0.0117
Metal Implied - Gov. Bond 12m	0.0040	0.0027	0.0036	0.0049
Metal Implied - Gov. Bond 18m	0.0037	0.0027	0.0024	0.0050

D. A Multivariate Analysis of Spreads

This section jointly studies the dynamics of our convenience yield measure and several arbitrage spreads that have been documented in the literature. We replicate these arbitrage spreads in Online Appendix B. The measures we consider are (1) the so-called 6-month spread (the difference between the yield on zero-coupon Treasury bills and the yield on notes and bonds that mature within 6 months), (2) the spread between principal and interest STRIPS, also called the STRIP spread, (3) the spread between on-the-run and off-the-run long-term bonds of equal maturity, and (4) the spread between notes and seasoned bonds with the same remaining maturity.

We find that the SPX option-implied convenience yield measure is exposed to considerably smaller idiosyncratic shocks than all other spreads in the analysis, and that the level of the SPX convenience yield today contains the vast majority of predictive information for its level in the future. This suggests that while some of the other spreads in our analysis could be thought of as market-specific measures of limits to arbitrage, our convenience yield measure can be thought of as measuring the overall scarcity of safe assets.[27]

To establish these results, we estimate a first order Vector Auto Regression (VAR) that in addition to our SPX convenience yield measure includes the bond arbitrage spreads above, as well as the DJX- and gold-implied convenience yields. The results are reported in Table 14.

The table shows that both option-implied interest rate series (SPX and DJX) are subject to considerably less idiosyncratic risk than others, while the metal convenience yields are subject to more.[28] As a result, if we are interested in using data on risky asset prices to

[27] We thank Eben Lazarus for providing this interpretation of our findings.

[28] This is also due to the fact that the convenience yields on precious metals are estimated with trade data only (as opposed to quote data) leading to noisier estimates.

Table 14
VAR(1) Analysis of Arbitrage Spreads

	djx	spx	lessthan6	metal	notesbonds	ontherun
djx	0.6798	0.0634	0.0032	-0.0585	0.0001	-0.0002
spx	0.2662	0.9103	-0.0400	0.4142	-0.0002	0.0001
lessthan6	-0.0066	-0.0323	0.5319	1.4711	-0.0018	-0.0015
metal	0.0005	-0.0002	0.0031	0.0223	0.0000	0.0000
notesbonds	1.6658	0.4304	0.2773	4.7072	0.7044	-0.0094
ontherun	0.0671	0.1230	-0.4992	-8.0279	-0.0037	0.4580
constant	2.7326	0.6319	-1.8320	41.4707	0.0078	-0.0281
R^2	0.8332	0.9403	0.3494	0.0126	0.5004	0.2147

infer a risk-free rate consistent with investors' cost of capital, the option-implied rates seem to be our best candidate. The R-squareds of predicting the in-sample SPX-implied rate is extremely high (94%), and substantially higher than that of the DJX (83%). Government bond arbitrages have intermediate R-squareds, while the metal series have by far the lowest. This implies that there does not seem to be a high degree of unpredictable, non-persistent noise in the SPX-implied rates. The SPX-implied rate seems most consistent with the intuitive notion that the cost of capital does not have extreme high frequency fluctuations.

E. Other Arbitrage Measures

In this appendix, we replicate several bond arbitrage measures from the literature and compare them to the dynamics of our estimated convenience yield. We consider four distinct categories of arbitrage spreads using government bond data. Two of them relate to zero coupon bond arbitrages, which can be computed without estimating (and interpolating) a yield curve, and two of them involve bonds with coupon payments that do require an estimated yield curve.

E.1. Zero Coupon Bond Arbitrages

6 month Spread

First, we consider the spread between notes/bonds that mature within the next 6 months and yields on Treasury bills that mature on the exact same date. Treasury bills are more liquid and therefore tend to have lower yields (Amihud and Mendelson (1991)). Because

Treasury securities pay coupons every 6 months, there are no intermediate coupon payments for either security used in constructing this spread. For each day, we compute the median of the continuously compounded yields to construct a daily time series.

