

Risk Free Interest Rates

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

November 17, 2018

Abstract

We estimate risk free rates unaffected by the convenience yield on safe assets by inferring them from risky options and futures prices. Our data provides time-varying estimates of the term structure of convenience yields from maturities of 1 month to 2.5 years at a minutely frequency. The convenience yield on government bonds equals about 40 basis points on average, is larger below 3 months maturity, and grows substantially during periods of financial distress. With our unique intraday estimates of the term structure of convenience yields, we estimate the high frequency response of convenience yields to monetary policy and quantitative easing. Convenience yields respond most strongly to central bank policy in the depths of the financial crisis, and both conventional and unconventional monetary stimulus reduce convenience yields. Additionally, a factor constructed from our measure of the term structure of convenience yields predicts excess bond returns even when controlling for commonly used factors in the literature. Finally, we study the dynamics of a large panel of other arbitrage spreads and find that our implied convenience yields predict other spreads and face the smallest idiosyncratic shocks, suggesting that they are highly informative measures of frictions in financial markets.

*We thank Jonathan Berk, Wenxin Du, Vincent Glode, Ben Hébert, Christian Opp, Krista Schwarz, and Jonathan Wright 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 can be reached at julesv@wharton.upenn.edu, diamondw@wharton.upenn.edu and grottm@wharton.upenn.edu.

1. Introduction

Arguably the most important variable in financial economics is the interest rate on a risk free investment. In frictionless asset pricing models, it is determined by the investor's time preference. The interest rate on a risk free investment tells us how much an investor is willing to pay today to obtain a payment in the future with absolute certainty. To measure investors' willingness to take risk, the returns on risky assets are compared to risk free interest rates, with the difference in returns ascribed to a risk premium. As a result, any attempt to measure the risk and time preferences of investors from asset prices requires a precise estimate of the return investors require on a risk free asset.

However, a recent literature has provided empirical evidence that the interest rates on safe assets such as U.S. Government debt are driven in part by forces other than the time preference of investors (Krishnamurthy and Vissing-Jorgensen (2012), Nagel (2016)). Safe assets earn a so-called "convenience yield" that pushes their yield below that implied by investors' time preference. This reflects the ease with which safe assets can be traded by uninformed agents, posted as collateral, or perform other roles similar to that of money.

This raises the question we address in this paper: what is the convenience yield free interest rate on a risk free investment? Answering this question has two main applications. First, the spread between such a rate and the observed yield on a safe asset precisely estimates the safe asset's convenience yield. Second, for researchers using frictionless models in which risk free rates are determined only by investors' time preference, it identifies the correct risk free to compare to such a model's predictions.

In this paper, we estimate risk free rates that are unaffected by the convenience yield on safe assets by inferring them from entirely from the prices of risky assets. Our benchmark rate is inferred from the put-call parity relationship in equity options prices (for which we obtain estimates minute by minute), and we also consider rates inferred from a storage arbitrage in the market for precious metals futures. 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 ¹. As a result, a risk free rate inferred only from risky

¹e.g. Frazzini and Pedersen (2014), Stein (2012), Caballero and Farhi (2018), and Diamond (2018).

asset prices should be a pure measure of time preference and not be affected by a convenience yield. Our risk free rate estimates do not require us to use a specific model of risk and return and rely simply on the risky assets we consider being priced in an internally consistent manner free of arbitrage. In any model where risky assets do not earn a convenience yield, we obtain a natural and robust estimate of the risk free rate of return reflecting investors' time preference and unaffected by a convenience yield. As a result, the spread between our estimated rate and that observed on government bonds is therefore a theoretically ideal measure of the convenience yield on safe assets. We find that the convenience yields on government bonds equals about 40 basis points on average, with a relatively flat average term structure across maturities beyond 3 months, and a somewhat higher average below 3 months of maturity. We further find that the convenience yield is strongly time varying and grows substantially during periods of financial distress.

Our minute level estimates of the term structure of convenience yields are ideal for an event study of the effects of monetary policy and quantitative easing. We find that monetary policy and QE have a nontrivial effect on convenience yields, particularly during the depths of the financial crisis. 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 unclear whether the effects of quantitative easing spill over beyond the prices of debt securities that are actually purchased. Under the “narrow” view of its transmission mechanism, quantitative easing should not spill over broadly into the discount rates at which the private sector can borrow despite the lowering of long-term treasury 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 (Diamond (2018) and Caballero and Farhi (2018)), which emphasizes that swapping reserves for long duration assets increases the overall supply of safe assets and should therefore reduce convenience yields across markets. Because our interest rates are inferred from assets distinct from the fixed income market, we can use our data to test whether there is a “narrow” or “broad” transmission mechanism and find support for the latter. In fact, 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.

