QE Auctions of Treasury Bonds*

Zhaogang Song† Johns Hopkins Carey Business School

Haoxiang Zhu‡ MIT Sloan and NBER

March 23, 2016

Abstract

This paper studies the cost of the Federal Reserve in implementing quantitative easing through QE auctions. Using a proprietary data set of QE auctions from November 2010 to September 2011, during which the Fed purchased $780 billion Treasury securities, we find that the Fed pays 0.71–2.73 cents per $100 par value above the secondary market ask prices. Moreover, we identify three economic determinants of auction costs across auctions and CUSIPs: cheapness (how much the bond market value is below a yield-curve-model-implied bond value), bond value uncertainty (measured by volatility), and bond scarcity/illiquidity (measured by specialness). A one standard deviation increase in these three measures increases the Fed's cost by 19, 5 and 1.8 cents per $100 par value, respectively. Finally, we document high concentration of profit margin and total profits among a small number of primary dealers.

Key Words: Auction, Federal Reserve, Quantitative Easing, Specialness, Treasury Bond

JEL classification: G12, G13

*For helpful discussions and comments, we thank Tobias Adrian, Jeremy Bulow, Hui Chen, Jim Clouse, Darrell Duffie, Michael Fleming, Glenn Haberbush, Jennifer Huang, Jeff Huther, Ron Kaniel, Leonid Kogan, Arvind Krishnamurthy, Haitao Li, Debbie Lucas, Laurel Madar, Andrey Malenko, Martin Oehmke, Jun Pan, Jonathan Parker, Tanya Perkins, Monika Piazzesi, Simon Potter, Tony Rodrigues, Ken Singleton, Adrien Verdelhan, Jiang Wang, Min Wei, Wei Xiong, Amir Yaron, Rob Zambarano, and Hao Zhou as well as seminar participants at the 2014 NBER Summer Institute (Asset Pricing), CKGSB, Tsinghua PBC, and MIT.

†Carey Business School, Johns Hopkins University, 100 International Drive, Baltimore, MD 21202. E-mail: zsong8@jhu.edu.

‡MIT Sloan School of Management and NBER, 100 Main Street E62-623, Cambridge, MA 02142. zhuh@mit.edu.

QE Auctions of Treasury Bonds

March 23, 2016

Abstract

This paper studies the cost of the Federal Reserve in implementing quantitative easing through QE auctions. Using a proprietary data set of QE auctions from November 2010 to September 2011, during which the Fed purchased $780 billion Treasury securities, we find that the Fed pays 0.71–2.73 cents per $100 par value above the secondary market ask prices. Moreover, we identify three economic determinants of auction costs across time and CUSIP: cheapness (how much the bond market value is below a yield-curve-model-implied bond value), bond value uncertainty (measured by volatility), and bond scarcity/illiquidity (measured by specialness). A one standard deviation increases in these three measures increases the Fed's cost by 19, 5 and 1.8 cents per $100 par value, respectively. Finally, we document high concentration of profit margin and total profits among a small number of primary dealers.

1 Introduction

One of the most significant events in the history of the U.S. Treasury market is the Federal Reserve's large-scale asset purchase programs of long-term Treasury securities since the 2008 financial crisis, commonly known as "quantitative easing" (QE).[1] The QE program was introduced in response to the tightening financial and credit conditions during the financial market turmoil, while the federal funds rate—the usual tool of monetary policy—was stuck at the zero lower bound. Through September 2011, the end of the sample period in our study, the Federal Reserve (Fed) purchased $1.19 trillion of Treasury bonds.[2] These purchases are equivalent to about 28% of the total outstanding stock of these securities at the beginning of the QE program of Treasury securities in March 2009, and about 15% of the total outstanding stock of these securities in September 2011.

While the objective of the QE program is to provide monetary policy stimulus to the economy, the outright purchase operations must be conducted in a manner that encourages competitive pricing to avoid excessive burdens on U.S. taxpayers (Potter (2013)). Therefore, a well-structured mechanism is needed to conduct these large amounts of outright purchases. The mechanism used by the Fed is QE auctions—a series of multi-object, multi-unit, and discriminatory-price auctions, conducted with primary dealers recognized by the Fed. To the best of our knowledge, this paper provides the first empirical analysis of QE auctions.[3]

In particular, we empirically characterize the outcomes of QE auctions and primary dealers' strategic behaviors, guided by auction theory. For example, given the sheer size of the Fed's purchase in a relatively short time window, do QE auctions implement the Fed's purchases efficiently, at reasonable costs?[4] What economic channels determine the auction price markup relative to prevailing secondary market prices across bonds and time? To what

[1] The large-scale asset purchase programs began with the purchasing of agency mortgage-backed securities and agency debt announced in November 2008. Since our study focuses on purchases of Treasury bonds, we shall use QE for purchase operations of Treasury securities throughout the paper.

[2] In this paper, we use "bonds" to refer to Treasury securities with maturity above one year, without distinguishing "Treasury notes" and "Treasury bonds."

[3] A large literature studies the effects of quantitative easing on interest rates, including Krishnamurthy and Vissing-Jorgensen (2011), Gagnon, Raskin, Remanche, and Sack (2011), Hancock and Passmore (2011), Swanson (2011), Wright (2012), D'Amico, English, Lopez-Salido, and Nelson (2012), Hamilton and Wu (2012), Christensen and Rudebusch (2012), Stroebel and Taylor (2012), Bauer and Rudebusch (2013), Li and Wei (2013), D'Amico and King (2013), and Meaning and Zhu (2011), among others. In a certain sense, these studies focus on the effect of QE on interest rates at the "macro" level, whereas we focus on the implementation mechanism of QE at the "micro" level.

[4] For instance, the purchase operations of the $600 billion of the "QE2" program from November 2010 to July 2011 "involve the Federal Reserve purchasing, over an eight-month period, more Treasury securities than the amount currently held by the entire U.S. commercial banking system" (Sack (2011)).

extent do dealers differ in their selling amount, profit margin, and aggregate profit? Answers to these questions would shed light not only on the economics of QE auctions, but also on the intermediation pattern of primary dealers in Treasury markets. These insights, in turn, are useful for understanding the effectiveness of the implementation of U.S. monetary policy.

The auction mechanism. The unique (reverse) auction mechanism used by the Fed to purchase Treasury securities is conducted on its FedTrade system and implemented by the Open Market Trading Desk at the Federal Reserve Bank of New York. Specifically, for each purchase operation, the Fed announces a *range* of total amount and a *maturity bucket* of the Treasury bonds to be purchased, but specifies neither the exact total amount nor the amount for individual bonds. Each operation is organized as a multi-object, multi-unit, and discriminatory-price auction, which allows participants to place multiple offers (up to nine on each bond) across all eligible securities simultaneously. Only primary dealers, the trading counterparties recognized by the Federal Reserve Bank of New York, are eligible to participate in the QE auctions directly, but investors can sell securities to the Fed through the primary dealers.

To evaluate the relative attractiveness of dealers' offers on different Treasury bonds (CUSIPs), the Fed uses the differences of the offered prices and internal spline-based benchmark prices. Therefore, from the perspective of the Fed, different CUSIPs become perfectly substitutable based on the spline benchmark prices, and the QE auction reduces to a regular single-object, multi-unit, and discriminatory-price auction. According to the Federal Reserve Bank of New York, this internal spline-based price is calculated from a spline model fitted through the prices of Treasury securities across CUSIPs (Sack (2011)). The Fed does not publish their spline-based prices, but this spline model is a standard practice in fixed-income markets for evaluating cheapness/richness of different bonds (see the yield curve arbitrage strategy studied in Duarte, Longstaff, and Yu (2007)). Most sophisticated investors in Treasury market should be able to implement a similar spline method to measure the bond values along the curve, though the parameters used could be different from those of the Fed's.

Our data includes the 139 purchase auctions of nominal Treasury securities from November 12, 2010 to September 9, 2011, with a total purchased amount of \$780 billion in par value. This amount includes the entire purchase of the "QE2" program, \$600 billion, as well as the \$180 billion reinvestment by the Fed of the principal payments from its agency debt and agency MBS holdings. The distinguishing feature of our study is the use of detailed data of each accepted offer, including dealers' identities, released by the Fed in accordance with the Dodd-Frank Wall Street Reform and Consumer Protection Act (Dodd-Frank Act), passed in July 2010. This allows us to study not only the auction outcomes at the aggregate level, but also the granular heterogeneity across different bonds and primary dealers.

The cost of the Fed. We measure the costs to the Fed by the difference between the accepted offer prices and the secondary market prices of the same bonds quoted closest to auctions closing time, all weighted by the accepted quantity. On average, the Fed pays 0.71–2.73 cents per $100 par value above the ask prices in the secondary market. Moreover, the auction price markup varies substantially across the 139 QE auctions, with the standard deviation 10-20 cents per $100 par value.

What economic channels drive the large time variation in the Fed's costs in QE auctions? To guide our empirical investigation, we use standard auction theory to capture the strategic considerations of dealers and the auction equilibrium. We identify three important economic channels that affect the auction price markup.

The first channel is the bond "cheapness" relative to a benchmark spline model of the yield curve. Cheapness is a distinguishing feature of QE auctions relative to issuance auctions of government securities (Cammack (1991), Nyborg and Sundaresan (1996), and Simon (1994)). To measure cheapness, we first fit a Svensson (1994) yield curve model at the beginning of each auction day. Then, for each eligible CUSIP, we calculate the difference between the fitted-yield-curve-implied price and the secondary market ask price, in cents per $100 par value. Although we do not observe the Fed's spline model, we expect our cheapness measure to be a significant determinant of auction cost as long as our spline model is sufficiently correlated with that of the Fed's. To the extent that the Fed's spline model is partly anticipated by dealers, theory suggests that bonds that are viewed to be cheaper by the Fed will incur the Fed a higher cost.

The second channel is bond value uncertainty that captures the winner's curse in theories of common value auctions (Wilson (1968); Ausubel, Cramton, Pycia, Rostek, and Weretka (2013)). We measure bond value uncertainty by the volatility of bond returns during the five days before the auction date.

The third channel is bond scarcity/illiquidity that reflects the dealers' costs of sourcing the bonds from the market or customers to deliver to the Fed. We measure scarcity/illiqudity by specialness, the degree to which a particular bond is more costly to borrow in the repo market, as well as bid-ask spread and outstanding balance of a particular bond.

To investigate the economic determinants of auction costs, we run panel regressions of the auction price markup on measures of the three economic channels discussed above. We find that QE auction costs load positively on all three economic channels. The effects are statistically and economically significant, controlling for variations in auction sizes. In particular, a one standard deviation increase in a bond's cheapness, volatility, and specialness (across time and CUSIP) increases the auction price markup by 19, 5, and 1.8 cents per $100 face value, respectively. Moreover, cheapness, as a measure unique to QE auctions, improves regression R^2 by 20% if CUSIP fixed effects are included, and by 80% if they are not.

