Journal of Financial Economics 128 (2018) 103–124

Contents lists available at ScienceDirect

Journal of Financial Economics

journal homepage: www.elsevier.com/locate/jfec

Quantitative easing auctions of Treasury bonds[☆]

Zhaogang Song[a], Haoxiang Zhu[b,c,*]

[a] Johns Hopkins University Carey Business School, 100 International Drive, Baltimore, MD 21202, USA
[b] MIT Sloan School of Management, 100 Main Street E62-623, Cambridge, MA 02142, USA
[c] National Bureau of Economic Research, 1050 Massachusetts Avenue, Cambridge, MA 02138, USA

ARTICLE INFO

Article history:
Received 4 August 2016
Revised 6 March 2017
Accepted 7 March 2017
Available online 16 February 2018

JEL classification:
G01
G12
G14
G18

Keywords:
Quantitative easing
Auction
Treasury bond
Federal Reserve

ABSTRACT

The Federal Reserve uses (reverse) auctions to implement its purchases of Treasury bonds in quantitative easing (QE). To evaluate dealers' offers across multiple bonds, the Fed relies on its internal yield curve model, fitted to secondary market bond prices. From November 2010 to September 2011, a one standard deviation increase in the cheapness of a Treasury bond (how much the market price of the bond is below a model-implied value) increases the Fed's purchase quantity of that bond by 276 million and increases the auction costs on that bond by 2.6 cents per $100 par value, controlling for standard covariates. Our results suggest that the Fed harvests gains from trades by purchasing undervalued bonds, but strategic dealers extract some profits because the Fed's relative values can be partly inferred from price data.

1. Introduction

One of the most significant events in the history of the US Treasury market is the Federal Reserve's large-scale asset purchase programs of long-term Treasury securities since the 2008–2009 financial crisis, commonly known as quantitative easing (QE).[1] Up to September 2011, the end of the sample period in our study, the Fed purchased $1.19 trillion of Treasury bonds, equivalent to about 28% of the total outstanding stock of these securities in March 2009, when these purchase operations began.[2] Although an extensive academic literature has studied the effect of QE on

☆ For helpful discussions and comments, we thank Bill Schwert (editor), Francis Longstaff (referee), Tobias Adrian, Jeremy Bulow, Hui Chen, Jim Clouse, Stefania D'Amico, 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, Loriana Pelizzon, 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 National Bureau of Economic Research (NBER) Summer Institute Asset Pricing meeting, the Third International Conference on Sovereign Bond Markets at New York University Leonard N. Stern School of Business, Cheng Kong Graduate School of Business, Tsinghua University PBC School of Finance, Massachusetts Institute of Technology, the 2016 European Central Bank and International Monetary Fund (ECB-IMF) Workshop on Money Markets, Monetary Policy Implementation and Market Infrastructures, and the 2016 Federal Reserve Bank of Atlanta Workshop on the Impact of Extraordinary Monetary Policy on the Financial Sector. The first version of the paper was written when Zhaogang Song was at the Board of Governors of the Federal Reserve System.

* Corresponding author at: MIT Sloan School of Management, 100 Main Street E62-623, Cambridge, MA 02142, USA.
E-mail addresses: zsong8@jhu.edu (Z. Song), zhuh@mit.edu (H. Zhu).

https://doi.org/10.1016/j.jfineco.2018.02.004
0304-405X/

[1] The large-scale asset purchase programs began with the purchasing of agency mortgage-backed securities (MBSs) and agency debt announced in November 2008. Given our focus and for brevity, we use "QE" to refer to purchases 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."

interest rates following the financial crisis, very few papers focus on the implementation mechanism of QE.[3]

Studying the implementation of QE generates insights into the behavior of dealers as intermediaries for monetary policy. In practice, the Fed purchases Treasury bonds not directly from end investors such as asset managers and insurance companies, but indirectly from about 20 primary dealers, who essentially act as an oligopoly of intermediaries between the Fed and the market. Hence, whether the Fed can achieve competitive pricing and avoid paying excessive execution costs (Potter, 2013) under the primary dealer system is not clear.[4]

The implementation of QE also provides new insights into the Fed's preference as a large player in financial markets. Among the large set of outstanding Treasury securities, the Fed has the discretion in choosing which ones to purchase. These securities are substitutes in principle, but they differ substantially in pricing and liquidity in the secondary market. Thus, the Fed's algorithm of selecting bonds, potentially conditional on market prices and other characteristics, reflects the Fed's preference in implementing monetary policy.

In this paper, we empirically study the auction mechanism of the Fed's purchase of Treasury securities during QE. We pay particular attention to the interaction between the Fed's preference and dealers' strategic behavior. According to public information, the Fed uses an internal (undisclosed) yield-fitting model to evaluate dealers' offers on different bonds. Under such an algorithm, the Fed would prefer bonds that appear to be undervalued (relative to model) and therefore harvest gains from trade. But because dealers are strategic, they could be able to extract high profits on those undervalued bonds. Consistent with these predictions, we find that bond cheapness, i.e., how much the market price of a Treasury security falls below a model-implied value, positively predicts the Fed's purchases quantities and costs in the time series and cross section, controlling for other covariates. We further find remarkable concentration of the dealers' profits among the top few. Evidence suggests that the top dealers are more responsive in delivering larger quantities of bonds that the Fed appears to prefer, and that the non-top dealers charge moderately higher prices.

[3] See Krishnamurthy and Vissing-Jorgensen (2011), Gagnon et al. (2011), Hancock and Passmore (2011), Swanson (2011), Wright (2012), D'Amico et al. (2012), Hamilton and Wu (2012), Christensen and Rudebusch (2012), Stroebel and Taylor (2012), Bauer and Rudebusch (2014), Li and Wei (2013), D'Amico and King (2013), Meaning and Zhu (2011), Eser and Schwaab (2016), and Schlepper et al. (2017), among others.