Because government bond yield movements are affected by the dynamics of convenience yields, a natural question that arises is whether the documented bond return predictability in the literature is related to the dynamics of the convenience yield or to movements in convenience-yield-free rates. We find that a forecasting factor constructed solely from the cross-section of convenience yields in the spirit of Cochrane and Piazzesi (2005) has substantial 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 even outperform the predictive power of conventional predictors in the literature in our sample. The results therefore suggest that a full explanation of the predictability in bond excess returns requires a model that features both a time varying premium related to safe assets (convenience yield) as well as another sources of excess return predictability (i.e. time-varying risk aversion or time-varying volatility).

Our paper contributes to several related literatures. First, it contributes to the empirical literature mentioned above on safe assets by providing a measure of the convenience yield that is tightly connected to theory and provides answers consistent with this theoretical rigor. Some existing proxies for the convenience yield are spreads between yields on two different safe assets, which may be an underestimate if both assets have some convenience yield. Others proxies are spreads between a safe asset and a (slightly) risky 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, which is consistent with the hypothesis that we have correctly estimated the true convenience yield on safe assets. We also find that our spread is almost identical to the LIBOR-treasury spread before the crisis but smaller after, consistent with the view that credit risk in LIBOR was considered negligible before the crisis but significant afterwards. Similarly, we estimate a smaller convenience yield than the 73 basis points reported by 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 high frequency, which is new in the literature.

Second, we use these unique features of the data to contribute to the literature on monetary policy and quantitative easing event studies. The baseline event study on quantitative easing (Krishnamurthy and Vissing-Jorgensen (2011)) did present spreads between different

yields on safe assets but was constrained to a 2 day event window by slow price discovery, while we are able to use estimates within an hour of all event timestamps. Existing work also does not have any risk free rates inferred from assets outside the fixed income market, so our data is ideal to test 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 circumvent.

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. Their results raise the question of whether such a zero-beta rate loads on other risk factors. The risk-free rate we estimate from options markets circumvents this problem and implies a considerably smaller spread than the estimates in their paper, so the spread we estimate measures the tightness of leverage constraints in any multi-factor generalization of their model.

Also close to our work is Hébert (2018), which presents a theoretical model in which arbitrage spreads are due to constraints on the trading of financial intermediaries. In his model, all constraints are assumed to be due to government policy, and he shows how to infer the government's normative objectives by observing multiple arbitrage spreads. Like the last section of our paper, he considers data on multiple arbitrage opportunities though with a more limited empirical scope.

The last part of paper relates to several existing papers (Amihud and Mendelson (1991), Krishnamurthy (2002), Musto, Nini, and Schwarz (2018), Du, Tepper, and Verdelhan (2018), Daves and Ehrhardt (1993)) that study individual arbitrages that we consider in our analysis across multiple asset classes. While these papers cannot make statements about the relative size and speed of convergence of different arbitrage related spreads that our multi-market

analysis allows for, they do provide additional institutional details about the frictions related to each arbitrage opportunity. The first such paper, Amihud and Mendelson (1991) documents a spread between maturity matched treasury notes and bills and relates it to measures of relative illiquidity. Krishnamurthy shows that spreads between repo rates makes it difficult for a levered investor to profit from the spread between on and off the run bonds. Musto et al. (2018) shows how the relative liquidity (measured using bid ask spreads and other proxies from the microstructure literature) of notes and bonds contributes to the spread between their yields. Du et al. (2018) shows how the size of the covered interest parity spread seems to be influenced by quarterly regulatory reviews of European bankings, suggesting that arbitrage spreads can be influenced by bank regulation as well as showing that they are correlated with some other arbitrage spreads. Daves and Ehrhardt (1993) shows that the spread between interest and principal STRIPS seems related to measures of their degree of illiquidity. Similar to us, Pasquariello (2014) constructs an aggregate index of multiple arbitrage spreads with the purpose of forecasting risky asset returns, while we use ours to measure segmentation between risk free rates in different markets. Finally, Golez, Jackwerth, and Slavutskaya (2018) use a combination of 3-month option and futures data on the S&P500 index to construct a funding illiquidity measure and find that this measure significantly affects the returns of leveraged managed portfolios by hedge funds.

The paper proceeds as follows. In 2, we show how we use the put-call parity relationship on European options to estimate risk free rates. In Section 3 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 4 we explore to what extent the dynamics of the term structure of the convenience yield adds to bond return predictability. In Sections 5-7 we explore mispricing measures from other markets and in Section 8 we explore how convenience yield and mispricing measures comove across various markets. Section 9 concludes.