Dealer profitability. The cost to the Federal Reserve is the profit of primary dealers. Taking advantage of dealer identities in our proprietary dataset, we study the granular heterogeneity of profitability across primary dealers in participating in QE auctions. We measure a dealer' profit margin as the prices of accepted offers less the secondary market prices, weighted by accepted quantities, and obtain the aggregate dollar profit as the product of profit margin and selling amount. This allows us to disaggregate dealers' profits one by one. We find a striking concentration of aggregate profits in a small number of dealers. Among the 20 primary dealers, the top five account for most of the aggregate profits, with large sale amounts and high profit margins (1.67 cents per $100 par value on average). The bottom few dealers sell similarly large amounts of Treasuries to the Fed, but they have low and even negative profit margins (-0.44 cents per $100 par value on average).

What contributes to dealers' heterogeneous profit margins? We run panel regressions of profit margins on cheapness, volatility, and specialness, measured for each winning dealer in each auction. Consistent with the regression results across CUSIPs and auctions, we find that a dealer's profit margin loads positively on the three economic channels. That is, dealers who sell more bonds that are cheap relative to the spline-model-implied price, are more volatile, or more special reap higher profit margins in participating in QE auctions. In particular, a one standard deviation increase in a bond's cheapness, volatility, and specialness increases the auction price markup by 1.4, 2.1, and 0.9 cents per $100 face value.

Related literature. The most relevant paper to our work is Han, Longstaff, and Merrill (2007), who study the Treasury's buyback auctions of long-term debt from March 2000 to April 2002. Our study of QE auctions differs in at least two important aspects. First, though both are discriminatory-price auctions and involve different CUSIPs, QE auctions have an explicitly announced mechanism, whereas the mechanism of the Treasury's buyback auctions is opaque. A transparent and explicit mechanism gives us some guideline to interpret the results. Second, our data include individual accepted offers and dealer identities on each

CUSIP in each auction, which are more comprehensive than the CUSIP-level aggregate auction data used by Han, Longstaff, and Merrill (2007). Dealer-level data enable us to look into the heterogeneity of dealers.

Our paper is also related to the large literature on Treasury issuance auctions, such as Cammack (1991), Simon (1994), and Nyborg and Sundaresan (1996), who use aggregate auction-level data, and Umlauf (1993), Gordy (1999), Nyborg, Rydqvist, and Sundaresan (2002), Keloharju, Nyborg, and Rydqvist (2005), Hortacsu and McAdams (2010), Kastl (2011), and Hortacsu and Kastl (2012), who use bid-level data.[5] QE auctions differ in that each auction involves multiple substitutable CUSIPs, whereas each issuance auction has a single CUSIP at a time. Moreover, our data include dealer identities on each CUSIP in each auction, allowing us to study the heterogeneity of dealers' equilibrium bidding strategies and profitability in a panel data analysis. Studies of issuance auctions either do not have dealer identities or only have masked identifiers.

In addition, a recent innovative paper Pasquariello, Roush, and Vega (2014) studies how the Fed's open market operations in normal times affect the Treasury market microstructure, using only the purchase amount. We differ by focusing on the auction mechanism of the Fed's open market operations for unconventional monetary policy and by using detailed offer-level data of primary dealers.

2 Institutional Background

From November 12, 2010 to September 9, 2011, the Federal Reserve conducted a series of 156 purchase auctions of U.S. Treasury securities, including Treasury notes, bonds, and Inflation-Protected Securities (TIPS). These auctions cover two Fed programs. The first, so-called QE2, is the $600 billion purchase program of Treasury securities, announced on November 3, 2010 and finished on July 11, 2011. The second program is the reinvestment of principal payments from agency debt and agency MBS into longer-term Treasury securities, announced on August 10, 2010, with a total purchase size of $180 billion over our sample period.[6] (We discuss our data and sample period shortly.) These actions are expected to maintain downward pressure on longer-term interest rates, support mortgage markets, and

[5]Theoretical and experimental studies of Treasury issuance auctions include Bikhchandani and Huang (1989), Chatterjee and Jarrow (1998), Goswami, Noe, and Rebello (1996), and Kremer and Nyborg (2004), among others.

[6]Principal payments from maturing Treasury securities are also invested into purchases of Treasury securities in auctions.

help to make broader financial conditions more accommodative, as communicated by the Federal Open Market Committee (FOMC).

The QE auctions are designed as a series of sealed-offer, multi-object, multi-unit, and discriminatory-price auctions. Transactions are conducted on the FedTrade platform. Direct participants of QE auctions only include the primary dealers recognized by the Federal Reserve Bank of New York, although other investors can indirectly participate through the primary dealers. In the first half of our sample period until February, 2, 2011, there were 18 primary dealers, including BNP Paribas Securities Corp (BNP Paribas), Bank of America Securities LLC (BOA), Barclays Capital Inc (Barclays Capital), Cantor Fitzgerald & Co (Cantor Fitzgerald), Citigroup Global Markets Inc (Citigroup), Credit Suisse Securities USA LLC (Credit Suisse), Daiwa Securities America Inc (Daiwa), Deutsche Bank Securities Inc (Deutsche Bank), Goldman Sachs & Co (Goldman Sachs), HSBC Securities USA Inc (HSBC), Jefferies & Company, Inc (Jefferies), J. P. Morgan Securities Inc (J. P. Morgan), Mizuho Securities USA Inc (Mizuho), Morgan Stanley & Co. Incorporated (Morgan Stanley), Nomura Securities International, Inc (Nomura), RBC Capital Markets Corporation (RBC), RBS Securities Inc (RBS), and UBS Securities LLC (UBS). On February 2, 2011, MF Global Inc (MF Global) and SG Americas Securities, LLC (SG Americas) were added to the list of primary dealers, making the total number of primary dealers 20 in the second half our sample period.[7]

Figure 1 describes the timeline of QE auctions, including pre-auction announcement, auction execution, and post-auction information release. To initiate the asset purchase operation, the Fed publishes a pre-auction announcement on or around the eighth business day of each month. The announcement includes an anticipated total amount of purchases expected to take place between the middle of the current month and the middle of the following month.[8] Most importantly, this announcement also includes a schedule of anticipated Treasury purchase operations, including operation dates, settlement dates, security types to be purchased (nominal coupons or TIPS), the maturity date range of eligible issues, and an expected range for the size of each operation. Therefore, the announcement identified the set of eligible bonds to be included as well as the minimum and maximum total par amount (across all bonds) to be purchased in each planned auction. While the purchase amount has

[7] See the website of the Federal Reserve Bank of New York for the historical list of primary dealers.

[8] This amount is determined by the part of the $600 billion purchases that are planned to be completed over the coming monthly period, and the sum of the approximate amount of principal payments from agency MBS expected to be received over the monthly period, and the amount of agency debt maturing between the seventh business day of the current month and the sixth business day of the following month. All the purchases are conducted as one consolidated purchase program.

Figure 1: Example of Timeline of QE Auctions

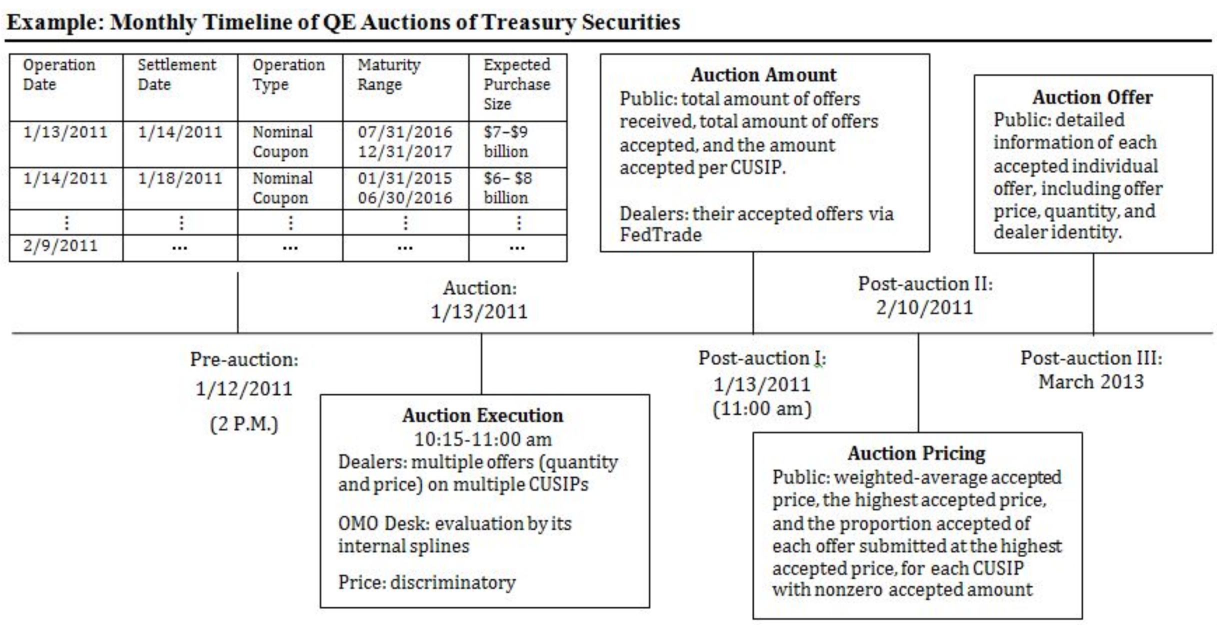

to reach the minimum expected size, the Fed reserves the option to purchase less than the maximum expected size.

On the auction date, each dealer submits up to nine offers per security or CUSIP, with both the minimum offer size and the minimum increment as $1 million. Each offer consists of a price-quantity pair, specifying the par value the dealer is willing to sell to the Fed at a specific price. The auctions happen mostly between 10:15am to 11:00am Eastern Time. Occasionally, the auctions happen between 10:40am and 11:30am, 11:25am and 12:05pm, and 1:15pm and 2:00pm.

In a discriminatory-price auction, offers are either accepted or rejected at the specified prices, and for each accepted offer, the dealer sells its offering amount of bonds to the Fed at its offer price. Since each auction involves a set of heterogeneous securities/CUSIPs, an algorithm is needed to compare the relative attractiveness of offers on different CUSIPs. To make this comparison, the Fed announces that it will compare each offer with a combination of the secondary market prices of similar securities at the close of the auction and its internal spline-based prices (Sack (2011)). Thus, the combination of secondary market prices and the Fed's internal spline-based prices makes different CUSIPs essentially perfect substitutes from the Fed's perspective. From the dealers' perspectives, however, information regarding the Fed's internal spline-based prices is valuable for their strategic bidding across different CUSIPs. Since a spline model of the yield curve is a standard tool in dealer banks, we expect that dealers have some information, if not complete information, about how the Fed evaluates the relative attractiveness of the CUSIPs.

Within a few minutes after the closing of the auction, the Fed announces the auction results publicly on the Federal Reserve Bank of New York website, including the total number of offers received, total number of offers accepted, and the amount purchased per CUSIP. At the same time, participating dealers receive their accepted offers via FedTrade. At the end of each scheduled monthly period, coinciding with the release of the next period's schedule, the Fed publishes certain auction pricing information. The pricing information released includes the weighted-average accepted price, the highest accepted price, and the proportion accepted of each offer submitted at the highest accepted price, for each security purchased in each auction. Finally, in accordance with the Dodd-Frank Act, detailed auction results including the offer price, quantity, and dealer identity for each accepted individual offer will be released two years after each quarterly auction period.