[4] In the context of Treasury issuance auctions, which also use the primary dealer system, Hortacsu et al. (2017) estimate that dealers make a surplus of about 3 basis points of the auction size and the total allocative inefficiency is about 2 basis points. A recent class-action lawsuit reinforces the concern of noncompetitive behaviors in Treasury auctions (see http://www.labaton.com/en/cases/upload/Treasuries-State-Boston-Complaint.PDF). Boyarchenko et al. (2015) argue that having a small number of dealers reduces information frictions and, hence, the issuance cost of the Treasury.

1.1. QE auctions mechanism and hypotheses

The purchases of Treasury securities in QE are conducted through a series of multi-object, multiunit, and discriminatory price (reverse) auctions, implemented by the Federal Reserve Bank of New York. We refer to these auctions as QE auctions. Before each purchase auction, the Fed announces a range of total amount and a maturity bucket of the Treasury bonds to be purchased, but it specifies neither the exact total amount nor the amount for individual bonds (CUSIPs). The primary dealers can submit multiple offers on any eligible bond. Holding a single auction for multiple bonds is faster than holding separate auctions for each bond.

To evaluate dealers' offers across different Treasury securities, the Fed needs an algorithm. The Federal Reserve Bank of New York states: "Offers will be evaluated based on their proximity to prevailing market prices at the close of the auction as well as on measures of relative value. Relative value measures are calculated using the Federal Reserve Bank of New York's proprietary model."[5] Sack (2011, footnote 6) further explains that "the methodology for comparing the relative value of the securities at the offered prices is based on a spline fitted through the prices of Treasury securities." Therefore, after adjusting for its internal benchmark prices, the Fed treats different Treasury securities as perfect substitutes.

The Fed's relative value based algorithm leads to an interesting trade-off. On the one hand, this algorithm helps the Fed identify undervalued securities for which the gains from trade between the Fed and the market are particularly large. In this sense, buying these securities is more efficient from the allocative point of view. On the other hand, the Fed's relative valuations of different bonds could be (partly) predictable by market participants, especially the primary dealers, even if the Fed's internal model is confidential. The reasons is twofold. First, fixed-income investors and dealers are familiar with yield curve models as a regular part of their business.[6] Second, primary dealers have years of experience interacting with the Fed through its permanent and temporary open market operations.[7] As standard auction theory predicts, if dealers are strategic and can (partly) predict which bonds the Fed views as cheaper, they can strategically extract higher profits on these bonds, for which gains from trades are larger to start with.

Based on this trade-off, we state our two main testing hypotheses:

Hypothesis 1. Bonds that appear cheaper (than model) are purchased more by the Fed.

[5] See https://www.newyorkfed.org/markets/longertermtreas_faq.html.

[6] See Duarte et al. (2007) for studies of the yield curve arbitrage strategy of fixed-income hedge funds.

[7] Relatedly, the Fed used external investment managers, including BlackRock, Goldman Sachs, PIMCO, and Wellington, to implement its agency MBS purchases in the early QE period, from December 2008 to March 2010. Through this interaction, these investment firms could learn valuable information about the Fed's operations. See https://www.newyorkfed.org/markets/mbs_faq.html for details.

Hypothesis 2. Cheaper bonds incur higher costs (relative to secondary market prices) to the Fed.

1.2. Empirical measure and evidence

Our empirical analysis employs a propriety data set that has the outcomes of the 139 purchase auctions of nominal Treasury securities from November 12, 2010 to September 9, 2011, with a total purchased amount of about $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 mortgage-backed security (MBS) holdings. The distinguishing feature of our study is the use of detailed data of each accepted offer, including the quantity, the price, and the dealer's identity. This allows us to study the granular heterogeneity across auctions, bonds, and dealers. Our secondary market data set is the New Price Quote System (NPQS) data, obtained from the Federal Reserve Bank of New York, that contain bond prices quoted in the secondary market at 8:40 a.m., 11:30 a.m., 2:15 p.m., and 3:30 p.m. every day. Additional data include the repo specialness, amount outstanding, and other characteristics of the bonds in our sample period.

Because our main hypotheses are based on the Fed's relative value method of evaluating offers, we begin by constructing a proxy for cheapness. Our baseline yield curve model is a piecewise cubic spline, used by McCulloch (1971,1975), Fisher et al. (1995), and Hu et al. (2013), among others. We fit the cubic spline model to the 8:40 a.m. prices for each QE auction day and define cheapness to be how much the market price of a bond falls below its model-implied value. (Similar results are obtained if the yield curve is fitted to the Svensson, 1994 model.) Although we do not observe the Fed's confidential yield curve model, we expect our proxy to be correlated with the Fed's. Potential noise in our measure would bias against finding significant results.

Our first hypothesis is that cheaper bonds are purchased more by the Fed. To test it, we run panel regressions of purchase quantity at the auction-CUSIP level on our proxy of cheapness and a set of covariates measured before the auction, including return volatility, specialness, outstanding balance, and bid-ask spread. We control for auction and CUSIP fixed effects and use double-clustered standard errors (Petersen, 2009). Consistent with our hypothesis, purchase quantities are significantly higher for cheaper bonds. Controlling for all else, a one standard deviation increase in a bond's cheapness in any auction [of about 23 basis points (bps)] increases the Fed's purchase of that bond in that auction by about $276 million in par value. This is a large magnitude, as the average purchase size of any bond in any QE auction is about $410 million.

The second hypothesis is that cheaper bonds incur higher costs to the Fed. At the auction-CUSIP level, we measure the costs to the Fed by the difference between the accepted offer prices and the secondary market prices, weighted by the accepted quantity. We hence use the 11:30 a.m. ask price as the closest ask for auctions closing at 11 a.m., 11:30 a.m., and 12:05 p.m., and we use the 2:15 p.m. ask price for auctions closing at 2:00 p.m. Because 134 of the 139 QE auctions close at 11 a.m., our measure of the cost to closest ask can be viewed as the realized costs of the Fed 30 min after the auction or interpreted as the profit of dealers in selling bonds to the Fed at 11 a.m. and then covering their short positions in the secondary markets 30 min later [see Hasbrouck (2007) for a discussion of realized cost]. In addition, we use the secondary market prices at 3:30 p.m. to compute the realized cost at the end of the day. On average, the Fed pays 0.71 cents per $100 par value above secondary market ask prices near the close of the auction and 2.73 cents above the 3:30 p.m. secondary market ask prices. The measured average costs are small to moderate compared with the cost of previous Treasury buyback operations and Treasury issuance auctions. The variations of these costs across time and CUSIPs are large, however, with standard deviations of about 11 cents and 22 cents per $100 par value, respectively.