2. Risk Free Interest Rates without Convenience Yields Estimated from Options Prices

In this section we propose a novel estimator of a term structure of convenience yield free interest rates from options prices. We apply our estimator to data from the CBOE on options quotes and summarize the results.

2.1. 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 $i = 1, \dots, N$. 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 value of the underlying S_t , where we need to adjust the latter for the present value of the cash flow (or convenience) that the security delivers.² Denote this present value of the cash flow (or convenience) by $\mathcal{P}_{t,T}$, then the relationship is given by:

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

This relationship provides two ways of obtaining the risk free interest rate $r_{t,T}$ implied by these markets.

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

$$p_i - c_i = \alpha + \beta K_i + \varepsilon_i \quad (2)$$

where the slope of the line is equal to:

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

²For 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)).

and where the intercept is equal to:

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

The continuously compounded risk free interest rate at time t for maturity T can then be computed from the slope coefficient as follows:

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

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:

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

where

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

and

$$\bar{K} = \frac{\sum_i K_i}{N} \quad (8)$$

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 $i \in 1, \dots, N$. This strategy will deliver the continuously compounded realized 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 $1, \dots, A$ where $A = \frac{N(N-1)}{2}$ and compute an implied risk-free rate for that strike pair. That is, $\forall i \in i = 1, \dots, N$ and $\forall j \in i = 1, \dots, -i, \dots, N$ for which $K_i > K_j$, we compute:

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

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

implied rates:

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

This estimator, which is also 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. That is, 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 . If one holds these positions till maturity, then the payoff is risk free and equal to $K_i - K_j > 0$. 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.

2.2. 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), the Dow Jones Index (DJX), between 2004 and 2018.³ 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). 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 using all strike prices for puts and call with a particular maturity. To compute daily estimates, we then take a median over the minute-level estimates in the day.

2.3. Results

We now describe the results for each of the underlying assets that we study.

³The VIX is another underlying asset that we have data on. The results will be added in the next draft of this paper.

2.3.1. Results S&P 500 Index (SPX)

We start with the results from the S&P 500 index. 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 NSS parameters estimated by Gürkaynak, Sack, and Wright, as well as the continuously compounded version of the LIBOR rate.

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

Zero Coupon Yields: 6 month maturity			
	Mean	St. Dev.	AR(1) (daily)
Option Implied SPX	0.0178	0.0174	0.9995
LIBOR Implied	0.0185	0.0173	0.9999
Government Bond	0.0142	0.0167	0.9998
LIBOR Implied - Option Implied SPX	0.0007	0.0021	0.9638
Option Implied SPX - Government Bond	0.0035	0.0022	0.9607
Zero Coupon Yields: 12 month maturity			
	Mean	St. Dev.	AR(1) (daily)
Option Implied SPX	0.0185	0.0171	0.9980
LIBOR Implied	0.0210	0.0160	0.9998
Government Bond	0.0148	0.0164	0.9997
LIBOR Implied - Option Implied SPX	0.0024	0.0026	0.9148
Option Implied SPX - Government Bond	0.0037	0.0021	0.8738
Zero Coupon Yields: 18 month maturity			
	Mean	St. Dev.	AR(1) (daily)
Option Implied SPX	0.0194	0.0167	0.9996
Government Bond	0.0157	0.0159	0.9996
Option Implied SPX - Government Bond	0.0037	0.0021	0.9774

The table shows that for all maturities the average yield on the SPX implied interest rates are above those of the corresponding government bonds 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 LIBOR rate and the SPX implied rate is positive for both the 6-month and 12-month maturities, equal to 7 basis points and 24 basis points respectively. For the 18-month maturity, a LIBOR rate is not available. Further, the LIBOR rate has the lowest volatility, and the SPX implied rate the highest. The autocorrelation of the spreads are high and typically above 0.9.

To better understand the variation and comovement in the three rates, we plot in Figures I, II and III the three interest rates for all three maturities.




FIGURE I

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

The three graphs shows 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 2009 and 2017 banks faced substantial credit risk, as measured by the spread between LIBOR and the SPX implied zero coupon yield.

As an additional way of summarizing our data, we present an average of daily Nelson-Svensson-Siegel yield curves fit to our SPX implied rates and compare it to the benchmark treasury yield curve of Gürkaynak et al. (2007). In addition, we plot a curve fit to constant




FIGURE II

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




FIGURE III

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

maturity treasury bill rates below. The average spread between our yield curve and the




FIGURE IV

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 V

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

treasury curve is remarkably flat. This treasury yield curve, which is fit to long maturity notes and bonds, over estimates short term treasury bill yields. If we compare to these bill

yields, we add roughly 25 basis points to our convenience yield estimate at short maturities. 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 large convenience yield it earns. In addition, 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 which occurred 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 away from data we actually observed, so it is an open question whether the convenience yields on 10 or 30 year bonds behave similarly.