3 Implications of Auction Theory for QE Auctions

QE auctions are multiple-object, multiple-unit and discriminatory-price auctions. To the best of our knowledge, this unique combination of institutional features is not yet addressed in existing auction models. In fact, even for a single-object, multiple-unit auction, multiple Bayesian-Nash equilibria can exist, so that no definitive theoretical predictions can be made about the equilibrium bidding strategies and auction outcomes (see, for example, Bikhchandani and Huang (1993); Back and Zender (1993); Ausubel, Cramton, Pycia, Rostek, and Weretka (2013)). The complications of multiple objects and internal spline-based prices involved in QE auctions make a thorough theoretical treatment of QE auctions much more challenging.

Instead of pursuing a full-fledged theory, which is beyond the empirical focus of this paper, we use the standard theory of single-unit auctions to guide our empirical analysis. Similar approach is employed in empirical studies of Treasury issuance auctions, such as Cammack (1991), Umlauf (1993), Gordy (1999), and Keloharju, Nyborg, and Rydqvist (2005).

Suppose there are N dealers participating in QE auctions. Consider a bond to be purchased by the Fed in QE auctions. Denote dealer i's valuation of the bond by v_i. In general, there are two components in v_i: (i) the common value component that captures the resale value of the bond in the secondary market; and (ii) the private value component that captures a dealer's idiosyncratic cost in obtaining the bond or private information about the bond. In practice, the common-value component is reflected (partly) in the secondary market quotes on various electronic trading platforms available to market participants, such as Bloomberg, BrokerTec, eSpeed, and TradeWeb. The private-value component is affected by a dealer's existing inventory, the cost of financing the bonds in the repo markets, and his relationship with customers and other dealers.

One classical implication from the common-value component is the winner's curse problem: because no dealer is absolutely certain about the resale value of a bond, a dealer worries about buying the bonds too expensively or selling it too cheaply. Applied to QE auctions, a more severe winner's curse problem implies that dealers will submit higher offers to the Fed, leading to a higher expected cost of the Fed. For simplicity, we will not formally reproduce this standard argument here (see Cammack (1991), Umlauf (1993), Keloharju, Nyborg, and Rydqvist (2005), and Han, Longstaff, and Merrill (2007) for more discussions). But the winner's curse channel predicts that a higher bond value uncertainty leads to a higher expected cost of the Fed in QE auctions.

Now, let us focus on the private value components. Since the implication from the common value component is clear, we will assume that the dealers' values $\{v_i\}$ for the bond are all private values, i.e., they are i.i.d. with a distribution function $F : [\underline{v}, \bar{v}] \to [0, 1]$. Note that "i.i.d." should be interpreted as conditional i.i.d., where the conditional information is the common value that depends on the information available to all dealers, such as secondary market price quotes. For simplicity, we normalize the common-value component as zero; this can always be done by shifting the support $[\underline{v}, \bar{v}]$ of $\{v_i\}$.

Suppose the Fed's private valuation for the bond is $v_0 \in [\underline{v}, \bar{v}]$. In the context of QE auctions, v_0 should be interpreted as the relative cheapness of the bond in question, compared to other eligible bonds.[9] Given v_0, among all dealers' offers $\{a_i\}$, the Fed picks the lowest one as long as it is no higher than v_0. If all the offers are higher than v_0, the Fed will not buy this bond. Again, this does not mean that v_0 is the reservation price of the Fed in the traditional sense; rather, the interpretation is that if the offer prices on a bond is too

[9]The Fed ranks offers on different bonds by the difference between the offer prices and its internal spline prices. Consequently, conditional on the offers on other bonds, the multi-object QE auction is equivalent to a single-object auction with the offer prices redefined as the difference between the original offer prices and the Fed's internal spline prices. In this sense, v_0 represents the cheapness of a bond relative to other bonds, based on the Fed's internal spline prices.

high relative to those on other bonds, the Fed will not buy this bond but other bonds. For simplicity, we assume that v_0 is common knowledge, corresponding to the fact that a spline model for the yield curve is a standard tool used by most fixed-income investors, especially the primary dealers.

We conjecture that a dealer's bidding strategy is a monotone increasing function $\beta(\cdot) : v \mapsto \beta(v)$. Since v_0 is known, a dealer cannot sell the bond at any price higher than v_0. Thus, without loss of generality, we can assume $\beta(v) = v$ if $v \geq v_0$.

Now focus on $v < v_0$. Dealer i wins the auction if $a_N < \min_{j \neq i} \beta(v_j)$, which happens with probability $[1 - F(\beta^{-1}(a_i))]^{N-1}$. So, dealer i's expected profit is

$$\Pi_i = (a_i - v_i)[1 - F(\beta^{-1}(a_i))]^{N-1}. \tag{1}$$

By the standard first-order condition, we can solve

$$\beta(v) = v + \frac{\int_{u=v}^{v_0} (1 - F(u))^{N-1} du}{(1 - F(v))^{N-1}}, \quad v \in [\underline{v}, v_0]. \tag{2}$$

Under this strategy, if the Fed accepts the best offer, the Fed's cost is $\min_i \beta(v_i)$. Thus, a higher offer $\beta(\cdot)$ implies a higher expected cost to the Fed.

In this highly stylized model, we have two comparative statics that motivate two economic determinants of the auction costs.

First, all else equal, $\beta(v)$ is increasing in v_0 for $v \in [\underline{v}, v_0]$, for all finite N. This predicts that the auction price is higher if the Fed's value v_0 of the bond is higher. Given the interpretation of v_0 as the relative bond cheapness, this implies that the auction price is higher if the bond looks "cheaper" based on the Fed's spline-implied bond value. Note that the dependence of auction price on v_0 has a premise of N being finite, i.e., the dealers have oligopoly power, which is standard in auction theory.[10]

Second, all else equal, $\beta(v)$ is increasing in v. Intuitively, this prediction says that the auction price is higher if dealers have higher private valuations on the bond. Note that such a prediction does not depend on N being finite. Even if there is perfect competition, dealers will make higher offer prices if the bond is more valuable to them either because of higher costs of acquiring the bond or because of other valuable use of the bonds, such as lending it to their customers. Although private values are unobservable, they may be approximated

[10] We also expect this prediction to hold in models where dealers do not have perfect knowledge about v_0, as long as the dealers' expected value of the bond is positively correlated with v_0. In this case, dealers still optimally offer higher prices on the "favorable" bond of the Fed.

by the scarcity or illiquidity of the bonds.

Overall, we have three conjectured economic determinants of the QE auction cost: (i) bond value uncertainty that captures the winner's curse problem due to the common-value component of the bond, (ii) bond cheapness that captures the Fed's valuation of the bond relative to a yield-curve-model-implied price; and (iii) dealers' private values in sourcing or owning the bond. We expect that the Fed's cost loads positively on each of the three measures.

4 Data and Summary Statistics

In this section, we describe the data and basic summary statistics of the outcomes of QE auctions. Summary statistics include the characteristics of bonds covered in QE auctions, auction size, number of winning offers, and number of winning dealers.

4.1 Data

Our sample period is from November 12, 2010 to September 9, 2011. We focus on this sample period because of data availability and the impact of monetary policy developments on the timeline of auctions. The Dodd-Frank Act, passed on July 21, 2010 by the U.S. Congress, mandates that the Federal Reserve should release detailed auction data to the public, but with a two-year delay after each quarterly operation period. As of January 2014, detailed data of dealer offers are only available from July 22, 2010 to December 31, 2011. We discard the period July 22, 2010–November 11, 2010 because no expected purchase sizes at the CUSIP level were announced by the Fed in this period. We also discard the period September 10, 2011–December 31, 2011 because the Fed announced its new policies on September 21, 2011, including the Maturity Extension Program and changes in the agency debt and agency MBS reinvestment policy. Thus, in our sample period, the Fed purchased only Treasury securities for its QE policy.

We focus on 139 auctions of nominal Treasury securities among the 156 purchase auctions between November 12, 2010 and September 9, 2011. The excluded 17 auctions of TIPS only account for 3% of the total purchases. Moreover, focusing on auctions of nominal bonds makes the bond characteristics, such as coupon rates and returns, comparable across CUSIPs. These 139 auctions were conducted on 136 days, with two auctions on November 29, 2010, December 20, 2010, and June 20, 2011, and only one auction on all the other days.

Our empirical analysis combines the auction data released by the Federal Reserve and three CUSIP-level datasets of Treasury securities, including the secondary market intraday price quotes, the specific collateral repo rates, and the outstanding quantity. The auction data include: (1) the expected total purchase size range, the total par value offered, and the total par value accepted for each auction; (2) the indicator of whether a CUSIP was included or excluded in the auctions; (3) for bonds included in the auctions, the par value accepted, the weighted average accepted price, and the least favorable accepted price for each CUSIP in each auction; and (4) the offered par value, offer (clean) price, and dealer identity for each accepted offer on each CUSIP in each auction.

To the best of our knowledge, the individual-offer data we use is the first set of auction data at the individual bid level that has been ever analyzed for U.S. Treasury securities. Previous studies of issuance auctions and buyback auctions of U.S. Treasury securities, including Cammack (1991), Simon (1994), Nyborg and Sundaresan (1996), and Han, Longstaff, and Merrill (2007), have used data only at the aggregate auction level or at the CUSIP level at best.[11]

Our secondary market price data contain indicative bid and ask quotes from the New Price Quote System (NPQS) by the Federal Reserve Bank of New York, as well as the corresponding coupon rate, original maturity at issuance, and remaining maturity, which are also used by D'Amico and King (2013). We choose the NPQS quotes because the these data also cover off-the-run securities, many of which are traded in QE auctions. (The BrokerTec data used in recent studies such as Fleming and Mizrach (2009) and Engle, Fleming, Ghysels, and Nguyen (2012) mainly contain prices of mostly on-the-run securities.) The NPQS data have four pairs of bid and ask quotes each day, which are the best bid and ask prices across different trading platforms of Treasury securities made at 8:40am, 11:30am, 2:15pm, and 3:30pm.

We obtain the CUSIP-level special collateral repo rates from the BrokerTec Interdealer Market Data that averages quoted repo rates across different platforms between 7am and 10am each day (when most of the repo trades take place). We then calculate the CUSIP-level repo specialness as the difference between the General Collateral (GC) repo rate and specific collateral repo rate, measured in percentage points. This specialness measure reflects the value of a specific Treasury security used as a collateral for borrowing (see Duffie (1996); Jordan and Jordan (1997); Krishnamurthy (2002); Vayanos and Weill (2008)). We also

[11] Studies of Treasury auctions in other countries have used bid-level data, such as Umlauf (1993), Gordy (1999), Nyborg, Rydqvist, and Sundaresan (2002), Keloharju, Nyborg, and Rydqvist (2005), Hortacsu and McAdams (2010), Kastl (2011), and Hortacsu and Kastl (2012).

obtain the outstanding par value of Treasury securities each day from the Monthly Statement of the Public Debt (MSPD) of the Treasury Department.