We run, at the auction-CUSIP level, panel regressions of the Fed's costs (relative to closest ask and 3:30 p.m. ask) on cheapness and four other covariates, controlling for auction and CUSIP fixed effects. Consistent with our second main hypothesis, cheapness is a significant predictor of the auction cost. A one standard deviation increase in cheapness (23 bps) increases the Fed's cost by 2.6 cents per $100 par value. Again, this is a large economic magnitude given the relative small average cost of the Fed.

In sum, the evidence on purchase quantities confirms the Fed's preference for buying undervalued bonds (relative to model) in implementing QE, which achieves more gains from trade. The evidence on the Fed's costs, however, implies that dealers extract more profits on more undervalued bonds, as they can (partly) predict which bonds the Fed views as cheaper.

Finally, we further study the granular heterogeneity of primary dealers, taking advantage of the available dealer identities in our data set. We find remarkable heterogeneity among dealers in their transaction volume and total profits in selling Treasury bonds to the Fed. Among the 20 primary dealers that participated in QE auctions, the top five handle about half of the total purchase amount ($780 billion) and extract all of the aggregate profits (relative to closest ask). The other dealers collectively incur a small loss.

To investigate the possible causes of this heterogeneity, we first compute, for each auction and each bond, each dealer's sales quantity and quantity-weighted price markup. Then, we regress the quantity and the price markup, at the auction-CUSIP-dealer level, on cheapness and other covariates, including the indicator function $1(\text{top five dealer})$ and a full set of interaction terms. These interaction terms would capture the additional sensitivity of the top five dealers' sales quantity and price markup to bond characteristics. The top five dealers' sales quantities are ten times as responsive to cheapness as the other dealers', but the non-top five dealers' price markups are 1.5 times as responsive to cheapness as the top five dealers. This evidence suggests that the top five dealers on average seem to have a large comparative advantage in finding and selling bonds that the Fed seems to prefer, whereas the other dealers on average have certain comparative advantage in strategically charging higher prices.

1.3. Related literature

To the best of our knowledge, ours is the first paper that studies the implementation mechanism of QE in the United States. Breedon and Turner (2016) estimate the aggregate cost of the Bank of England's QE auctions of gilts in the time series. Relative to the Bank of England's mechanism that does not involve an internal model, the unique model-based approach of the Fed allows us to provide novel insights into the strategic interaction between the central bank and the primary dealers. Our results also provide new evidence on the heterogeneity of dealers.

In the literature of US Treasury markets, the Treasury's buyback operations from March 2000 to April 2002 seems to be the closest match to QE auctions. Han et al. (2007) conduct a comprehensive analysis of the buyback operations and find that the Treasury incurs an average cost of 4.38 cents per $100 par value. Pasquariello et al. (2014) study how the Fed's permanent open market operations (POMOs) conducted from 2001 to 2007 affect Treasury market liquidity. Their data contain only the aggregate purchase quantities in POMOs, not security-specific quantities or transaction prices.

This paper is also related to the empirical literature on Treasury issuance auctions, including Cammack (1991), Nyborg and Sundaresan (1996), Goldreich (2007), Lou et al. (2013), and Fleming and Liu (2014), among others. Relative to these studies, the Fed's fairly transparent QE auction mechanism allows us to identify cheapness as a new and significant predictor of purchase quantities and costs. Moreover, the dealer-level data enable us to look into the heterogeneity of dealers and link it back to how dealers respond to (proxies of) Fed's relative valuations, a unique feature of QE auctions.

The rest of the paper proceeds as follows. Section 2 provides institutional details of QE auctions. In Section 3, we present a theoretical framework for QE auctions and derive testable predictions. Section 4 describes the data. Sections 5 and 6 present regression results of the Fed's purchase quantities and costs, respectively. Section 7 presents evidence on the heterogeneity of dealers. Section 8 concludes.

2. Institutional background of QE auctions

From November 12, 2010 to September 9, 2011, the Federal Reserve conducted a series of 156 purchase auctions of US Treasury securities, including nominal Treasury securities and Treasury inflation protected securities (TIPS). These auctions cover two Fed programs. The first, commonly referred to as 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 MBSs into longer-term Treasury securities, announced on August 10, 2010, with a total purchase size of $180 billion over our sample period.[8] These programs aim to maintain downward pressure on longer-term interest rates, support mortgage markets, and 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, multiunit, and discriminatory price auctions. Transactions are conducted on the FedTrade platform. Direct participants of QE auctions include only the primary dealers recognized by the Federal Reserve Bank of New York, although other investors can indirectly participate through the primary dealers.[9]

Fig. 1 describes the time line of a typical QE auction: pre-auction announcement, auction execution, and post-auction information release. To initiate the asset purchase operation, the Fed makes a pre-auction announcement on or around the eighth business day of each month. The announcement includes the total amount of purchases expected to take place between the middle of the current month and the middle of the following month.[10] Most important, this announcement contains a schedule of upcoming purchase operations, including operation dates, settlement dates, security types to be purchased (nominal or TIPS), the maturity date range of eligible issues, and an expected range for the size of each operation. Therefore, the announcement identifies 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 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. The minimum offer size and the minimum increment of offer size are both $1 million. Each offer is a price-quantity pair, specifying the par value the dealer is willing to sell to the Fed at a specific

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

[9] In the first half of our sample period until February, 2, 2011, there were 18 primary dealers: 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. After the acquisition of Merrill Lynch in 2009, Bank of America is rebranded as Bank of America Merrill Lynch (BAML). See https://www.newyorkfed.org/markets/primarydealers#primary-dealers for the current and historical lists of primary dealers.

[10] This amount is determined by the planned purchase amount over the coming monthly period, as part of the $600 billion total purchases, the approximate amount of principal payments from agency MBSs 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.