STRIP Spread

Second, we consider the spread between two types of STRIPS (Separate Trading of Registered Interest and Principal of Securities) constructed respectively from interest and principal payments on U.S. government debt. These securities pay identical cash flows and are backed by the full faith and credit of the U.S. government, so any difference between the yields on coupon vs principal STRIPS identifies an arbitrage. In general, whichever of the principal or interest STRIP that has a higher supply outstanding tends to have a lower yield. At short maturities, interest STRIPS are in larger supply while at long maturities principal STRIPS are.[29] Because all principal and interest payments happen on a regular 6-month schedule, there are enough overlapping bonds to consider only spreads between interest and principal strips that mature on exactly the same day. We present averages of both the level as well as the value of this spread across all maturity matched pairs of coupon and principal STRIPS below.

E.2. Coupon Bond Arbitrages

Because the two spreads we study in the previous subsection are between pairs of zero coupon securities with the exact same maturity, no assumptions were required regarding the shape of the yield curve to construct them. This is not true for the two arbitrage spreads that we consider next. The reason is that these next two measures relate to government bonds that make coupon payments for which no exact matching security may exist. As a result, we compute these spreads by comparing a bond's true yield to the yield implied by fitting a yield curve to all Treasury bonds. To estimate this yield curve, we estimate a parametric model following Svensson (1994), and Gürkaynak et al. (2007). A Nelson-Siegel-Svensson (NSS) instantaneous forward rate τ periods in the future is assumed to have the functional form

$$f(\tau) = \beta_0 + \left(\beta_1 + \beta_2 \frac{\tau}{\tau_1}\right) \exp\left(-\frac{\tau}{\tau_1}\right) + \beta_3 \frac{\tau}{\tau_2} \exp\left(-\frac{\tau}{\tau_2}\right).$$

Given parameters $(\beta_0, \beta_1, \beta_2, \tau_1, \tau_2)$, this forward rate function uniquely implies a zero coupon yield curve that can be used to price any risk-free bond. To estimate the parameters, we use data from GovPX between 3pm and 4pm of each day and consider the price of all off-the-run

[29]The reason is that all bonds of all maturities pay coupon payments every 6 months and contribute to the coupon-related supply.

notes and bonds. Let y_i be the yield to maturity of bond i, D_i be the duration of bond i, and $y_i(\beta_0, \beta_1, \beta_2, \tau_1, \tau_2)$ be its yield to maturity implied by the NSS yield curve. We estimate the parameters of the yield curve for each day by minimizing

$$\sum_i \frac{1}{D_i} \left(y_i - y_i(\beta_0, \beta_1, \beta_2, \tau_1, \tau_2)\right)^2$$

where the sum i goes over all bond quotes between 3 and 4pm that day.

On the Run Spread

We use the NSS yield curve to compute an implied yield for the most recently issued bond of each maturity, called the on-the-run bond, and take its difference from the true yield on that bond. On-the-run bonds tend to be more liquid than off-the-run bonds and therefore trade at a lower yield as shown below. The spread between on- and off-the-run bonds is related to the timing of the Treasury auction cycle. This is particularly true for the yield on the on-the-run 10-year bond. We plot the spread between this on-the-run 10-year bond yield and yield implied by the NSS yield curve in Figure E.1.

Notes vs Bonds

We also use the NSS yield curve to consider the relative spread between Treasury notes (which by definition have a maturity less than 10 years after issuance) and bonds (which mature more than 10 years after issuance), that have less than 10 years of maturity left, following (Musto, Nini, and Schwarz (2018)). For each note and bond that mature between 3 and 10 years from the day on which the security is traded, we compute the spread between the security's actual yield to maturity and the yield implied by the estimated NSS yield curve. We then take the median of this spread across all notes, and the median of the spread across all bonds on each day and compute a daily difference between these two medians. As the above authors show, this spread is small in normal times but spikes during the financial crisis.