To further study the dynamics of this credit risk measure and the convenience yield of government bonds, we plot in Figures VI, VII and VIII 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.




FIGURE VI

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.




FIGURE VII

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 VIII

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

Note that for all three maturities, both spreads exhibit large variation, and they both go up during the crisis and (just as with the results for government bond arbitrages in the previous section) have since been reduced to levels closer to zero.

Finally, if no-arbitrage conditions hold perfectly, the R-squared of the regression in equation 2 equals 1. As such, this measure of fit 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:

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

$$= \frac{\beta^2 \text{var}(K)}{\beta^2 \text{var}(K) + \text{var}(\varepsilon)} \quad (12)$$

$$= \frac{1}{1 + \frac{\text{var}(\varepsilon)}{\beta^2 \text{var}(K)}} \quad (13)$$

$$(14)$$

Rewriting this equation, we find:

$$\frac{1}{R^2} - 1 = \frac{\text{var}(\varepsilon)}{\beta^2 \text{var}(K_i)} \approx \frac{\text{var}(\varepsilon)}{\text{var}(K_i)}. \quad (15)$$

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):

$$\sigma(\hat{r}_{t,T}) \equiv \sigma\left(\frac{1}{T} \ln(\beta_{OLS})\right) \approx \sqrt{\frac{1}{NT^2} \left(\frac{1}{R^2} - 1\right)} \quad (16)$$

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 IX 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 16, 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 typical standard error is in the




FIGURE IX

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

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

2.3.2. Results Dow Jones Industrial Index (DJX)

Next, we repeat the analysis that we did for the S&P 500 index for options on the Dow Jones industrial index (DJX). In Table 2 we summarize the results for the median estimator (Estimator 2). The regression-based estimator gives highly comparable results. As for the SPX, we find that the implied interest rate is on average higher than the government bond yield by about 40 basis points, which is invariant to maturity. Also, it is generally lower than the Libor rate.

Table 2
Summary Statistics of DJX Option Implied Interest Rates 2004-2018

Zero Coupon Yields: 6 month maturity			
	Mean	St. Dev.	AR(1) (daily)
Option Implied DJX	0.0184	0.0171	0.9906
LIBOR Implied	0.0186	0.0171	0.9999
Government Bond	0.0144	0.0166	0.9998
LIBOR Implied - Option Implied DJX	0.0002	0.0029	0.6816
Option Implied DJX - Government Bond	0.0040	0.0029	0.6756
Zero Coupon Yields: 12 month maturity			
	Mean	St. Dev.	AR(1) (daily)
Option Implied DJX	0.0190	0.0168	0.9961
LIBOR Implied	0.0211	0.0159	0.9998
Government Bond	0.0150	0.0163	0.9997
LIBOR Implied - Option Implied DJX	0.0021	0.0028	0.8623
Option Implied DJX - Government Bond	0.0040	0.0023	0.7787
Zero Coupon Yields: 18 month maturity			
	Mean	St. Dev.	AR(1) (daily)
Option Implied DJX	0.0197	0.0164	0.9982
Government Bond	0.0159	0.0158	0.9996
Option Implied DJX - Government Bond	0.0039	0.0021	0.8875

Given how comparable the results for the DJX are to the SPX we only plot the implied continuously compounded interest rate the 1-year maturity as an illustration in Figure X. The graphs very much the same pattern, though the implied rates are somewhat noisier than the ones implied by the SPX. Next we repeat the efficiency analysis of Figure IX but




FIGURE X

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

now for DJX. The results are summarized in Figure XI where we plot the standard error of the OLS estimate of equatio 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 over all the minute-level observations, those estimates have smaller standard errors.

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 the table below.

Table 3

Summary Statistics of DJX 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




FIGURE XI

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

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

3. Convenience Yields, Monetary Policy and Quantitative Easing

We 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. Our unique high frequency measure of the term structure of convenience yields is ideal for this purpose and broadens the set of questions that can be studied using high frequency event studies. Existing event studies on the effects of quantitative easing (Krishnamurthy Vissing-Jorgensen (2011)) use two-day event windows because of issues related to slow price discovery. While price




FIGURE XII

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

discovery in treasury bonds themselves is quite fast, more illiquid bonds such as agency debt, corporate debt, or mortgage back securities have posed an issue for high frequency event studies. Because our box rate estimates seem to have price discovery roughly as fast as treasuries, we are able to measure spreads between different risk free rates using a considerably shorter event window than would otherwise be possible. 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.