4.2 Summary statistics: Bonds covered in QE auctions

What Treasury securities does the Fed buy and what is the allocation of purchasing quantities across different securities? Table 1 reports the maturity distribution of planned purchases in QE Auctions of Treasury debt over our sample period, announced on November 3, 2010 by the Fed. Only 6% of planned purchase amounts have a maturity beyond 10 years. According to the Fed, this maturity distribution has an average duration of between 5 and 6 years for the securities purchased. The Fed does not purchase Treasury bills, STRIPS, or securities trading in the when-issued market.[12]

Table 1: Maturity Distribution of Planned Purchases in QE Auctions

Maturity Sectors for QE Auctions of Treasury Securities								
	Nominal Coupon Securities							TIPS
Maturity Sector (Years)	1.5–2.5	2.5–4	4–5.5	5.5–7	7–10	10–17	17–30	1.5–30
Percentage	5%	20%	20%	23%	23%	2%	4%	3%

Note: This table shows the maturity distribution of planned purchases in QE Auctions of Treasury debt over our sample period (November 12, 2010–September 9, 2011), announced on November 3, 2010 by the Fed. The on-the-run 7-year note will be considered part of the 5.5- to 7-year sector, and the on-the-run 10-year note will be considered part of the 7- to 10-year sector.

Panel A of Table 2 presents descriptive statistics on the number of bonds for the 139 QE auctions (of nominal Treasury securities) in our sample period. The number of eligible bonds in an auction varies between 15 and 36, with a mean of 26. On average, one CUSIP is excluded in each auction, with the minimum and maximum number of excluded CUSIPs being 0 and 4, respectively. This leaves the number of eligible (included) bonds between 13 and 34, averaging 25 per auction. Among these included bonds, in each auction the Fed purchases between 2 and 26 bonds, with an average of 14 bonds. On average, in each auction, 11 eligible bonds are not purchased by the Fed. Across all 139 auctions, only 186 CUSIPs have ever been purchased by the Fed, among the 215 eligible CUSIPs.

According to Fed communications to the public, excluded bonds are those trading with heightened specialness in the repo market, or cheapest to deliver into the front-month Trea-

[12]See http://www.newyorkfed.org/markets/lttreas_faq_101103.html for details.

Table 2: Summary Statistics of Bonds in QE Auctions

A: Number of Bonds in QE Auctions					
	Eligible Bonds	Excluded Bonds	Included Bonds	Included Bonds (not purchased)	Included Bonds (purchased)
Mean	26	1	25	11	15
Std	6	1	6	7	5
Min	15	0	13	0	3
Q-25%	20	1	19	5	12
Q-50%	27	1	26	9	15
Q-75%	31	2	29	16	19
Max	36	4	34	28	27
B: Regression of Inclusion Dummy					
Maturity	-0.078**			-0.066**	
	(-3.445)			(-2.837)	
CP		0.072		0.072	
		(1.374)		(1.393)	
Specialness			-7.376**	-7.639*	
			(-2.683)	(-2.554)	
N	3,127	3,127	3,127	3,127	
R^2	0.030	0.030	0.054	0.058	
CUSIP FE	No	No	No	No	
Auction FE	Yes	Yes	Yes	Yes	

Note: Panel A of this table presents descriptive statistics, including mean, standard deviation (std), minimum (min), maximum (max), and quartiles of the number of bonds covered in the 139 QE auctions of nominal Treasury securities, from November 12, 2010 to September 9, 2011. Panel B reports panel Logit regressions of whether a bond is included in an auction (yes= 1, no= 0), with z-statistics in the parentheses. The auction fixed effects are included. The right-hand variables are Specialness of a bond (the general repo rate minus special repo rate on the bond) in percentage points, the coupon rate CP in percentage points, and the Maturity in years. Significance levels: ** for $p < 0.01$, * for $p < 0.05$, and $^{+}$ for $p < 0.1$, where p is the p-value.

sury futures contracts.[13] Presumably, these bonds are excluded to avoid creating or exacerbating supply shortages in repo and futures markets. To formally explore the Fed's decision to include or exclude certain CUSIPs, we estimate a panel logit regression as follows:

[13]The Fed also excludes CUSIPs by the size limit for purchase amount per security according to the percentage of the outstanding issuance and the Fed's existing holdings of this security. See the website of Federal Reserve Bank of New York for details (http://www.newyorkfed.org/markets/lttreas_faq_101103.html). We do not study this criterion as the Fed purchase rarely hit the size limit in our sample. In addition, communications with the Fed confirm that primary dealers have almost perfect foresight about which securities will be excluded before the auction.

$$ln[I_{tj}/(1 - I_{tj})] = \sum_t \alpha_t D_t + \beta_1 \cdot Specialenss_{tj} + \beta_2 \cdot Maturity_{tj} + \beta_3 \cdot CP_j + \epsilon_{it}, \quad (3)$$

where I_{tj} is the probability of bond j being included in the t-th auction, $Maturity_{tj}$ is the time to maturity measured in years, and CP_j is the coupon rate in percentage points. Because our objective here is to check what type of bonds the Fed includes in an auction, we include auction fixed effects D_t.

Panel B of Table 2 reports the results from the panel Logit regression (3), as well as regressions with only one of $Specialenss_{tj}$, $Maturity_{tj}$, and CP_j on the right-hand side. We observe that the regression coefficient of Specialness is significantly negative, confirming that the Fed did exclude bonds trading at heightened specialness. Moreover, the Fed also tends to exclude bonds with longer maturity on average, which is intuitive because, for a fixed auction and hence a maturity bucket, longer-maturity bonds tend to be closer to on-the-run.

4.3 Summary statistics: Auction size, offer, and dealer

In this section, we present summary statistics of auction size, number of winning offers, and number winning dealers. Panel A of Table 3 shows that the par amounts of submitted offers vary between \$4 and \$43 billion, averaging about \$21 billion per auction, while the offer amount accepted by the Fed varies between \$0.7 and \$8.9 billion, averaging \$5.6 billion per auction. The ratio between submitted and accepted offer amounts (offer-to-cover) is on average 4.2, with a range of 1.7 to 26.2. The average expected minimum and maximum auction sizes are \$4.6 billion and \$6.2 billion, and the accepted offer amount always falls between the expected minimum and maximum auction sizes. In addition, the offer amount per included bond is \$0.84 billion on average, while the accepted offer amount per accepted bond is \$0.41 billion.

Panel B of Table 3 presents summary statistics on the number of winning offers and dealers. We observe that the number of winning offers ranges between 8 and 326, with a mean of 103, whereas the number of winning dealers ranges between 4 and 20, with a mean of 16. (In our sample period, the total number of primary dealers was 18 before February 2, 2011 and 20 afterwards.) As a result, per auction, each winning offer has an average size of \$0.07 billion, and each winning dealer sells \$0.36 billion to the Fed on average.

Table 3: Auction Size, Offer, and Dealer

A: Auction Size

	Offer Amount ($ billion)	Purchase Amount ($ billion)	Offer to Cover Ratio	Announced Min Size ($ billion)	Announced Max Size ($ billion)	Offer Amount per Included Bond ($ billion)	Purchase Amount per Accepted Bond ($ billion)
Mean	20.76	5.59	4.22	4.58	6.2	0.84	0.41
Std	8.81	2.33	2.87	1.98	2.48	0.31	0.29
Min	4.13	0.72	1.67	0.50	1	0.12	0.04
Q - 25%	13.53	3.16	3.06	2.75	3.5	0.59	0.23
Q - 50%	21.04	6.68	3.66	5	7	0.83	0.36
Q - 75%	27.70	7.30	4.64	6	8	1.05	0.52
Max	42.88	8.87	26.18	7	9	1.76	2.57

B: Number of Winning Offers and Dealers

	Number of Winning Offers	Number of Winning Dealers	Purchase Amount Per Winning Offer	Purchase Amount Per Winning Dealer
Mean	103	16	0.07	0.36
Std	54	3	0.06	0.17
Min	8	4	0.01	0.06
Q - 25%	57	15	0.04	0.24
Q - 50%	102	17	0.05	0.38
Q - 75%	137	18	0.08	0.45
Max	326	20	0.49	1.38

Note: This table presents descriptive statistics on the auction size (in Panel A) and the number of winning offers and dealers (in Panel B) for the 139 QE auctions of nominal Treasury securities, from November 12, 2010 to September 9, 2011. We report mean, standard deviation (Std), minimum (Min), Maximum (Max), and quartiles at the auction level.

5 Costs of the Federal Reserve

In this section, we study the cost of the Federal Reserve in executing QE auctions. We first present summary statistics of the auction costs and then study what economic channels affect auction costs using panel regressions.

5.1 How much does the Fed pay?

Following the literature,[14] we measure the auction cost as the auction price markup, namely the difference between the auction price and secondary market price on the days the auctions are executed.[15] Let $p_{t,j,d,o}$ and $q_{t,j,d,o}$ be the o-th winning offer price and par value from dealer d on CUSIP j in auction t, respectively, and let $P_{t,j}$ be the secondary market price of the CUSIP j at the time auction t is closed. Then the auction price markup of auction t is

$$Markup_t = \frac{\sum_{j,d,o} (p_{t,j,d,o} - P_{t,j}) \cdot q_{t,j,d,o}}{\sum_{j,d,o} q_{t,j,d,o}} \tag{4}$$

That is, the auction price markup is the average, weighted by the amount of each accepted offer, of the differences between the offer price and the corresponding secondary market price of the bond for that offer at the time the auction is closed.

Which secondary-market prices do we use for $P_{t,j}$? Recall that we have bid and ask prices quoted at 8:40am, 11:30am, 2:15pm, and 3:30pm from the NPQS data. Since the auction close time is one of 11:00am, 11:30am, 12:05pm, and 2:00pm, we define the Contemporaneous Ask price for each purchased bond in each auction as the 11:30am ask price for the first three auction closing times and the 2:15pm ask price for the last auction closing time. Using this Contemporaneous Ask, we compute the auction price markup based on (4), denoted as $Markup_t^c$. We also compute alternative measures of the auction price markup using the 2:15pm ask price, the 3:30pm ask price, and the average of the 2:15pm and 3:30pm ask prices. The costs measured relative to these benchmark prices, especially the 3:30pm ask, capture more of the so-called realized costs that can be interpreted as the revenue of dealers

[14] See, for example, Han, Longstaff, and Merrill (2007), Cammack (1991), Nyborg and Sundaresan (1996), and Hortacsu and Kastl (2012), among others.

[15] We also measure the cost of purchasing a bond as the difference between the worst price accepted by the Fed (also known as the stop-out price) and the corresponding secondary market price. This cost measure quantifies the maximum price the Fed is willing to tolerate to achieve its minimum purchase amount. The cost measure based on stop-out price is around 2 cents per $100 par value higher than that based on the weighed average price, and the correlation between the two measures is as high as 99%. Because of their high correlation, we only report results based on the average purchase price.

Table 4: Auction Cost

	Quantity-Weighted Mean	Quantity-Weighted Stanard Deviation	Min	Max
Cost to Contemporaneous Ask	0.71	11.05	-43.12	46.46
Cost to 2:15 PM Ask	2.73	21.89	-92.08	79.19
Cost to 3:30 PM Ask	2.15	26.50	-125.13	102.74
Cost to Average PM Ask	2.44	23.31	-108.60	83.28

Note: This table presents summary statistics of the auction cost (in cents per $100 par value) to the Fed in the 139 QE auctions from November 12, 2010 to September 9, 2011. The auction cost (auction price markup) is computed as the average, weighted by the amount of each accepted offer, of the differences between the offer price and the corresponding secondary market price of the bond for that offer at the time the auction is closed. We use the Contemporaneous Ask, the 2:15pm ask, and the 3:30pm ask, as well as the average of the last two as the secondary market price. Contemporaneous Ask of a bond equals the 11:30am secondary market ask price for the auctions closed at 11:00am, 11:30am, 12:05pm, and the 2:15pm ask price for the auctions closed at 2:00pm. The quantity-weighted mean and standard deviation are computed as the mean and standard deviation of the auction price markup across the 139 QE auctions, weighted by auction sizes.

by entering a sales contract with the Fed first and then covering the short positions at the subsequent market price (Hasbrouck (2007)).