E.3. Data

As mentioned previously, our U.S. Treasury security prices come from the GovPX database, which reports trades and quotes from the inter-dealer market for U.S. Treasuries. We use indicative quotes, which provide the most frequent measure of bond prices on GovPX from 3 to 4pm on each day. In addition, we have data from Tradeweb on the prices of STRIPS, which are zero coupon bonds created by separating the principal and interest payments on Treasury securities. This database provides quotes 2 times a day, and we restrict ourselves to quotes at 3pm. Whenever using quote data, we take the midpoint of the bid and ask as the price measure.

Table 15
Summary Statistics of Government Bond Arbitrages 2004-2018

	Mean (in bp)	St. Dev.
6 month Spread	6.4388	6.9544
STRIP Spread	3.8696	8.8144
On the Run Spread		
All Bonds	0.4945	1.9194
10 year Bond	2.1587	2.4044
Notes vs Bonds	0.3576	0.9573

E.4. Results

First, we present summary statistics on the four above-mentioned government bond arbitrages in Table 15 and we plot them in Figures E.1, E.2 and E.3.

Several patterns appear across all of these arbitrage spreads. First, both their level and volatility generally increase during the financial crisis period of late 2008 and early 2009. Second, most spreads are smallest in the later part of our sample, suggesting that government bond markets are now even more integrated than they were before the crisis. Third, some spreads (such as the 10-year on the run spread) seem to be driven in part by idiosyncratic factors such as the Treasury auction cycle. That is, the regular spikes in Figure E.1 correspond to auction cycle dates.

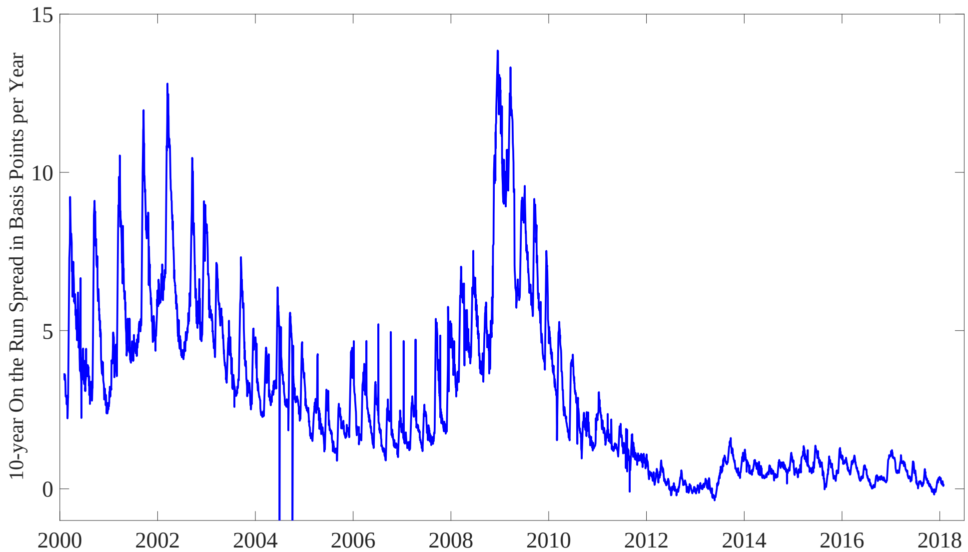

FIGURE E.1
Ten Year On-the-Run Spread in Basis Points per Year.