Finding an effect of central bank policy on convenience yields provides quantitative empirical evidence in favor of a view that liquidity is an important transmission mechanism of monetary policy. An idea going back to the LM curve of the IS-LM model (Hicks) 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 time preference and changes in convenience yields. Our contribution is to present direct evidence of monetary policy on convenience yields because we can focus on a short event window around monetary policy announcements, while existing evidence merely shows that the nominal interest rates are correlated with spreads between different interest rates.

Our results on quantitative easing provide evidence that disciplines our understanding of its transmission mechanism, 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 unclear whether what is bought or what is sold in the transaction determines its effects. One view based on evidence in (Krishnamurthy Vissing-Jorgensen JPE) 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 its transmission mechanism, quantitative easing should not spill over broadly into the interest rates at which the private sector can borrow despite lowering treasury yields. A “broad” view of the transmission mechanism of quantitative easing is common in much of the theoretical literature (Diamond, Caballero Farhi), which emphasizes that swapping reserves for long duration assets increases the overall supply of safe assets and should therefore reduce convenience yields. Because our box rates are inferred from equity option prices, an asset class quite distinct from the fixed income market, our data is ideal for testing whether there is a “narrow” or “broad” transmission mechanism.

3.1. 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. For Q.E. 3, which happened after their paper, we follow the dates in Di Maggio et al. 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. 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. 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 and then fit Nelson-Siegel-Svensson yield curves before and after to these median yields before and after each event. In particular, 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 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 below on 3-month, 12-month, and 30-month yields. We report results on treasury yields, box yields, and the convenience yield equal to 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, the first round of quantitative easing in which our data ranges from November 2008 to March 2009, box yields fell considerably more than treasury yields. In Figure XIII below, we show that 12 and 30 month box yields respectively fell by 88 and 86 basis points, 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 fall only 2 basis points while box yields fall 37 basis points leading to a 36 basis point reduction in the convenience yield. The lack of response in short term government rates is likely due to the fact that 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 speaks strongly against a narrow transmission of quantitative easing, in which asset prices outside of narrowly defined fixed

income markets do not 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 the cost of capital for private firms that issue risky securities 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. This relative scarcity is consistent with our finding that the convenience yield is largest at the shortest maturities.




FIGURE XIII

Effect of QE 1 on government bond yields, box yields, and convenience yields (i.e. the difference between the two) across various 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 XIV. Summing up across all 8 event dates in this period, we find that 3, 12, and 30-month treasury yields fell respectively by 3, 4 and 11 basis points. 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 lead to a 3 basis point decrease in the 3 month convenience yield and respectively a 7 and 3 basis point increase in the 12 and 30 month convenience yield. The aggregate effect of all Q.E. 2 and 3 announcements, while it does decrease treasury yields, is of considerably smaller magnitude

than the effect of Q.E. 1. In particular it seems if anything 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 result 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 too 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 XIV

Effect of QE 2 and 3 on government bond yields, box yields, and convenience yields (i.e. the difference between the two) across various maturities.

3.2. 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 until 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 Kuttner, Nakamura Steinsson). We use the first trade more than 10 minutes before and the first trade more than 20 minutes after each announcement in order 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 after each announcement. We then regress the change in each yield around an announcement on our measure of the monetary shock associated to that announcement and use our estimated regression coefficient to predict the effects of a 100 basis point surprise increase in the federal funds rate in order to report 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. Our first results below (Figure XIV) show the overall results in all of our samples. A 100 basis point rate increase respectively leads to a 54, 88, and 74 basis point increase in box yields. 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 for the 3, 12, and 30 month maturities. 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. Like with quantitative easing, this implies that monetary policy does not have asset price effects only narrowly in the market for money-like securities but spills over into unrelated asset classes like equity options.

Next, we present results in Figure XVI 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




FIGURE XV

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.

of monetary policy on convenience yields. That said, it does seem robustly true that rates in the equity options market move in the same direction as treasury yields and with reasonably large magnitudes, suggesting that monetary policy broadly decreases risk free rates even outside the narrowly defined market for safe, money-like securities.

4. Bond Return Predictability

A literature as early as Fama and Bliss (1987) has focused on the predictability of government bond returns using information contained in the term structure of bond returns. This large 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 time preference and convenience yield. In particular, we




FIGURE XVI

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.

have defined the convenience yield $cy_{t,n}$ as the difference between the implied yield inferred from S&P500 options and the yield on government bonds:

$$cy_{t,n} = y_{t,n}^{box} - y_{t,n}^{gov}, \quad (17)$$

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

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

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

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

can then be written as the difference between two return components, the one related to changes in the box rate and

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

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 time preference?