Table 4 presents summary statistics of the auction cost $Markup_t$, in cents per $100 par value, to the Fed for the 139 auctions in our sample. The average cost over all purchase auctions, weighted by auction size, is 0.71 cents measured relative to the Contemporaneous Ask and reaches up to 2.73 cents measured to the 2:15pm and 3:30pm ask prices in the secondary Treasury market. To put the magnitude of $Markup_t$ into perspective, the weighted average bid-ask spread for the purchased bonds (weighted by the par value purchased) is 2.56 cents per $100 par value during our sample period. Therefore, the average QE auction cost is about 28% to 107% of the average bid-ask spread. These results suggest that the Fed does not suffer large market-impact costs in purchasing the huge amount ($780 billion) of Treasury securities in QE auctions in our sample period.

It is also informative to compare the QE auction cost to the costs reported in the literature of other Treasury security auctions. Han, Longstaff, and Merrill (2007) report that the Treasury's buyback auctions from March 2000 to April 2002 incur an average cost of 4.38 cents per $100 par value, which is about 70% of the average bid-ask spread of the auctioned bonds. Given that the average par value per auction of these buyback operations is $1.5 billion, which is only about 1/3 of the $5.6 billion (from Table 3) in the QE auctions, the Fed's QE auction mechanism seems to be comparable to the Treasury's buyback auction mechanism in executing large purchases effectively. The average cost in QE auctions also compares well with those in issuance auctions of Treasury securities estimated by prior studies. For example, among others, Goldreich (2007) estimates that the average issuance cost of Treasury notes and bonds from 1991 to 2000 is about 3.5 cents per $100 par value, while Cammack (1991) and Nyborg and Sundaresan (1996) provide similar estimates for T-bill issuance auctions.

While the average auction cost may appear small or moderate, the aggregate dollar cost is about $55 million at 0.71 cents per $100 par value, and 213 million at 2.73 cents per $100 par value. Moreover, the dispersion of costs across auctions is substantially larger. Figure 2 shows that $Markup_t^c$ is strongly time-varying. Across the 139 QE auctions, the minimum and maximum of $Markup_t^c$ are -43.12 cents and 46.46 cents per $100 par value, respectively. Table 4 reports that the quantity-weighted standard deviations of the QE auction price markup is 11.05 cents per $100 par value, if measured relative to the Contemporaneous Ask. This standard deviation is 26.5 cents per $100 par value, if the cost is measured relative to the 3:30pm ask prices. These large variations of the auction price markup naturally leads to the question of what economic channels drive them, which we investigate in the next subsection.

5.2 What determines auction costs?

What economic forces explain the variations of the QE auction cost? Our analytic framework in Section 3 suggests three economic determinants of the equilibrium behaviors in QE auctions and hence the auction price markup: bond cheapness, bond value uncertainty, and bond scarcity/illiquidity. We first describe empirical measures of these three economic channels and then present the regression results.

5.2.1 Empirical measures

To measure the bond cheapness, we first back out, on each auction day, a smooth zero-coupon yield curve based on observed market prices. This yield curve is then used to price all bonds purchased in the auction on the auction day. A bond's cheapness is calculated as its theoretical price based on the fitted curve minus the market ask price.

Various estimation methods can be employed to fit the zero-coupon yield curve from Treasury securities. Similar to Gurkaynak, Sack, and Wright (2007) and Hu, Pan, and Wang (2013), we use the the Svensson (1994) functional form, which is an extension of Nelson and

Figure 2: Time Series of Auction Cost

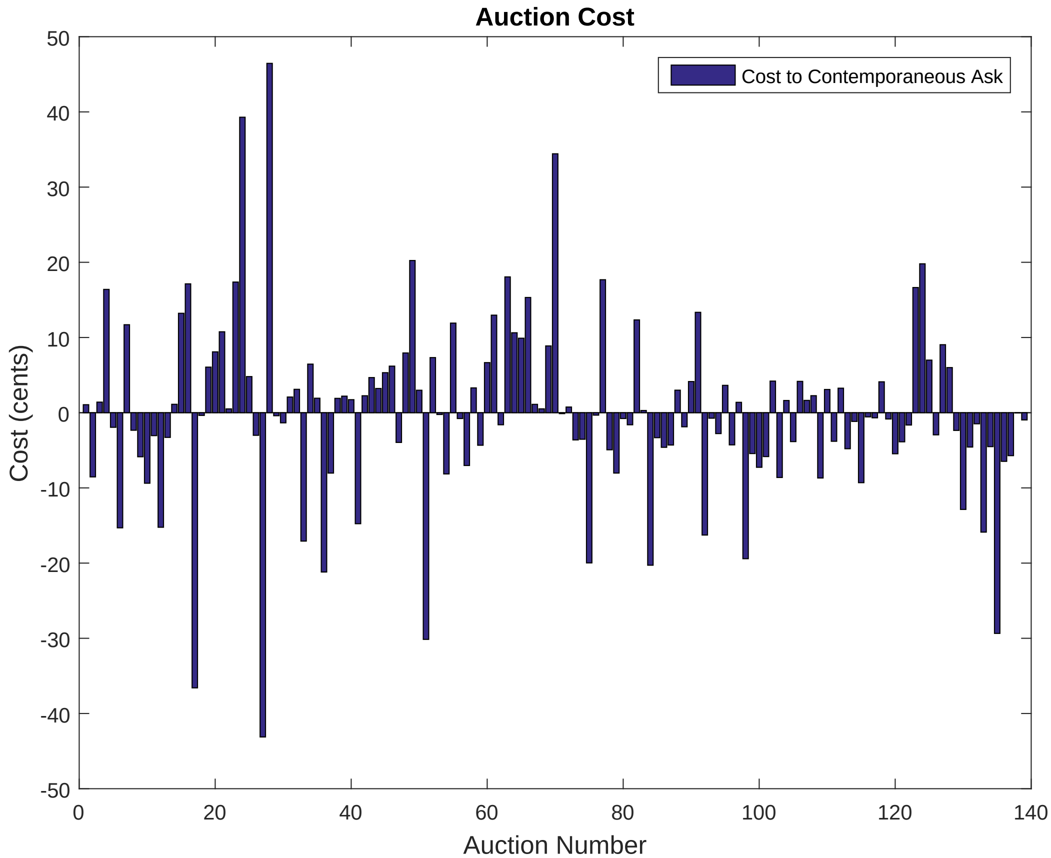

Note: This figure presents the time series of the auction cost (in cents per $100 par value) to the Fed in the 139 QE auctions from November 12, 2010 to September 9, 2011. The auction cost (auction price markup) is computed as the average, weighted by the amount of each accepted offer, of the differences between the offer price and the corresponding Contemporaneous Ask of the bond for that offer. Contemporaneous Ask of a bond equals the 11:30am secondary market ask price for auctions closed at 11:00am, 11:30am, 12:05pm, and the 2:15pm ask price for auctions closed at 2:00pm.

Siegel (1987) and which assumes that the instantaneous forward rate follows the following specification:

$$f(m) = \beta_0 + \beta_1 \exp\left(-\frac{m}{\tau_1}\right) + \beta_2 \frac{m}{\tau_1} \exp\left(-\frac{m}{\tau_1}\right) + \beta_3 \frac{m}{\tau_2} \exp\left(-\frac{m}{\tau_2}\right), \qquad (5)$$

where m is the time to maturity, and β_0, β_1, β_2, β_3, τ_1, and τ_2 are parameters to be estimated. These parameters must satisfy certain regularity conditions, including $\beta_0 > 0$, $\beta_0 + \beta_1 > 0$,

$\tau_1 > 0$, and $\tau_2 > 0$. See Svensson (1994), Gurkaynak, Sack, and Wright (2007), and Hu, Pan, and Wang (2013) for details on the interpretations of these parameters in terms of the yield curve.

For each set of parameters, $\theta \equiv (\beta_0, \beta_1, \beta_2, \beta_3, \tau_1, \tau_2)$, we compute the corresponding zero-coupon yield curve by integrating the forward rates based on (5), which can then be used to price any outstanding Treasury security with specific coupon rates and maturity dates. To estimate the yield curve, we choose the parameters to minimize the weighted sum of the squared deviations between the actual market prices of Treasury securities and the Svensson model implied prices, with the weights chosen as the inverse of the duration of each individual bond. The market prices we use are the mid-quotes at 8:40am that are available before the auction is closed (note that most of the auctions are close at 11:00 am, as discussed above). Specifically, let P_t^j be the market price of the bond j available on auction day t, $j = 1, \cdots, N_t$. We choose the parameter θ as

$$\theta_t = \arg\min_{\theta} \sum_{j=1}^{N_t} \left[\left(P_t^j(\theta) - P_t^j \right) / D_t^j \right]^2, \tag{6}$$

where $P_t^j(\theta)$ is the model-implied price of bond j based on the Svensson model in 5, and D_t^j is the bond duration. With the parameter estimate θ_t, we can compute the model-implied price as $P_t^j(\theta_t)$ and define the bond cheapness as

$$Cheapness_t^j = P_t^j(\theta_t) - P_t^j. \tag{7}$$

The higher is $Cheapness_t^j$, the lower the raw market price is relative to the theoretical price based on the model estimated using all available bond prices, and the cheaper the bond j is according to its quoted market price.

To measure bond value uncertainty, we follow the literature to use the pre-auction volatility VOL_t^j computed as the standard deviation of daily returns of bond j during the five trading days prior to the auction date t. The higher is VOL_t^j, the high uncertainty market participants face about the fundamental value of this bond.

Finally, we consider three measures of bond scarcity/illiqudity: specialness, outstanding balance, and bid-ask spread. Specialness is the difference between the general repo rate and special repo rate on the CUSIP, in percentage points. The higher the specialness, the more costly it is for a dealer to locate or finance the bond in the repo market. Following Han, Longstaff, and Merrill (2007), we define outstanding Balance (OB_t^j) as the total outstanding par value of the bond j, normalized by the total purchase amount of auction t, to control for the potential proportionality of auction size to the outstanding amount. Bid-ask spread ($Bid - Ask_t^j$) is the difference between the ask and bid quotes of bond j, normalized by the mid-quote, denoted in basis points. We use the 8:40am NPQS quotes when computing Bid-Ask, similar to the calculation of Cheapness.