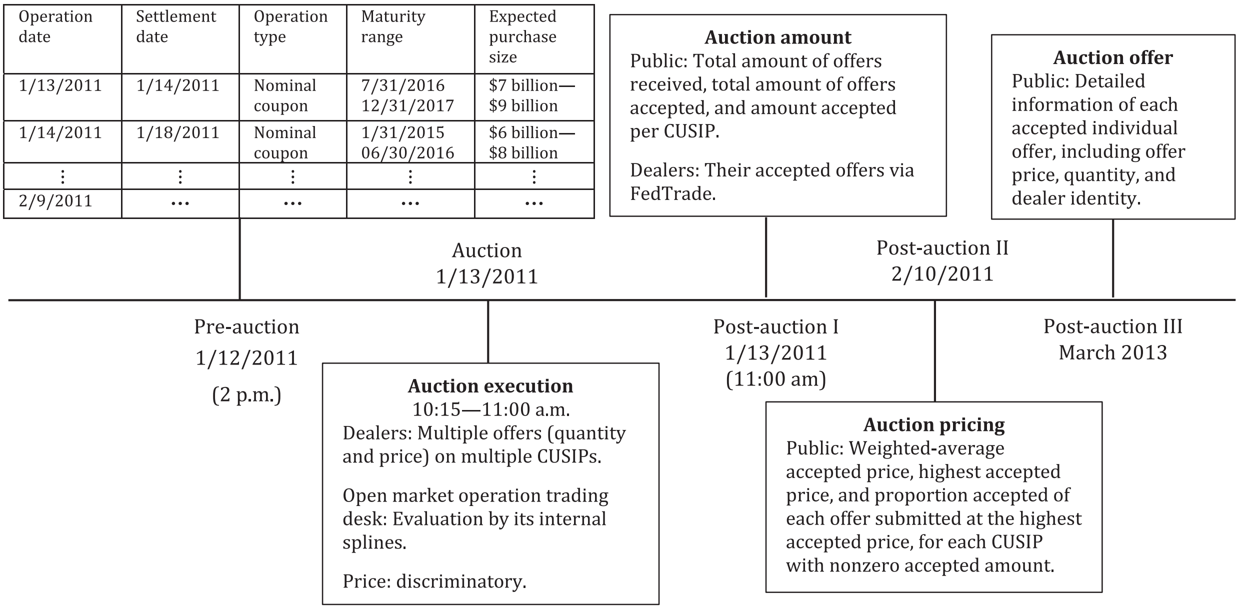

Fig. 1. Example of a quantitative easing (QE) auction time line.

price. The auctions happen mostly between 10:15 a.m. and 11:00 a.m. Eastern Time Zone. Very rarely, the auctions happen between 10:40 a.m. and 11:30 a.m., 11:25 a.m. and 12:05 p.m., and 1:15 p.m. and 2:00 p.m. Out of the 139 auctions we analyze, one is closed at 11:30 a.m., one is closed at 12:05 p.m., and three are closed at 2 p.m. The remaining 134 auctions are all closed at 11:00 a.m.

Within a few minutes after the closing of the auction, the Fed publishes the auction results on the Federal Reserve Bank of New York's website, including the total number of offers received, total number of offers accepted, and the amount purchased per CUSIP.[11] 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, for each security purchased in each auction, the weighted average accepted price, the highest accepted price, and the proportion accepted of each offer submitted at the highest accepted price. Finally, in accordance with the Dodd–Frank Wall Street Reform and Consumer Protection Act of 2010, detailed auction results including the offer price, quantity, and dealer identity for each accepted individual offer are released two years after each quarterly auction period.

The unique feature of QE auctions is that each auction involves a set of heterogeneous securities. Therefore, an algorithm is needed to rank offers on different CUSIPs. To make this ranking, the Fed compares each offered price with a benchmark price of the offered bond calculated from its internal spline-based yield curve model, fitted to the secondary market prices of Treasury securities (Sack, 2011). Thus, after adjusting for these benchmark prices, different CUSIPs become perfect substitutes from the Fed's perspective.

Evaluating dealers' offers on different bonds based on model-implied prices introduces an interesting trade-off. On the one hand, conditional on filling the desired purchase amount, the Fed would naturally prefer bonds that trade at a discount in the secondary market relative to otherwise similar bonds. That is, the Fed behaves like a rational investor in buying undervalued securities based on its internal yield curve model. On the other hand, sophisticated investors and primary dealers could have some information about the Fed's yield curve algorithm, even though the Fed does not publish it. After all, fitting yield curve models is a routine practice by sophisticated fixed-income investors and dealers for evaluating the relative cheapness or richness of different bonds. Moreover, dealers could gain information about the Fed's relative valuations through years of interactions with the Fed. This information may encourage dealers to bid strategically in QE auctions.

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 has not been addressed in existing auction models. 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 illustrate how dealers' information about the Fed's yield curve model can affect the auction outcome, including purchase quantities and costs.[12]

[11] See https://www.newyorkfed.org/markets/pomo_landing.html.

[12] Even for a single-object, multiple-unit auction, multiple Bayesian–Nash equilibria can exist, so that no definitive theoretical predictions can

Suppose N dealers are 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: (1) the common value component that captures the resale value of the bond in the secondary market and (2) 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 (at least 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 trading network 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 (Wilson, 1968; Ausubel et al., 2014). Applied to QE auctions, a more severe winner's curse problem implies that dealers submit higher priced offers to the Fed, leading to a higher expected cost of the Fed. For simplicity, we do not formally reproduce this standard argument here (see Cammack, 1991; Umlauf, 1993; Keloharju et al., 2005; Han et al., 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 component. Because the implication from the common value component is clear, we assume that the dealers have pure values $\{v_i\}$ for the bond, where $\{v_i\}$ are independent and identically distributed (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., in which 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, which can always be done by shifting the support $[\underline{v}, \bar{v}]$ of $\{v_i\}$.