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 premium related to safe assets (convenience yield).

5. Other Arbitrage Measures: Government Bonds

In the next few sections, we construct several additional arbitrage measures so that we can compare their dynamics to those of the convenience yield.

	$rx_{t+1,2}^{gov}$	$rx_{t+1,2}^{gov}$	$rx_{t+1,2}^{gov}$	$rx_{t+1,2}^{box}$	$rx_{t+1,2}^{box}$	$rx_{t+1,2}^{box}$
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

5.1. Constructing Risk Free Arbitrages

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.

5.1.1. 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 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.⁴ 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.

5.1.2. Coupon Bond Arbitrages

Because the two spreads we study in the previous section 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 which 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 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

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

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

⁴The reason is that all bonds of all maturities pay coupon payments every 6 months and contribute to the coupon-related supply.

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 XVII.

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 et al. (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.

5.2. Data

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. The database provides prices for bills, notes, and bonds at an intraday frequency. 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 4
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

5.3. Results

First, we present summary statistics on the four above-mentioned government bond arbitrages in Table 4 and we plot them in Figures XVII, XVIII and XIX.

Several patterns appear across all of these arbitrage spreads. First, both their level and volatility generally increases 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 XVII correspond to auction cycle dates. Finally and most importantly, the magnitude of these spreads are uniformly smaller than the spreads we estimate in this paper from several other asset classes. This last fact is consistent with the idea that all government bonds benefit from some degree of specialness due to their safety and liquidity, so that all of their yields lie strictly below the risk-free rate that is implied by risky asset prices.




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




FIGURE XVIII
Notes/Bonds Spread in Basis Points per Year.



FIGURE XIX
Median Absolute STRIP Spread.

6. Other Arbitrage Measures: Exchange Rates

6.1. The No-Arbitrage Relation

We now test the no-arbitrage hypothesis in the FX markets. To construct a risk-free strategy for this market, we use the well-known covered interest parity relationship. We compare for a US-based agent at time t two alternative strategies. The first alternative is to invest in a riskless asset denominated in dollars with a time to maturity T . The other is to exchange money into foreign currency, invest into the riskless asset denominated in that foreign currency (with the same time to maturity T) and buy a promise to exchange the money back into dollars at a predetermined rate at T . Covered Interest Parity (CIP) is a no-arbitrage relationship that states that both strategies should earn the same return. More formally, denote by $r_{t,T}$ and $r_{t,T}^*$ the continuously-compounded riskless interest rate at time t which matures at date T in dollars and foreign currency. CIP predicts that

$$e^{Tr_{t,T}} = e^{Tr_{t,T}^*} \frac{S_t}{F_{t,T}}, \quad (21)$$

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

We construct the *cross-currency basis*, in logs, as

$$x_{t,T} = r_{t,T} - r_{t,T}^* - \frac{1}{T} \ln(S_t/F_{t,T}) \quad (22)$$

$\frac{1}{T} \ln(S_t/F_{t,T})$ is the continuously-compounded “forward premium”.

6.2. Data

Our data set is made of all futures trades made between January 2010 and January 2018 on the Chicago Mercantile Exchange (CME) regarding six major currency pairs: British pound (GBP), Euro (EUR), Canadian dollar (CAD), Japanese yen (JPY), Swiss franc (CHF), and New Zealand dollar (NZD). Every day, more than \$100 billion are traded over the CME FX markets, which makes the CME the world’s largest regulated marketplace for foreign currency trading.⁵ Daily spot exchange-rate quotes are from the TrueFX dataset, which offers historical tick-by-tick market data for dealable interbank foreign exchange rates for each millisecond. For spot exchange rate quotes, we take the mid-point between the bid and the ask rates.

We exclude futures with less than 30 days to maturity. We match daily and forward rates at the closest millisecond and keep only trades which happen in a range of 3 seconds one from the other. We construct the forward premium, i.e. $\frac{1}{T} \ln(S_t/F_{t,T})$, and take a daily median of it.

We then merge the forward premia with daily curves of interest rates in different countries. For the US, we use the parameters estimated by Jonathan Wright for the Nielson-Siegel-Svensson (NSS) model fitted on US government-bond yields. Similarly, for the Euro area, we use the NSS parameters estimated by the Bundesbank on German bonds. Those parameters characterize the entire yield curve for all maturities.