Table 5: Summary Statistics of Empirical Measures

	Cheapness	VOL	Specialness	OB	Bid-Ask
Mean	-1.398	0.370	0.024	5.940	3.059
Standard Deviation	2.042	0.237	0.020	5.963	1.173
Skewness	-1.543	1.290	2.580	6.328	0.623
Kurtosis	3.953	5.193	14.138	48.084	2.552
Min	-7.233	0.043	-0.002	1.858	0.786
Max	0.059	1.328	0.148	53.52	6.007

Note: This table reports summary statistics of empirical measures of the economic forces that underlie the QE auctions. For each bond in each auction, Cheapness is the difference between the model price implied from the fitted zero-coupon yield curve based on the Svensson model and the actual market price (in cents per $100 dollar face value), VOL is the standard deviation of daily returns of the bond during the five trading days prior to the auction date (in percentage points), Specialness is the difference between the general repo rate and special repo rate on the bond (in percentage points), OB is the total outstanding par value of the bond normalized by the total purchase amount of the auction, and Bid-Ask is the difference between the secondary market ask and bid quotes of a bond normalized by the mid-quote (in basis points). We then compute the weighted averages of these five measures across bonds, with the weight as the purchase amount of each bond, for each of the 139 nominal QE auctions from November 12, 2010, to September 9, 2011.

Table 5 reports the summary statistics of these five empirical measures across auctions. We first calculate these measures for each CUSIP in each auction, and then compute the weighted averages of these five measures across bonds, with the weight as the purchase amount of each bond, within each auction. We observe that the purchased bonds by the Fed have a cheapness of -1.40 cents on average, implying that the Fed bought bonds that are more expensive than the Svensson model price in general. The daily pre-auction volatility is 0.37% on average, and the mean specialness is about 2.4 basis points. The mean of (normalized) outstanding balance is 5.94, suggesting that the purchase amount of QE auctions is 1/6 of the outstanding balance on average. The average bid-ask spread is about 3 basis points for the bonds purchased by the Fed. In addition, the skewness is positive for all variables but Cheapness, and there is excessive kurtosis for Specialness and outstanding balance.

5.2.2 Regression results

In this section, we study how, if any, QE auction cost depends on bond cheapness, value uncertainty, and scarcity/illiquidity by regressing the auction price markup on empirical measures of the three economic channels. As mentioned above, to be conservative, we mainly use the auction price markup measured to the Contemporaneous Ask in our baseline results. For each bond j in each auction t, the markup is defined as

$$Markup_{t,j} = \frac{\sum_{d,o} (p_{t,j,d,o} - P_{t,j}) \cdot q_{t,j,d,o}}{\sum_{d,o} q_{t,j,d,o}}, \tag{8}$$

where we use Contemporaneous Ask for $P_{t,j}$.

Our empirical analysis is based on the following panel regression of the auction cost for each bond in each auction on empirical measures of the bond's cheapness, value uncertainty, and scarcity/illiquidity:

$$\begin{aligned} Markup_{t,j} &= \sum_j \alpha_j D_j + \beta_1 \cdot \text{Cheapness}_{tj} + \beta_2 \cdot \text{VOL}_{tj} \\ &\quad + \beta_3 \cdot \text{Specialness}_{tj} + \beta_4 \cdot \text{OB}_{tj} + \beta_5 \cdot \text{Bid-Ask}_{tj} \\ &\quad + \beta_5 \cdot \text{Size}_t^E + \epsilon_{it}, \end{aligned} \tag{9}$$

where CUSIP fixed effect D_j controls for potential unobservable effects generic to individual bonds. We do not include auction fixed effects because our main objective is to study the time variations of costs across auctions, similar to the literature (such as Han, Longstaff, and Merrill (2007), Nyborg and Sundaresan (1996), and Hortacsu and Kastl (2012), among others). Nonetheless, in order to control for the potential size effect at the aggregate auction level, we include the expected auction size Size_t^E, computed as the mean of the announced minimum and maximum of intended purchase amount divided by the number of included CUSIPs, in \$billions of par value.[16]

Columns (1)–(5) of Table 6 report results of panel regressions of the auction price markup $\text{Markup}_{t,j}^c$ on each of the empirical measures of bond cheapness, bond value uncertainty, and bond scarcity/illiquidity separately, as special cases of regression 9. We find that Cheapness, VOL, Specialness, and outstanding balance are all highly significant, implying that the QE

[16]To control for the possibility that dealers learn about the bond values or the Fed's internal spline-based prices over time, we also include the auction number in the OLS regression as an explanatory variable (Milgrom and Weber (1982), Ashenfelter (1989), and Han, Longstaff, and Merrill (2007)). We do not find any time trend in the data.

auction cost depends on all three economic channels. The result in Column (6) on the multivariate panel regression (9) further confirms the significance of Cheapness, VOL, and Specialness (the significance of outstanding balance is weaker).

The effects are also economically large. For example, according to column (6), a two dollar increase in Cheapness, which is roughly one standard deviation of Cheapness across auctions and CUSIPs in our sample,[17] increases the auction price markup by about 19 cents (= 200 × 0.095) per \$100 face value; a 0.23 percentage point increase in VOL, which is roughly one standard deviation of the pre-auction volatility across auction and CUSIP in our sample, increases the auction price markup by about 5 cents (= 0.23 × 22.04) per \$100 face value; a 0.03 percentage point increase in specialness, which is roughly one standard deviation of the specialness across auction and CUSIP in our sample, increases the auction price markup by about 1.8 cents (= 0.03 × 60.56) per \$100 face value.[18]

Using relative cheapness to rank offers is a unique feature of QE auctions, compared to all other auctions of Treasury securities (see Han, Longstaff, and Merrill (2007), Nyborg and Sundaresan (1996), and Hortacsu and Kastl (2012), among others). Therefore, it is natural to check how much extra explanatory power is provided by Cheapness in explaining variations of QE auctions costs. Accordingly, we report the result on a panel regression of $\text{Markup}^c_{t,j}$ on all empirical measures except Cheapness in Column (7). Comparing the results in Columns (6) and (7), we see that adding Cheapness into regressions brings an additional 2.7% into the R^2, which is about 20% (= $(0.155 - 0.128)/0.128$) improvement relative to the explanatory power of all other empirical measures and CUSIP fixed effects.

The improvement in explanatory power by Cheapness is even larger if we drop the CUSIP fixed effects. Specifically, columns (8)–(9) report results of the same panel regressions without any fixed effect. From column (8), we find that Cheapness, VOL, and Specialness are all highly significant in explaining variations of auction costs, similar to the baseline result in column (6). Comparing results in columns (8) and (9), adding Cheapness into regressions brings an additional 2.8% of the R^2, which is about 82% (= $(0.062 - 0.034)/0.034$) improvement relative to the explanatory power of all other empirical measures. Therefore, the bond cheapness, corresponding to a unique feature of QE auctions, is a major determinant of the

[17] If we first aggregate Cheapness into a single time series by auction, then its standard deviation is about 2 cents per \$100 par value. In other words, the dispersion of Cheapness is much higher in the cross section than in the time series.

[18] We also run OLS regressions of the average cost of each auction on weighted average versions of the explanatory variables. Almost none of the variables are significant, suggesting that even if dealers have any propensity to charge a higher spread on any type of bonds, the Fed substitutes into cheaper offers on other bonds according to the ranking by its internal spline-based prices, leading to comparable costs across bonds.

Table 6: Panel Regression of Auction Cost across CUSIP/Auction

Explanatory Variables	(1) Markup^{c}	(2) Markup^{c}	(3) Markup^{c}	(4) Markup^{c}	(5) Markup^{c}	(6) Markup^{c}	(7) Markup^{c}	(8) Markup^{c}	(9) Markup^{c}	(10) $\text{Markup}^{3:30}$
Cheapness	0.075**					0.095**		0.018**		0.165**
	(5.035)					(6.308)		(6.636)		(5.889)
VOL		17.530**				22.042**	18.913**	16.612**	7.696**	35.511**
		(5.226)				(6.910)	(5.735)	(6.739)	(4.759)	(4.563)
Specialness			58.607**			60.561**	75.112**	53.539**	52.384**	138.660**
			(3.365)			(3.744)	(4.279)	(5.761)	(5.945)	(3.264)
OB				-0.115*		-0.067	-0.077	0.002	-0.113+	-0.440*
				(-2.282)		(-1.246)	(-1.288)	(0.048)	(-1.778)	(-2.358)
Bid-Ask					-0.347	-0.761	-0.104	-0.126	-0.972**	-0.516
					(-0.615)	(-1.362)	(-0.184)	(-0.709)	(-4.791)	(-0.385)
Size^{E}						2.706	2.520	2.683+	4.104*	4.507
						(1.466)	(1.320)	(1.884)	(2.046)	(1.242)
N	1,776	1,776	1,776	1,776	1,776	1,776	1,776	1,776	1,776	1,776
R^2	0.104	0.114	0.092	0.086	0.086	0.155	0.128	0.062	0.034	0.120
CUSIP FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	Yes
Auction FE	No	No	No	No	No	No	No	No	No	No

Note: This table reports results of panel regressions of the auction price markup, for each purchased bond in each of the 139 QE auctions of nominal Treasury securities executed from November 12, 2010 to September 9, 2011. The auction price markup $\text{Markup}^{c}_{t,j}$ is computed as the average, weighted by the amount of each accepted offer, of the differences between the offer price and the corresponding Contemporaneous Ask (the 11:30am secondary market ask price for the auctions closed at 11:00am, 11:30am, 12:05pm, and the 2:15pm ask price for the auctions closed at 2:00pm) of the bond for that offer. $\text{Markup}^{3:30}_{t,j}$ is measured relative to the 3:30pm ask price. For each bond in each auction, Cheapness is the difference between the model price implied from the fitted zero-coupon yield curve based on the Svensson model and the actual market price (in cents per $100 dollar face value), VOL is the standard deviation of daily returns of the bond during the five trading days prior to the auction date (in percentage points), Specialness is the difference between the general repo rate and special repo rate on the bond (in percentage points), OB is the total outstanding par value of the bond normalized by the total purchase amount of the auction, and Bid-Ask is the difference between the secondary market ask and bid quotes of a bond normalized by the mid-quote (in basis points). Expected Auction Size is computed as the mean of the announced minimum and maximum of the total auction amount to be purchased divided by the number of included CUSIPs in $billions of par value. Robust t-statistics that correct for serial correlation in the residuals clustered at the CUSIP level are reported in parentheses. Significance levels: ** for $p < 0.01$, * for $p < 0.05$, and $^{+}$ for $p < 0.1$, where p is the p-value.

auction cost variation.

Finally, we run panel regression (9) using the auction price markup measure based on the 3:30pm ask price, as reported in column (10). As in the baseline result of column (6), Cheapness, VOL, and Specialness are all highly significant in explaining variations of the QE auction costs. Moreover, the economic magnitude of Cheapness, VOL and specialness in column (10) are roughly twice as large as those in column (6). Since $\text{Markup}^{c}_{t,j}$ and $\text{Markup}^{3:30}_{t,j}$ differ by the price movement between the auction and the end of day, this last result suggests that bonds that are cheaper relative to model, more volatile or more special experience a greater degree of price decline between the auctions and the end of day.

6 Dealer Profitability

The cost to the Federal Reserve is the profit of primary dealers. In this section, we study the granular heterogeneity of profitability across primary dealers in participating in QE auctions, taking advantage of the availability of dealer identities for each accepted offer in our proprietary dataset.