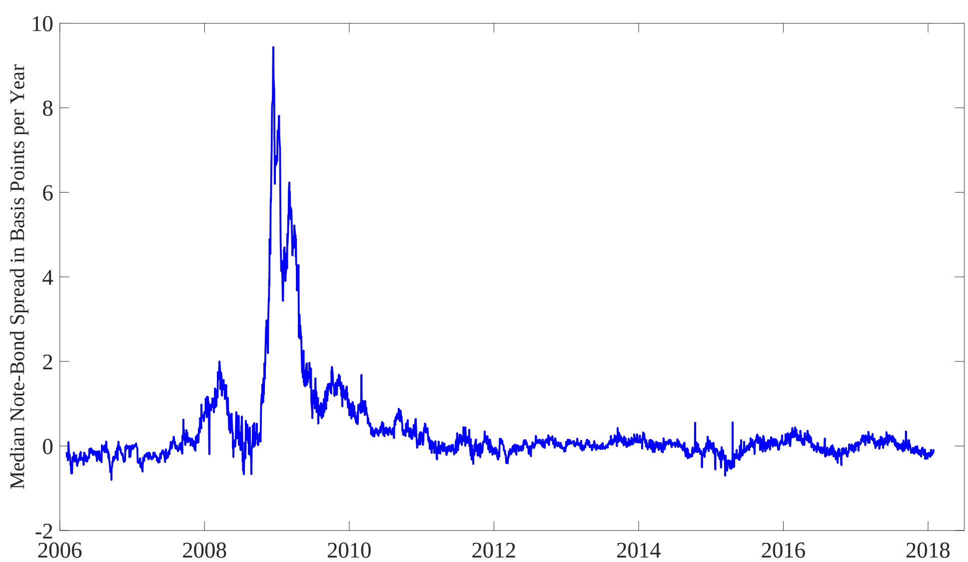

FIGURE E.2
Notes/Bonds Spread in Basis Points per Year.

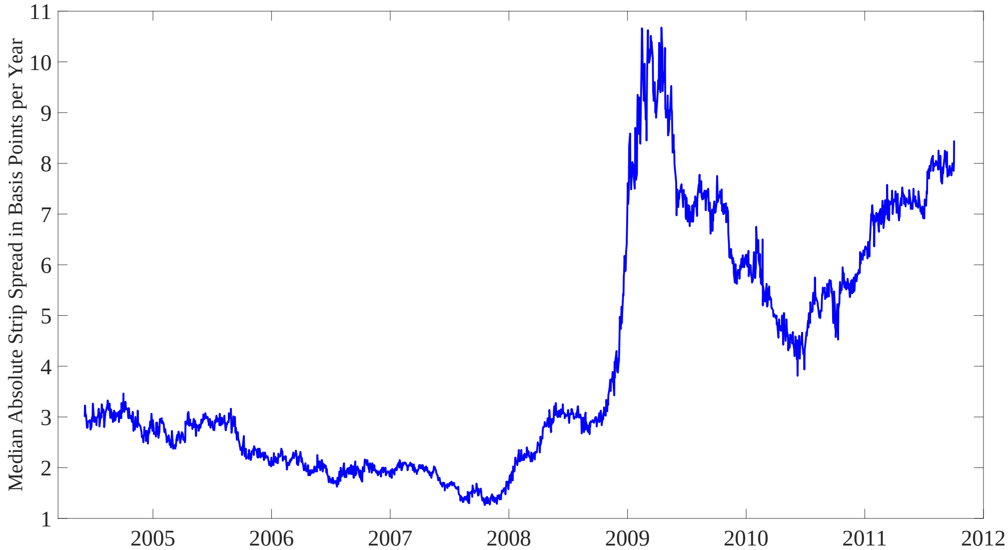

FIGURE E.3
Median Absolute STRIP Spread.