⁵More information can be found at <https://www.cmegroup.com/trading/fx/why-trade-fx-futures-and-options.html>

With respect to the other countries, data on constant-maturity zero-coupon yield curves are downloaded from local central banks' websites, except from Japan whose rates are from Bloomberg. We then linearly interpolate constant-maturity zero-coupon yields to match the maturity of the forward contract.

6.3. Results

On average the *cross-currency basis* with respect to the US dollar is close to zero. Table 5 reports key summary statistics for the annualized cross-currency basis of the 6 currencies mentioned above. Cross-country heterogeneity is visible. Canadian dollar, Euro, and British pound have an average CIP deviation between 0.7 and 2.1 basis points. Their volatility is also very low. On the other hand, Japanese yen, Swiss Franc, and New Zealand dollar have a cross-currency basis which is both larger on average (from 10 to 19 basis points) and more volatile.

Table 5
Summary Statistics – Covered Interest Parity

Statistic	Mean	St. Dev.	25th Percentile	75th Percentile	Autocorrelation (1)
GBP	0.68	12.98	−6.82	10.29	0.94
EUR	−1.58	14.63	−10.76	7.61	0.90
CAD	2.14	11.05	−4.30	9.68	0.90
JPY	−15.68	17.64	−24.34	−4.20	0.93
CHF	−18.64	25.27	−31.07	−2.16	0.89
NZD	−10.36	12.33	−17.87	−3.02	0.72

Notes: This table reports key summary statistics for the annualized cross-currency basis, measured in basis points (1 hundreth of a percentage point), for the 6 major currencies. According to Covered Interest Parity, the cross-currency basis must be zero. We construct the *cross-currency basis*, in logs, as

$$x_{t,T} = r_{t,T} - r_{t,T}^* - \frac{1}{T} \ln(S_t/F_{t,T})$$

using high-frequency futures data from the CME between 2010 and 2018.




FIGURE XX

Deviation from Covered Interest Rate Parity. The figure plots the time series for the annualized cross-currency basis, measured in basis points (1 hundredth of a percentage point), for six major currencies. The *cross-currency basis*, in logs, is $x_{t,T} = r_{t,T} - r_{t,T}^* - \frac{1}{T} \ln(S_t/F_{t,T})$.

Figure XX plots the cross-country mean of the CIP deviation for the period post financial crisis. The series exhibit a strong autocorrelation, as also evidenced by the statistics in Table 5.

7. Other Arbitrage Measures: Commodity Markets

7.1. Constructing Risk Free Assets

To construct a risk free asset in commodity markets we use the no-arbitrage relationship between the futures price and the spot price ($F_{t,T}$ and the spot price S_t):

$$F_{t,T} = S_t \exp((r_{t,T} + c_{t,T})T). \quad (23)$$

where $r_{t,T}$ is the 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 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, silver and platinum. The risk-free rate is then computed as:

$$r_{t,T} = \frac{1}{T} \ln \left(\frac{F_t}{S_t} \right). \quad (24)$$

7.2. Data

Our data set contains all futures trades made between May 2007 and January 2018 on the Chicago Mercantile Exchange (CME) regarding three precious metals: gold, silver and platinum. Unlike the CBOE data, which also contains quotes, the database we purchased only contains trades. In a future draft of this paper we intend to extend this evidence to quotes as well.

7.3. Results

In Table 6 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 estimated by Jonathan Wright. The table shows that the estimated convenience yield for government bonds relative to metal-implied interest rates is the same as for our previous

Table 6
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

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. 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.

Finally, as an illustration of the comovement, we plot in figure XXI the implied 12 month gold rate and compare it to corresponding government bond yield (other maturities follow a similar pattern). To reduce estimation noise induced by using trades instead of quotes, we plot the monthly median rate for both series. The plot confirms that the spread between the two rates is positive, and that this spread varies substantially over time, with movements that are not shared by several of the other markets.

8. A Multivariate Analysis and an Aggregate Index

This section presents a multivariate analysis of multiple arbitrage spreads analyzed. This allows us to make progress on several questions that previous studies focusing on individual arbitrages cannot answer. First, we are able to show a relatively strong common component in the size of all arbitrage opportunities by principal component analysis (pca). The first




FIGURE XXI

12 month zero coupon yields implied by gold spot and futures prices and implied by government bond prices.

principal component of the arbitrages we study (plotted in Figure XXII explains 33.62% of the variance in the cross-section of spreads we consider. We interpret this first principal component as a summary of the overall degree of mispricing in financial markets and can therefore be used as an aggregate arbitrage index. Second, we estimate a first order Vector Auto Regression (VAR) to study how each spread evolves after shocks to other spreads occur. We find that the SPX implied rate is our cleanest, least noisy measure of the risk-free rate of return, since all of our other arbitrage spreads seem to feature greater degrees of idiosyncratic risk and volatility.