Suppose that 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 with other eligible bonds, based on the Fed's internal benchmark prices, as discussed in Section 2. Consequently, 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 benchmark prices. In this sense, v_0 represents the cheapness of a bond relative to other bonds. We expect dealers to have some information about the Fed's relative valuations. To model the information that dealers have on the Fed's yield curve algorithm, we assume for simplicity that v_0 is common knowledge for simplicity, i.e., dealers have full information on the Fed's preference of bonds. (If dealers observe only a noisy signal of v_0 and, hence, have a probability distribution over v_0, the qualitative implications of the model would not change.)

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 does not buy this bond. Again, this does not mean that v_0 is the reservation price of the Fed in the usual sense. Instead, the interpretation is that if the offer prices on a bond are too high relative to those on other bonds, the Fed buys other bonds instead of this one.

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

We now consider $v < v_0$. Dealer i wins the auction if $a_i < \min_{j \neq i} \beta(v_j)$, 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, the Fed accepts the lowest offer if and only if $\min_i \beta(v_i) \leq v_0$, which happens with probability

$$\begin{aligned} 1 - P(\min_i v_i > \beta^{-1}(v_0)) &= 1 - (1 - P(v_i < \beta^{-1}(v_0)))^N \\ &= 1 - (1 - F(\beta^{-1}(v_0)))^N, \end{aligned} \tag{3}$$

which is increasing in v_0. That is, the Fed is more likely to buy a bond if it values the bond higher.

Moreover, if the Fed accepts the best offer, the Fed's cost is $\min_i \beta(v_i)$. All else equal, $\beta(v)$ is increasing in v_0 for $v \in$[<uline> v </uline> , v_0], for all finite N. This predicts that the auction price and, hence, markup are higher if the Fed's value v_0 of the bond is higher. Intuitively, this is because dealers have market power and behave strategically, which is standard in auction theory. This implies that the auction price is higher if the bond looks cheaper based on the Fed's internal model-implied bond value.

4. Data and measurement

In this section, we describe the data and the empirical measures of various determinants of auction outcomes.

4.1. Auction data

Our sample period is from November 12, 2010 to September 9, 2011, which is the only time period during which the Fed purchased only Treasury securities and for which the detailed data of dealer offers are available.[13]

Table 1 reports the maturity distribution of planned purchases of Treasury debt over our sample period,

be made about the equilibrium bidding strategies and auction outcomes (see, e.g., Bikhchandani and Huang, 1993; Back and Zender, 1993; Ausubel et al., 2014). 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.

[13] The Dodd–Frank Act, enacted on July 21, 2010, mandates that the Federal Reserve should release detailed auction data to the public with a

Table 1
Maturity distribution of planned purchases in quantitative easing (QE) auctions.
This table shows the maturity distribution of planned purchases 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 seven-year note is considered part of the 5.5–7-year sector, and the on-the-run ten-year note is considered part of the 7–10-year sector. TIPS = Treasury Inflation Protected Securities.

	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

announced on November 3, 2010 by the Fed. During this period, the Fed conducted 139 auctions of nominal Treasury securities and 17 auctions of TIPS. Because TIPS account for only 3% of the total purchases in terms of par value, in this paper we focus on the 139 auctions of nominal securities. These 139 auctions were conducted on 136 days, with two auctions on each of November 29, 2010, December 20, 2010, and June 20, 2011 and only one auction on all the other days. Only 6% of planned purchase amounts have a maturity beyond ten years. According to the Fed, this maturity distribution has an average duration between five and six years for the securities purchased. The Fed does not purchase Treasury bills, Separate Trading of Registered Interest and Principal Securities (STRIPS), or securities traded in the when-issued market.[14]

Our primary data set provides detailed outcomes of QE auctions: the expected total purchase size range, the total par value offered, and the total par value accepted for each auction; the indicator of whether a CUSIP was included or excluded in the auctions; 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 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, we are the first to analyze the implementation mechanism of QE in the US using bid-level data. In the context of US Treasury issuance auctions, Hortacsu et al. (2017) is so far the only other study we are aware of that uses bid-level data, although their data are not publicly available.[15] Moreover, QE auctions differ from Treasury issuance auctions in several ways, especially the use of a model to evaluate offers on multiple bonds.

Table 2 presents descriptive statistics on the number of nominal Treasury securities for the 139 QE auctions in our sample period. The number of eligible bonds in an auction varies between 15 and 36, with a mean of 26. Because a small number of bonds are excluded from auctions, the number of eligible (included) bonds is between 13 and 34, averaging 25 per auction.[16] Among these included bonds, in each auction the Fed purchases between three and 27 bonds, with an average of 15 bonds. On average, 11 eligible bonds are not purchased by the Fed in any particular auction. Across all 139 auctions, 186 CUSIPs have ever been purchased by the Fed, among the 215 eligible CUSIPs.

Panel A of Table 3 shows that the par amounts of submitted offers vary between $4 billion and $43 billion, averaging about $21 billion per auction. The offer amount accepted by the Fed varies between $0.7 billion 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, and the accepted offer amount per accepted bond is $0.41 billion.

two-year delay after each quarterly operation period. In consequence, detailed dealer offers are available from July 22, 2010. We discard the period July 22, 2010–November 11, 2010 because no orderly expected purchase sizes at the auction level were announced by the Fed in this period. Moreover, on September 21, 2011, the Fed announced the Maturity Extension Program and changes in the agency debt and agency MBS reinvestment policy, which lead to purchases of agency MBS and sales of short-term Treasury securities in addition to purchases of long-term Treasury securities thereafter. Therefore, to avoid potential compounding effects due to other policy operations, we discard the period starting from September 21, 2011 to focus on a clean period of only Treasury bond purchases. In addition, we discard the period September 10–20, 2011 because the monthly planned operation of this September is interrupted by the policy change on September 21, 2011.

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

[15] Previous empirical studies of issuance auctions and buyback auctions of US Treasury securities, including Cammack (1991), Simon (1994), Nyborg and Sundaresan (1996), and Han et al. (2007), have used data only at the aggregate auction level or at the CUSIP level at best. For other countries, however, studies of government debt auctions have used bid-level data, such as Umlauf (1993), Gordy (1999), Nyborg et al. (2002), Keloharju et al. (2005), Hortacsu and McAdams (2010), Kastl (2011), and Hortacsu and Kastl (2012). Theoretical and experimental studies of Treasury issuance auctions include Bikhchandani and Huang (1989), Chatterjee and Jarrow (1998), Goswami et al. (1996), and Kremer and Nyborg (2004), among others.