For each dealer d, we compute his aggregate profit as

$$\text{Aggregate Profit}_d = \sum_{t,j,o} \left(p_{t,j,d,o} - P_{i,j}\right) \cdot q_{t,j,d,o}, \tag{10}$$

and profit margin as

$$\text{Margin}_d = \frac{\sum_{t,j,o} \left(p_{t,j,d,o} - P_{i,j}\right) \cdot q_{t,j,d,o}}{\sum_{t,j,o} q_{t,j,d,o}}, \tag{11}$$

where we use the Contemporaneous Ask for the secondary market price $P_{t,j}$ for CUSIP j and auction t.

Table 7 summarizes the dealer-by-dealer profitability across the 139 QE auctions in our sample period. Column (1) lists the rankings according to dealers' aggregate profits, while the dealer identities are provided in column (2). Columns (3), (4), and (5) provide the aggregate profit in \$ millions, the profit margin in cents per \$100 par value, and the aggregate offer amount in \$ billions of par value sold to the Fed, respectively. In addition, we also report the total number of auctions in column (6), the number of winning offers per auction in column (7), and the total number of CUSIPs sold in column (8).

From columns (3)–(5), we observe a striking concentration of aggregate profits and aggregate amounts of accepted offers among dealers. The top five dealers, Morgan Stanley,

Table 7: Dealer Profitability

(1) Rank by Aggregate Profit	(2) Dealer Identity	(3) Aggregate Profit ($ million)	(4) Profit Margin (Cents per $ 100)	(5) Aggregate offer amount ($ billion)	(6) Total Number of Auctions	(7) Number of Winning Offers per Auction	(8) Total Number of Winning CUSIPs
1	Morgan Stanley	22.67	1.80	125.88	128	6	129
2	Goldman Sachs	19.30	1.34	143.61	138	12	147
3	BNP Paribasl	7.61	1.57	48.49	131	32	174
4	J.P. Morgan	6.45	1.96	32.91	131	5	137
5	RBC	2.97	1.66	17.92	103	4	127
6	Merrill Lynch	2.26	1.35	16.78	114	5	130
7	Daiwa	1.77	2.04	8.71	89	3	81
8	HSBC	1.67	0.74	22.55	117	3	114
9	Jefferies	1.22	0.94	12.92	116	3	108
10	Nomura	0.68	0.36	19.08	116	3	117
11	Deutsche Bank	0.57	0.24	24.24	113	5	125
12	Cantor Fitzgerald	0.36	0.62	5.80	101	3	96
13	Mizuho	-0.19	-0.41	4.74	90	3	80
14	SG Americas	-0.21	-0.56	3.77	51	2	45
15	MF Global	-0.34	-1.53	2.21	40	2	44
16	RBS	-0.63	-0.11	59.27	120	5	120
17	Credit Suisse	-0.72	-0.09	82.97	132	6	140
18	UBS	-0.8	-0.31	25.61	106	4	117
19	Barclays Capita	-2.77	-0.4	69.57	131	7	144
20	Citigroup	-6.50	-1.31	49.61	124	5	135
Average		2.77	0.50	38.83	110	6	116

Note: This table summarizes the dealer profitability across the 139 QE auctions of nominal Treasury securities from November 12, 2010 to September 9, 2011. For each dealer, we compute the profit margin as the average (weighted by the amount of each accepted offer) of the differences between the offer price and the corresponding Contemporaneous Ask (the 11:30am secondary market ask price for the auctions closed at 11:00am, 11:30am, 12:05pm, and the 2:15pm ask price for the auctions closed at 2:00pm) for the bond of that offer, for each dealer across all auctions. The aggregate profit of each dealer is computed as the product between the profit margin and total offer amount the dealer sold to the Fed. We rank the dealers according to their aggregate profits, and report the corresponding rankings, dealer identities, aggregate profits in $ millions, profit margins in cents per $100 par value, aggregate offer amount in $ billions of par value sold to the Fed, total number of auctions, number of winning offers per auction, and total number of CUSIPs sold.

Goldman Sachs, BNP Paribas, J.P. Morgan, and RBC make a combined profit of $59 million, even larger than the $55.37 million total profit of all dealers. These five dealers also accounted for about half the $776.6 billion total purchase amount. Furthermore, the bottom five dealers, RBS, Credit Suisse, UBS, Barclays Capital, and Citigroup accounted for about 40% of the $776.6 billion total purchase amount, comparable to the top five dealers. However, their profit margins are much lower. In fact, the bottom five dealers all have negative profit margins, averaging about −0.44 cents per $100 par value, whereas the top five dealers all have positive profit margins, averaging 1.67 cents per $100 par value.

In sum, through the window of QE auctions, we show that the intermediation profits in the Treasury market is concentrated in only a few dealers. To the best of our knowledge, this is the first formal empirical evidence on the concentrated intermediation activity for U.S. Treasury market, though similar patterns have been documented in other over-the-counter markets such as municipal bond, corporate bond, and asset-backed securities (Li and Schurhoff (2014), Di-Maggio, Kermani, and Song (2015), and Hollifield, Neklyudov, and Spatt (2014)). Furthermore, we document substantial heterogeneity in the profit margin among dealers of similar intermediation activity. Hence, dealers' profit margins seem to be related to economic channels that are beyond the regular intermediation activity.

To investigate what economic channels explain the variations of dealer profit margins, we run the the following panel regression of the profit margin of each dealer in each auction on empirical measures of bond cheapness, value uncertainty, and scarcity/illiquidity:

$$\begin{aligned}\text{Margin}_{t,d} \;=\; & \sum_d \alpha_d D_d + \beta_1 \cdot \text{Cheapness}_{t,d} + \beta_2 \cdot \text{VOL}_{t,d} \\ & + \beta_3 \cdot \text{Specialness}_{t,d} + \beta_4 \cdot \text{OB}_{t,d} + \beta_5 \cdot \text{Bid-Ask}_{t,d} \\ & + \beta_5 \cdot \text{Size}^E_t + \epsilon_{it}, \qquad (12)\end{aligned}$$

where the profit margin $\text{Margin}_{t,d}$ of dealer d in auction t is computed as

$$\text{Margin}_{t,d} = \frac{\sum_{j,o} \left(p_{t,j,d,o} - P_{i,j}\right) \cdot q_{t,j,d,o}}{\sum_{j,o} q_{t,j,d,o}}, \qquad (13)$$

and the Contemporaneous Ask for the secondary market price $P_{t,j}$ for CUSIP j of auction t. Moreover, the empirical measures $\text{Cheapness}_{t,d}$, $\text{VOL}_{t,d}$, $\text{Specialness}_{t,d}$, $\text{OB}_{t,d}$, and $\text{Bid-Ask}_{t,d}$ are calculated as the average of the respective measures of each purchased CUSIP for each winning dealer d in each auction t, weighted by this dealer's purchase amount of this CUSIP in auction t. Therefore, $\text{Cheapness}_{t,d}$ measures the average cheapness of the bonds dealer d

sold to the Fed in auction t, and similarly for other measures. We include the dealer fixed effect D_d to control for potential unobservable effects generic to dealers, but do not include auction fixed effects because our main objective is to study the variations of costs across auctions, similar to (9). We also control for the potential auction size effect by including the expected auction size Size_t^E as a control variable.

From columns (1)–(6) of Table 8, both univariate and multivariate regressions show that the dealer profit margin $\text{Margin}_{t,d}^c$ depends positively on Cheapness, VOL, and Specialness, all statistically significant. That is, for any fixed dealer, the dealer makes a higher profit if he manages to sell to the Fed bonds that are cheaper relative to model, are more volatile, or are more special. In particular, from column (6), a two dollar per \$100 par value increase in Cheapness, which is roughly one standard deviation of Cheapness across auction and dealer in our sample, increases the dealer profit margin by about 1.4 cents ($= 200 \times 0.007$) per \$100 face value. A 0.24 percentage point increase in VOL, which is roughly one standard deviation of the pre-auction volatility across auction and dealer in our sample, increases the dealer profit margin by about 2.1 cents ($= 0.24 \times 8.74$) per \$100 face value. A 0.02 percentage point increase in specialness, which is roughly one standard deviation of the specialness across auction and dealer in our sample, increases the dealer profit margin by about 0.91 cents ($= 0.02 \times 45.60$) per \$100 face value. Therefore, the economic effects of VOL and Specialness are quite large.

The significant explanatory power of Cheapness, VOL, and Specialness is further confirmed in panel regressions without dealer fixed effects, reported in column (8), and in panel regressions measuring the profit margin relative to the 3:30pm ask price in the secondary Treasury market, reported in column (10). The coefficients in column (8) are almost identical to those in column (6), suggesting that the difference in dealer profitabilities is entirely explained by the types of bonds they sell to the Fed.

7 Conclusion

The large scale purchase of Treasury securities by the Federal Reserve, commonly known as quantitative easing (QE), is one of the most significant events in the U.S. Treasury market. The Fed uses QE auctions—multi-object, multi-unit, and discriminatory-price auctions—to implement its QE policy. This paper provides the first empirical analysis of the outcomes of QE auctions. Our proprietary dataset contains all accepted offers with dealer identities of QE auctions executed from November 2010 to September 2011, with a total purchase

Table 8: Panel Regression of Dealer Profitability across Dealer/Auction

Explanatory Variables	(1) $Margin^c_{t,d}$	(2) $Margin^c_{t,d}$	(3) $Margin^c_{t,d}$	(4) $Margin^c_{t,d}$	(5) $Margin^c_{t,d}$	(6) $Margin^c_{t,d}$	(7) $Margin^c_{t,d}$	(8) $Margin^c_{t,d}$	(9) $Margin^c_{t,d}$	(10) $Margin^{3:30}_{t,d}$
Cheapness	0.002**					0.007**		0.007**		0.005*
	(3.543)					(3.299)		(3.289)		(2.464)
VOL		1.476^+				8.739**	5.662**	8.301**	5.205**	1.991
		(1.786)				(4.879)	(4.167)	(4.764)	(4.058)	(1.346)
Specialness			42.046**			45.596**	42.234**	47.279**	43.811**	29.013^+
			(4.174)			(3.262)	(3.105)	(3.548)	(3.391)	(1.940)
OB				-0.113**		-0.086*	-0.134**	-0.083*	-0.131**	-0.226**
				(-3.175)		(-2.385)	(-3.630)	(-2.465)	(-3.920)	(-5.380)
Bid-Ask					-4.476**	-5.102	-9.469**	-3.939	-8.263**	3.115
					(-2.907)	(-1.638)	(-3.600)	(-1.327)	(-3.344)	(0.958)
$Size^E$						0.183	0.268	0.153	0.245	0.387
						(0.744)	(1.074)	(0.652)	(1.024)	(1.285)
N	1,720	1,720	1,720	1,720	1,720	1,720	1,720	1,720	1,720	1,720
R^2	0.006	0.005	0.011	0.006	0.007	0.028	0.023	0.022	0.018	0.020
Dealer FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	Yes
Auction FE	No	No	No	No	No	No	No	No	No	No

Note: This table reports results of panel regressions of the dealer profit margin, for each winning dealer in each of the 139 QE auctions of nominal Treasury securities executed from November 12, 2010 to September 9, 2011. The dealer profit margin $Margin^c_{t,d}$ is computed as the average of the differences between the offer price and the corresponding Contemporaneous Ask (the 11:30am secondary market ask price for the auctions closed at 11:00am, 11:30am, 12:05pm, and the 2:15pm ask price for the auctions closed at 2:00pm) of all purchased CUSIPs for each winning dealer d in each auction t, weighted by this dealer's purchase amount on the corresponding CUSIP in auction t, while $Markup^{3:30}_{t,j}$ is measured relative to the 3:30pm ask price. The explanatory variables, $Cheapness_{t,d}$, $VOL_{t,d}$, $Specialness_{t,d}$, $OB_{t,d}$, and $Bid\text{-}Ask_{t,d}$ are calculated similarly. For each bond in each auction, Cheapness is the difference between the model price implied from the fitted zero-coupon yield curve based on the Svensson model and the actual market price (in cents per $100 dollar par value), VOL is the standard deviation of daily returns of the bond during the five trading days prior to the auction date (in percentage points), Specialness is the difference between the general repo rate and special repo rate on the bond (in percentage points), OB is the total outstanding par value of the bond normalized by the total purchase amount of the auction, and Bid-Ask is the difference between the secondary market ask and bid quotes of a bond normalized by the mid-quote (in basis points). Expected Auction Size is computed as the mean of the announced minimum and maximum of the total auction amount to be purchased divided by the number of included CUSIPs in $billions of par value. Robust t-statistics that correct for serial correlation in the residuals clustered at the dealer level are reported in parentheses. Significance levels: ** for $p < 0.01$, * for $p < 0.05$, and $^+$ for $p < 0.1$, where p is the p-value.

amount of $780 billion Treasury securities.