F. Select responses of box rate and government yields to FOMC announcements

Panel A: March 22, 2005

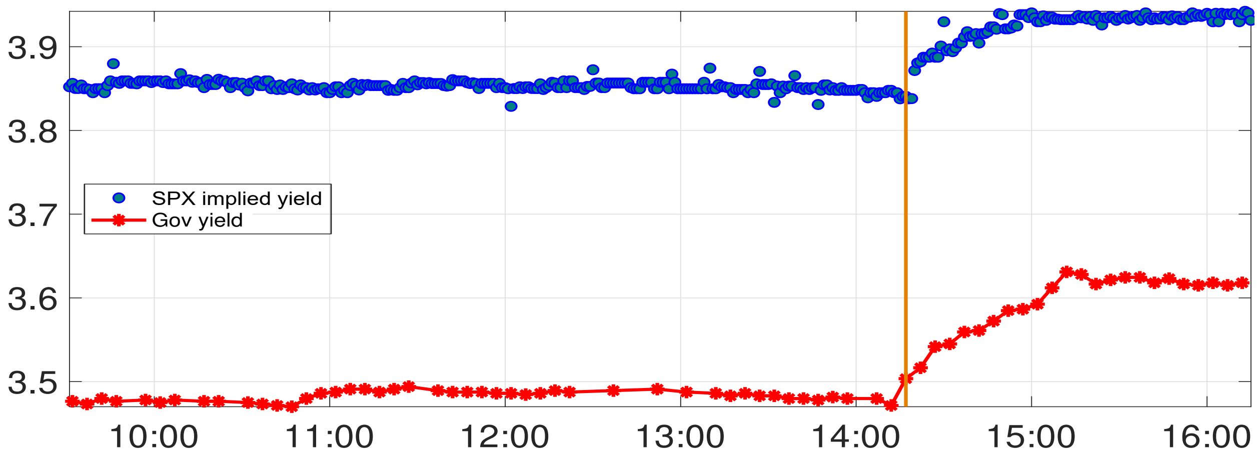

Panel B: June 29, 2006

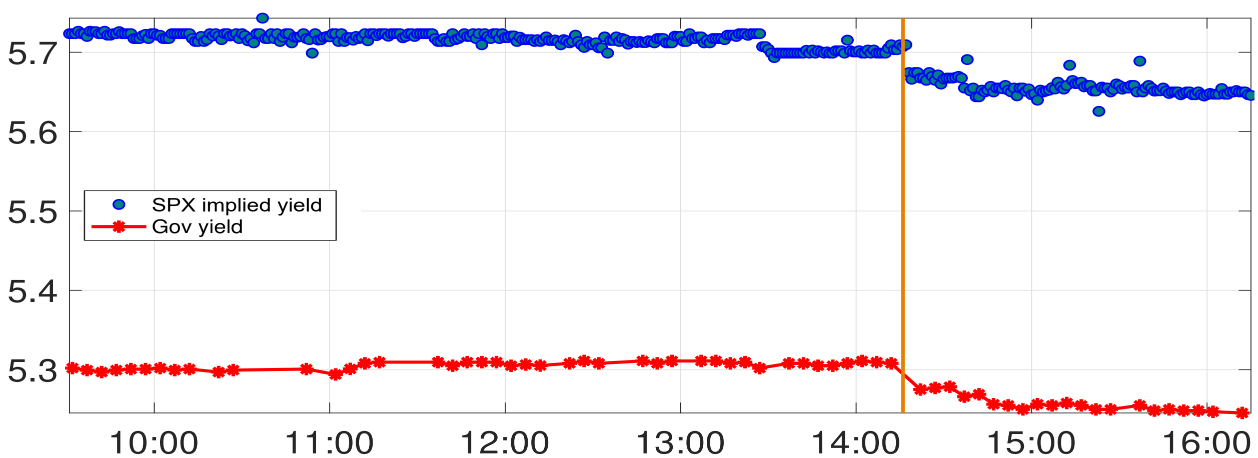

The figure plots the maturity-matched box rates and government yields around the release of the March 22, 2005 and June 29, 2006 FOMC announcements. The maturities are 452 days for Panel A and 541 days for Panel B. The vertical line represents the time of the release. Yields are in %.