[16] According to the Fed's communications to the public, excluded bonds are those trading with heightened specialness in the repo market or the cheapest to deliver into the front-month Treasury futures contracts. These bonds are excluded to avoid creating or exacerbating supply shortages in repo and futures markets. Unreported regressions reveal that the Fed is less likely to include a bond into QE auctions if the bond has a higher specialness, reflecting its tight supply in the repo market. The Fed also excludes a CUSIP if additional purchase of that CUSIP would push the Fed's total holding of it above a fixed per-security size limit. 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.

Table 2
Summary statistics of bonds in quantitative easing (QE) auctions.
This table presents descriptive statistics about the number of bonds covered in the 139 QE auctions of nominal Treasury securities, from November 12, 2010 to September 9, 2011.

Statistic	Eligible bonds	Excluded bonds	Included bonds	Included bonds (not purchased)	Included bonds (purchased)
Mean	26	1	25	11	15
Standard deviation	6	1	7	7	5
Minimum	15	0	13	0	3
25% quartile	20	1	19	5	12
50% quartile	27	1	26	9	15
75% quartile	31	2	29	16	19
Maximum	36	4	34	28	27

Table 3
Summary statistics of auction sizes, offers, and dealers.
This table presents descriptive statistics on the auction size (Panel A) and the number of winning offers and dealers (Panel B) for the 139 quantitative easing auctions of nominal Treasury securities, from November 12, 2010 to September 9, 2011. We report mean, standard deviation, minimum, maximum, and quartiles at the auction level.

Panel A: Auction size

Statistic	Offer amount (billions of dollars)	Purchase amount (billion of dollars)	Offer-to-cover ratio	Announced Minimum size (billions of dollars)	Announced maximum size (billions of dollars)	Offer Amount per included bond (billions of dollars)	Purchase Amount per accepted bond (billions of dollars)
Mean	20.76	5.59	4.58	4.58	6.2	0.84	0.41
Standard deviation	8.81	2.33	2.87	1.98	2.48	0.31	0.29
Minimum	4.13	0.72	1.67	0.50	1	0.12	0.04
25% quartile	13.53	3.16	3.06	2.75	3.5	0.59	0.23
50% quartile	21.04	6.68	3.66	5	7	0.83	0.36
75% quartile	27.70	7.30	4.64	6	8	1.05	0.52
Maximum	42.88	8.87	7	7	9	1.76	2.57

Panel B: Number of winning offers and dealers

Statistics	Number of winning offers	Number of winning dealers	Purchase amount per winning offer (billions of dollars)	Purchase amount per winning dealer (billions of dollars)
Mean	103	16	0.07	0.36
Standard deviation	54	3	0.06	0.17
Mininum	8	4	0.01	0.06
25% quartile	57	15	0.04	0.24
50% quartile	102	17	0.05	0.38
75% quartile	137	18	0.08	0.45
Maximum	326	20	0.49	1.38

Panel B presents summary statistics on the number of winning offers and dealers. The number of winning offers ranges between eight and 326, with a mean of 103, and the number of winning dealers ranges between four 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 afterward.) 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.

4.2. Data on bond prices and bond characteristics

Our secondary market price data contain indicative bid and ask quotes from the New Price Quote System 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). There are four pairs of bid and ask quotes each day at 8:40 a.m., 11:30 a.m., 2:15 p.m., and 3:30 p.m., which are the best bid and ask prices across different trading platforms of Treasury securities. We choose these NPQS quotes because they cover off-the-run securities that are targeted in QE auctions. By contrast, the BrokerTec data used in recent studies such as Fleming and Mizrach (2009) and Engle et al. (2012) mainly contain prices of newly issued on-the-run securities. Moreover, these price quotes are important sources for the Fed's internal yield curve fitting algorithm and benchmark prices. Hence, our bond cheapness measure using NPQS prices as inputs should be reasonably correlated with the Fed's true preference.

We obtain the CUSIP-level special collateral repo rates from the BrokerTec Interdealer Market Data that averages quoted repo rates across different platforms between 7 a.m. and 10 a.m. 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; D'Amico et al., 2013). We also obtain the outstanding par value of Treasury securities each day from the Monthly Statement of the Public Debt (MSPD) of the Treasury Department.

4.3. Empirical measures

To capture the unique feature of QE auctions involving relative valuations of different CUSIPs, we need an empirical measure of bond cheapness. The exact benchmark prices used by the Fed to evaluate dealer offers are confidential and unobservable. Nonetheless, we construct a proxy of bond cheapness by applying a popular spline-based yield-fitting method to the NPQS data that are also used by the Fed (Sack, 2011). Following McCulloch (1971, 1975), Fisher et al. (1995), and Hu et al. (2013), we assume that the instantaneous forward interest rate at maturity m is a cubic spline with knot points on $(\tau_0, \tau_1, \ldots, \tau_k)$:

$$f(m,\theta) = a^i\left(\frac{m-\tau_{i-1}}{\tau_i-\tau_{i-1}}\right)^3 + b^i\left(\frac{m-\tau_{i-1}}{\tau_i-\tau_{i-1}}\right)^2 + c^i\left(\frac{m-\tau_{i-1}}{\tau_i-\tau_{i-1}}\right) + d^i, \quad \text{for } m \in [\tau_{i-1}, \tau_i], \tag{4}$$

where $\theta = \left\{\left(a^i, b^i, c^i, d^i\right), i = 1, 2, \ldots, k\right\}$ is the set of cubic spline parameters. Additional restrictions are imposed on the parameters so that both f and its first derivative are continuous at the connecting knot points over the k subintervals, and the forward rates are positive at maturities of zero and infinity.