We find that, on average, the Fed pays 0.71–2.73 cents per $100 par value above the secondary market ask prices. We empirically identify there economic channels that affect the Fed's cost in QE auctions. The first is bond cheapness, the extent to which the market price of a bond is below a yield-curve-model-implied bond value. The second is bond value uncertainty, measured by bond volatility. The third is bond scarcity or illiquidity, measured by the specialness of a bond. The Fed's cost loads positively on each measure. In the data, a one standard deviation increase in cheapness, volatility and specialness increases the Fed's cost by 19, 5, and 1.8 cents per $100 par value, respectively.

On the flip side, we find that the top five dealers account for most of the aggregate profits. Panel regressions suggest that these dealers make higher profits by selling to the Fed bonds that look cheaper (according to a yield-curve model), are more volatile, or are more special.

References

Ashenfelter, O. 1989. How auctions work for wine and art. *Journal of Economic Perspectives* 3: 23–36.

Ausubel, L., Cramton, P., Pycia, M., Rostek, M., and Weretka, M. 2013. Demand Reduction and Inefficiency in Multi-Unit Auctions. Working paper. .

Back, K., and Zender, J. 1993. Auctions of divisible goods: On the rationale for the Treasury experiment. *Review of Financial Studies* 6: 733–764.

Bauer, M., and Rudebusch, G. 2013. The Signaling Channel for Federal Reserve Bond Purchases. *International Journal of Central Banking* forthcoming.

Bikhchandani, S., and Huang, C. 1989. Auctions with Resale Markets: A Model of Treasury Bill Auctions. *Review of Financial Studies* 2: 311–340.

Bikhchandani, S., and Huang, C. 1993. The economics of Treasury securities markets. *Journal of Economic Perspectives* 7: 117–134.

Cammack, E. 1991. Evidence on bidding strategies and the information in Treasury bill auctions. *Journal of Political Economy* 99: 100–130.

Chatterjee, A., and Jarrow, R. 1998. Market manipulation, price bubbles, and a model of the US Treasury securities auction market. *Journal of Financial and Quantitative Analysis* 33: 255–290.

Christensen, J., and Rudebusch, G. 2012. The Response of Interest Rates to US and UK Quantitative Easing. *Economic Journal* 122: 385–414.

D'Amico, S., English, W., Lopez-Salido, D., and Nelson, E. 2012. The Federal Reserve's Large-Scale Asset Purchase Programs: Rationale and Efects. *Economic Journal* 122: 415–226.

D'Amico, S., and King, T. 2013. Flow and Stock Effects of Large-Scale Treasury Purchases: Evidence on the Importance of Local Supply. *Journal of Financial Economics* 108: 425–448.

Di-Maggio, M., Kermani, A., and Song, Z. 2015. The Value of Trading Relationships in Turbulent Times. *working paper*.

Duarte, J., Longstaff, F., and Yu, F. 2007. Risk and Return in Fixed Income Arbitrage: Nickels in Front of a Steamroller?. *Review of Financial Studies* pp. 769–811.

Duffie, D. 1996. Special repo rates. *Journal of Finance* 51: 493–526.

Engle, R., Fleming, M., Ghysels, E., and Nguyen, G. 2012. Liquidity, Volatility, and Flights to Safety in the U.S. Treasury Market: Evidence from a New Class of Dynamic Order Book Models. *Federal Reserve Bank of New York Staff Reports* no. 590.

Fleming, M., and Mizrach, B. 2009. The Microstructure of a U.S. Treasury ECN: The BrokderTech Platform. *Federal Reserve Bank of New York Staff Reports* no. 381.

Gagnon, J., Raskin, M., Remanche, J., and Sack, B. 2011. The Financial Market Effects of the Federal Reserve's Large-Scale Asset Purchases. *International Journal of Central Banking* 7-1: 3–44.

Goldreich, D. 2007. Underpricing in discriminatory and uniform-price Treasury auctions. *Journal of Financial and Quantitative Analysis* 42 (2): 443–466.

Gordy, M. 1999. Hedging Winner's Curse with Multiple Bids: Evidence from the Portuguese Treasury Bill Auction. *Review of Economics and Statistics* 81: 448–465.

Goswami, G., Noe, T., and Rebello, M. 1996. Collusion in Uniform-Price Auctions: Experimental Evidence and Implications for Treasury Auctions. *Review of Financial Studies* 9: 757–785.

Gurkaynak, R., Sack, B., and Wright, J. 2007. The U.S. Treasury Yield Curve: 1961 to the Present. *Journal of Monetary Economics* 54(8): 2291–2304.

Hamilton, J. D. D., and Wu, J. C. 2012. The Effectiveness of Alternative Monetary Policy Tools in a Zero Lower Bound Environment. *Journal of Money, Credit and Banking* 44: 3–46.

Han, B., Longstaff, F., and Merrill, C. 2007. The U.S. Treasury Buyback Auctions: The Cost of Retiring Illiquid Bonds. *Journal of Finance* 62: 2673–2693.

Hancock, D., and Passmore, W. 2011. Did the Federal Reserve's MBS Purchase Program Lower Mortgage Rates?. *Journal of Monetary Economics* 58-5: 498–514.

Hasbrouck, J. 2007. Empirical Market Microstructure. *Oxford University Press, New York.*

Hollifield, B., Neklyudov, A., and Spatt, C. 2014. Bid-ask spreads, trading networks, and the pricing of securitizations: 144a vs. registered securitizations. *Working paper.*

Hortacsu, A., and Kastl, J. 2012. Informational Advantage: A Study of Canadian Treasury Auctions. *Econometrica* 80 (6): 2511–2542.

Hortacsu, A., and McAdams, D. 2010. Mechanism Choice and Strategic Bidding in Divisible Good Auctions: An Empirical Analysis of the Turkish Treasury Auction Market. *Journal of Political Economy* 118 (5): 833–865.

Hu, X., Pan, J., and Wang, J. 2013. Noise as Information for Illiquidity. *Journal of Finance* 68: 2223–2772.

Jordan, B., and Jordan, S. 1997. Special repo rates: An empirical analysis. *Journal of Finance* 52: 2051–2072.

Kastl, J. 2011. Discrete Bids and Empirical Inference in Divisible Good Auctions. *Review of Economic Studies* 78: 978–1014.

Keloharju, M., Nyborg, K., and Rydqvist, K. 2005. Strategic behavior and underpricing in uniform price auctions: Evidence from Finnish Treasury auctions. *Journal of Finance* 60: 1865–1902.

Kremer, I., and Nyborg, K. 2004. Underpricing and Market Power in Uniform-Price Auctions. *Review of Financial Studies* 17: 849–877.

Krishnamurthy, A. 2002. The bond/old-bond spread. *Journal of Financial Economics* 66: 463–506.

Krishnamurthy, A., and Vissing-Jorgensen, A. 2011. The Effects of Quantitative Easing on Interest Rates. *Brookings Papers on Economic Activity* (Fall).

Li, C., and Wei, M. 2013. Term Structure Modeling with Supply Factors and the Federal Reserve's Large-Scale Asset Purchase Programs. *International Journal of Central Banking* 9(1): 3–39.

Li, D., and Schurhoff, N. 2014. Dealer Networks. *Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series* 2014-95.

Meaning, J., and Zhu, F. 2011. The Impact of Recent Central Bank Asset Purchase Programmes. *BIS Quarterly Review* December: 73–83.

Milgrom, P., and Weber, R. 1982. A theory of auctions and competitive bidding. *Econometrica* 50: 1089–1122.

Nelson, C. R., and Siegel, A. F. 1987. Parsimonious Modeling of Yield Curves. *Journal of Business* 60: 473–490.

Nyborg, K., Rydqvist, K., and Sundaresan, S. 2002. Bidder behavior in multiunit auctions-Evidence from Swedish Treasury auctions. *Journal of Political Economy* 110: 394–424.

Nyborg, K., and Sundaresan, S. 1996. Discriminatory versus uniform Treasury auctions: Evidence from when-issued transactions. *Journal of Financial Economics* 42: 63–104.

Pasquariello, P., Roush, J., and Vega, C. 2014. Government Intervention and Strategic Trading in the U.S. Treasury Market. Working paper.

Potter, S. 2013. The Implementation of Current Asset Purchases. *Remarks at the Annual Meeting with Primary Dealers (New York City).*

Sack, B. 2011. Implementing the Federal Reserve's Asset Purchase Program. *Remarks at Global Interdependence Center Central Banking Series Event, Federal Reserve Bank of Philadelphia*.

Simon, D. 1994. Markups, quantity risk, and bidding strategies at Treasury coupon auctions. *Journal of Financial Economics* 35: 43–62.

Stroebel, J., and Taylor, J. 2012. Estimated Impact of the Federal Reserve's Mortgage-Backed Securities Purchase Program. *International Journal of Central Banking* 8-2.

Svensson, L. 1994. Estimating and Interpreting Forward Rates: Sweden 1992-1994. *NBER Working Paper* 4871.

Swanson, E. T. 2011. Let's Twist Again: A High-Frequency Event-Study Analysis of Operation Twist and Its Implications for QE2. Working paper Federal Reserve Bank of San Francisco.

Umlauf, S. 1993. An empirical study of the Mexican Treasury bill auction. *Journal of Financial Economics* 33: 313–340.

Vayanos, D., and Weill, P.-O. 2008. A search-based theory of the on-the-run phenomenon. *Journal of Finance* 63: 1361–1398.

Wilson, R. 1968. The Theory of Syndicates. *Econometrica* 36: 119–132.

Wright, J. H. 2012. What does Monetary Policy do to Long-Term Interest Rates at the Zero Lower Bound?. *The Economic Journal* 122 (November): 447–466.