One key finding in our analysis is that our overall measure of mispricing increases greatly during the U.S. financial crisis and to a lesser degree during the European financial crisis. This reflects the fact that all of the spreads we consider are large during the U.S. crisis while only the option implied rates we consider seem to be seriously exposed to the European crisis. This is reflected in the loadings we report below of each individual series on our principal components. The SPX and DJX implied rates load more heavily on the first principal component than other series, and our higher principal components seem to be exposed to the U.S. crisis but not the European crisis. One interpretation of this finding is that European banks play a more important role in derivatives markets than they do




FIGURE XXII

Aggregate Arbitrage Index. We compute the first principal component across all the arbitrage spreads we have computed across the asset classes including the bond arbitrages, the option arbitrages and commodities.

in the market for other asset classes, perhaps due to less stringent capital requirements on derivatives trades (consistent with findings in Du et al. (2018)).

Table 7
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
$\hat{R}^2$	0.8332	0.9403	0.3494	0.0126	0.5004	0.2147

The VAR analysis also shows that the option implied series are subject to considerably less idiosyncratic risk than others, while the metal convenience yields are subject to more.⁶ As a result, if we are interested in using data on risky asset prices to infer the risk-free rate consistent with investors' time preference, the option implied rates seem to be an ideal candidate. The R-squareds of predicting the in-sample SPX implied rate is extremely high, while that for the DJX is quite high as well. 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. If we tend to think that investors have a relatively stable time preference that does not fluctuate daily, then this series seems to be a very good empirical proxy for it. Another important finding in the vector autoregression is that the option implied rates seem to predict each other, while many series seem to respond to the notes/bonds spread. This seems to be due to the fact that the notes/bond spread spikes dramatically during the U.S. financial crisis but is otherwise very small and that the Box rates are responsive to both the U.S. and European financial crises. The other series we consider seem to have more idiosyncratic variation in their spreads, suggesting that these markets may be relatively more segmented from overall financial conditions.

9. Conclusion

We have constructed and analyzed a novel panel of riskless rates of return that are free of the convenience yield on safe assets. We have presented several important applications of this novel data set including event studies of the effects of central bank policy and bond predictability. More generally, we wish to advocate for its widespread use in the empirical asset pricing (and intermediary asset pricing literatures). For example, our data is important for the accurate measurement of risk premia on stocks and credit instruments, as it prevents researchers from inadvertently confusing the convenience yield on safe assets with compensation for risk in the traditional asset pricing sense of the word.

⁶This is also due to the fact that the convenience yields on precious metals are estimated with trade data only leading to noisier estimates.

References

- Amihud, Y., Mendelson, H., 1991. Liquidity, maturity, and the yields on u.s. treasury securities. *Journal of Finance* 46, 1411–1425.
- Bernanke, B., Kuttner, K., 2005. What Explains the Stock Market’s Reaction to Federal Reserve Policy? *Journal of Finance* 60, 1221–1257.
- 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.
- Daves, P. R., Ehrhardt, M. C., 1993. Liquidity, reconstitution, and the value of u.s. treasury strips. *Journal of Finance* 48, 315–329.
- Diamond, W., 2018. Safety Transformation and the Structure of the Financial System. working paper .
- Du, W., Tepper, A., Verdelhan, A., 2018. Deviations from covered interest rate parity. *The Journal of Finance* 73, 915–957.
- Frazzini, A., Pedersen, L., 2014. Betting against beta. *Journal of Financial Economics* 111, 1–25.
- Golez, B., Jackwerth, J., Slavutskaya, A., 2018. Funding illiquidity implied by s& p 500 derivatives. Working paper.
- 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.

- He, Z., Krishnamurthy, A., 2013. Intermediary asset pricing. *American Economic Review* 103, 732–770.
- Hébert, B., 2018. Externalities as arbitrage. Working paper, Stanford University.
- Krishnamurthy, A., 2002. The bond/old-bond spread. *Journal of Financial Economics* 66, 463–506.
- 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.
- 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.
- Nakamura, E., Steinsson, J., 2018. High Frequency Identification of Monetary Non-Neutrality: The Information Effect. *Quarterly Journal of Economics* 133, 1283–1330.
- Pasquariello, P., 2014. Financial market dislocations. *Review of Financial Studies* 27, 1868–1914.
- Rigobon, R., Sack, B., 2004. The Impact Of Monetary Policy On Asset Prices. *Journal of Monetary Economics* 51, 1553–1575.

- 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.