For each set of parameters θ, we compute the corresponding zero-coupon yield curve by integrating the forward rates based on Eq. (4), 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 parameter θ by

$$\theta_t = \arg\min_{\theta}\left\{\sum_{j=1}^{N_t}\left[\left(P_{tj}(\theta) - P_{tj}\right)/D_{tj}\right]^2 + \lambda\int_0^{\tau_k}\left[f''(x,\theta)\right]^2 dx\right\}, \tag{5}$$

where P_{tj} is the (clean) market price of bond j on auction day t, $j = 1, \ldots, N_t$; $P_{tj}(\theta)$ is the model-implied (clean) price of bond j based on the spline model in Eq. (4); and D_{tj} is the duration of bond j. The first term $\sum_{j=1}^{N_t}\left[\left(P_{tj}(\theta) - P_{tj}\right)/D_{tj}\right]^2$ is the inverse duration-weighted sum of the squared deviations between the actual market prices of Treasury securities and the spline model-implied prices. The parameter θ is chosen primarily to minimize this deviation.

We need to make several choices in fitting the model. (In the appendix, we conduct robustness checks of the main results by varying these inputs.)

Smoothness parameter λ: The second term in Eq. (5), $\lambda\int_0^{\tau_k}\left[f''(x,\theta)\right]^2 dx$, penalizes the roughness of the fitting curve. A small penalty coefficient λ prioritizes the goodness of fit of the yield curve at the cost of less smoothness. A large λ delivers the opposite. Similar to Hu et al. (2013), we find that our results are robust to different choices of λ. Consequently, we set λ to zero in our baseline specification but show robustness with positive λ.

Set of knots, or maturities, $(\tau_0, \tau_1, \ldots, \tau_k)$: In fitting a spline curve with maturities up to ten years, Hu et al. (2013) use the three naturally important maturities in the Treasury market, two, five, and ten years, as the knots. Because QE auctions cover securities up to 30 years, we use 20 and 30 years as two additional knots. This setup involves 20 parameters overall, with four for each subinterval. But eight degrees of freedom are taken away due to the smoothness conditions at the 2-, 5-, 10-, and 20-year maturity junctions, so there are 12 free parameters in our baseline spline curve model. In robustness checks, we consider a larger set of knots.

Set of securities in the fit of the yield curve: In our baseline model, we include all outstanding Treasury securities except those with maturity less than one year and the most recently issued on-the-run securities. The yields on Treasury bills can reflect idiosyncratic supply or demand fluctuations beyond the risk-free rate (see Hu et al., 2013; Gurkaynak et al., 2007). On-the-run securities tend to have additional convenience yield tied to their use in the repo market and, hence, are more expensive than otherwise similar off-the-run securities.

Choice of market prices: We use the midpoint of bid-ask prices at 8:40 a.m. in the NPQS data. By construction, the 8:40 a.m. prices are obtained before the auction time and hence not affected by the auction outcomes.

With the parameter estimate θ_t, we can compute the model-implied price as $P_{tj}(\theta_t)$ and define the bond cheapness as

$$Cheapness_{tj} = \frac{P_{tj}(\theta_t) - P_{tj}}{P_{tj}}, \tag{6}$$

expressed in basis points. The cheapness measure is calculated for all eligible bonds in QE auctions, which include on-the-run securities but exclude Treasury bills. Our hypothesis is that bonds with a higher cheapness measure are more attractive to the Fed. We caution that correlated measures do not mean identical measures, and some difference can arise from different choices of the knots, the penalty function, the sampling times of market prices, and the set of included securities, among others. Again, we conduct robustness checks along all these dimensions and find similar results.

Fig. 2 illustrates the fitting of yield curve on two dates during the sample period. Bond cheapness as in our measure is positive whenever the model-implied yield ("x") is below the market yield ("o") and negative otherwise. (The solid line is the par yield implied from the fitted model, which is not used in the calculation of cheapness but plotted as a comparison.) According to our fitted yield curve, between 40% and 65% of bonds are deemed to be cheaper than model on any given day in our sample. Overall, our model fits the yields reasonably well. The root mean squared fitting error in yield, defined as

$$\sqrt{\frac{1}{N_t}\sum_{j=1}^{N_t}(y_t^j(\theta) - y_t^j)^2}, \tag{7}$$

where $y_t^j(\theta)$ is the model-implied yield and y_t^j is the market yield, has a mean of 4.8 bps and a standard deviation

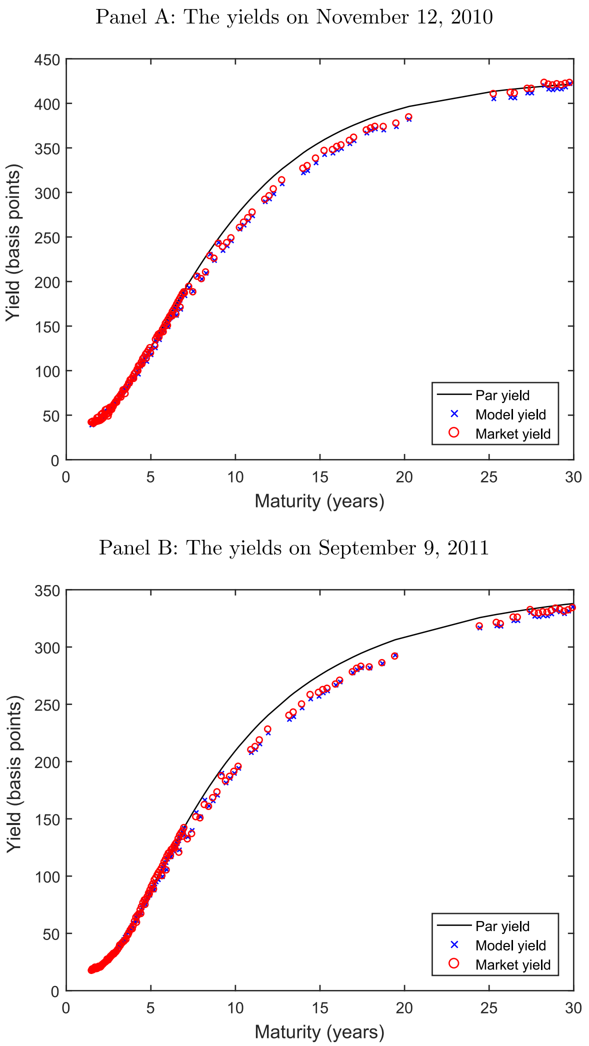

Fig. 2. Examples of fitted yield curves.
This figure plots the fitted bond yields ("x"), market bond yield ("o"), and the par-coupon yield (solid line), on November 12, 2010 (Panel A) and September 9, 2011 (Panel B).

of 1.1 bps. These fitting errors are moderate compared with those in term structure models of bond yield. For instance, Piazzesi (2010) reports that the standard level, slope, and curvature factors together deliver average absolute yield-fitting errors between 5 bps and 11 bps, depending on the maturity. Piazzesi (2010) further comments that additional structures imposed in affine term structure models tend to increase the fitting errors.