Panel A: October 31, 2007

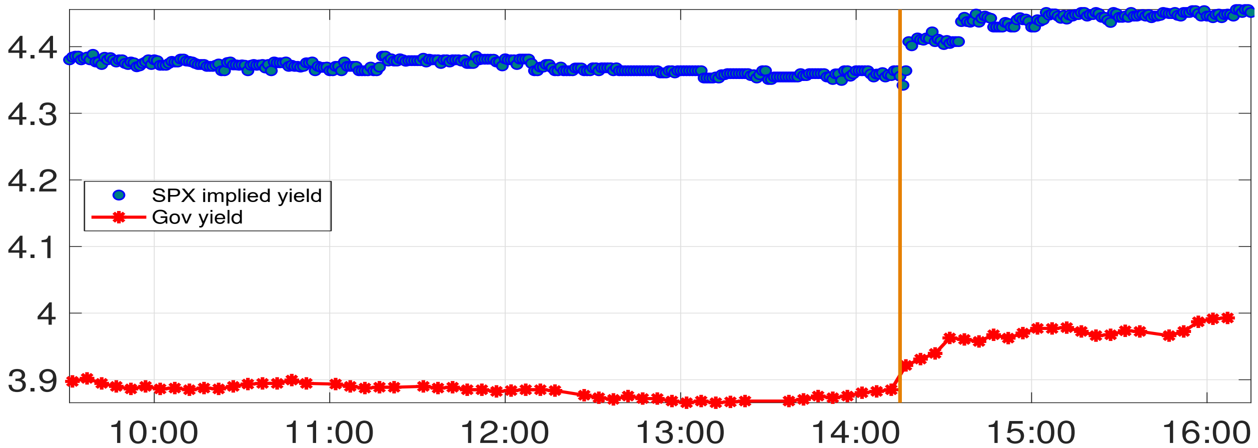

Panel B: September 23, 2009

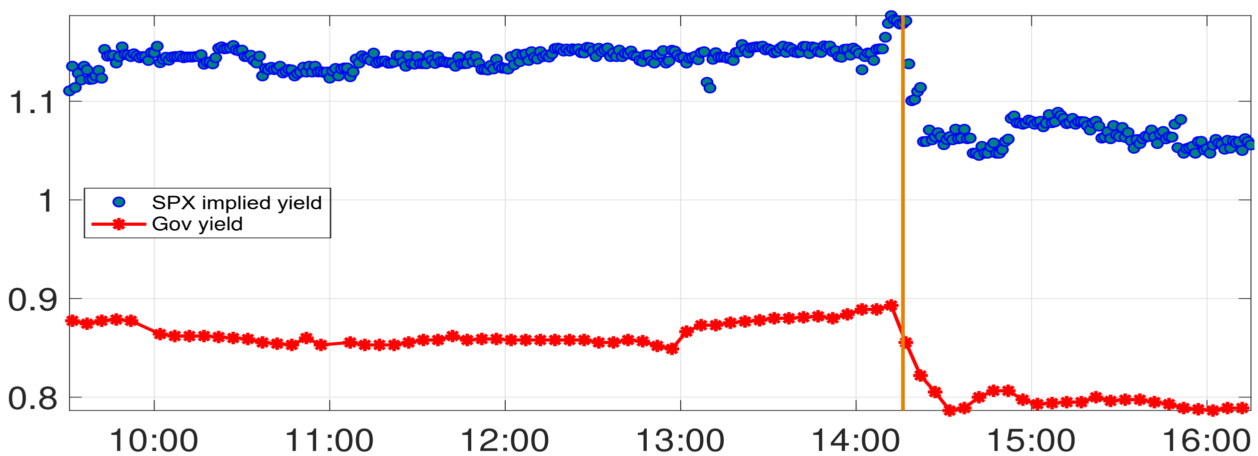

The figure plots the maturity-matched box rates and government yields around the release of the October 31, 2007 and September 23, 2009 FOMC announcements. The maturities are 598 days for Panel A and 633 days for Panel B. The vertical line represents the time of the release. Yields are in %