Although we do not observe the Fed's confidential yield curve model or dealers' information about it, we expect our bond cheapness measure to be positively correlated with both, for two reasons. First, the cubic spline model is a standard yield curve-fitting model used by institutional investors, broker-dealers, and the Fed in conducting QE auctions. Second, our NPQS data on the market prices of Treasury bonds are obtained from the Fed and primarily consist of quotes of primary dealers. Therefore, our model and data are similar to those used by the Fed and the dealers. Table 4 reports the time series correlations of the daily 2-, 5-, 10-, and 30-year par yields from our fitted spline yield curve with the corresponding par yields obtained from Barclays and those from Gurkaynak et al. (2007). The correlations are very high, around 0.98 or 0.99. These high correlations provide suggestive evidence that the cheapness measure from our yield curve model is reasonably correlated with that from the Fed's or dealers' yield curve models.

Table 4
Yield curve correlations with Barclays and Gurkaynak et al. (GSW, 2007)
This table reports time series correlations of the daily 2-, 5-, 10-, and 30-year par yields from our fitted spline yield curve with those from Barclays and from Gurkaynak et al. (GSW, 2007) over our sample period.

	2 year	5 year	10 year	20 year	30 year
Barclays	0.987	0.992	0.989	0.984	0.979
GSW	0.986	0.992	0.992	0.986	0.981

Table 5
Glossary of empirical measures at the auction-CUSIP level.

Measure	Description
Cheapness	Difference between the benchmark price implied from the fitted yield curve model and the actual market mid-price, normalized by the market mid-price (unit is basis points); see Eq. (6)
VOL	Pre-auction volatility; standard deviation of daily returns of the bond during the five trading days prior to the auction date (unit is percentage points)
Specialness	Difference between the general collateral repo rate and special repo rate on the bond (unit is percentage points)
OB	Outstanding balance; total outstanding par value of a particular bond (unit is \$10 billions)
Bid-Ask	Difference between the secondary market ask and bid quotes of a bond normalized by the mid-quote (unit is basis points)

While the cheapness measure is the most important one for our purpose, we do include other covariates that can affect the auction outcome. For example, a standard determinant of auction costs is the winner's curse (see Section 3). Following the literature, we measure winner's curse by the pre-auction volatility VOL_{tj}, computed as the standard deviation of daily returns of bond j for the five days of $t-1, t-2, \ldots, t-5$. In addition, we consider three measures of bond scarcity and illiquidity: specialness, outstanding balance, and bid-ask spread. As a measure for the scarcity of a bond, specialness is the difference between the general collateral repo rate and special repo rate on bond j and day t, in the unit of percentage points. The repo rates are recorded before 10 a.m. and, hence, are available before the auction starts.[17] Following Han et al. (2007), we define outstanding balance (OB_{tj}) as the total outstanding par value of the bond j that are not in STRIPS form, in the unit of \$10 billions, as of the day before the auction day.[18]

[17] If the special repo rate for a security is missing on one day, we use the observation on the previous day as a replacement, to avoid losing data.

[18] A Treasury bond can be stripped into multiple zero-coupon bonds, each corresponding to a single cash flow of the original Treasury bond.

Table 6
Summary statistics of empirical measures.
Panel A reports basic summary statistics of five empirical measures, computed across auctions and CUSIPs, all equal weighted. Panel B reports the correlation matrix of these five empirical measures by pooling observations across auctions and CUSIPs. Both the summary statistics and correlations are computed for auction-CUSIP pairs that are included (but not necessarily purchased) in quantitative easing auctions (“All included”) and for auction-CUSIP pairs that are purchased (“All purchased”).

Panel A: Basic summary statistics

Measure	Mean	Standard deviation	Minimum	25th percentile	50th percentile	75th percentile	Maximum
All included							
Cheapness	8.608	20.967	−66.087	−2.529	4.246	14.227	109.395
VOL	0.338	0.230	0.025	0.179	0.293	0.426	1.430
Specialness	0.026	0.038	−0.070	0.010	0.018	0.033	1.183
OB	2.990	1.438	0.329	2.218	3.014	3.638	7.160
Bid-Ask	3.510	1.571	0.732	2.341	3.074	4.446	9.193
All purchased							
Cheapness	12.146	22.879	−43.304	−0.433	7.220	17.931	109.395
VOL	0.370	0.233	0.027	0.218	0.324	0.456	1.430
Specialness	0.025	0.033	−0.060	0.008	0.016	0.032	0.532
OB	3.038	1.473	0.329	2.218	3.162	3.681	7.160
Bid-Ask	3.534	1.647	0.732	2.305	3.076	4.588	9.193

Panel B: Pooled correlation

	Cheapness	*VOL*	*Specialness*	*OB*	*Bid-Ask*
All included					
Cheapness	1				
VOL	0.594	1			
Specialness	−0.149	−0.0870	1		
OB	−0.402	−0.106	−0.190	1	
Bid-Ask	0.396	0.378	0.0992	−0.614	1
All purchased					
Cheapness	1				
VOL	0.606	1			
Specialness	−0.170	−0.148	1		
OB	−0.465	−0.216	−0.186	1	
Bid-Ask	0.405	0.412	0.0728	−0.622	1

Bid-ask spread ($Bid_{tj} - Ask_{tj}$) is the difference between the ask and the bid quotes of bond j, divided by the mid-quote, in the unit of basis points. We use the 8:40 a.m. NPQS quotes when computing bid-ask spread, similar to the calculation of cheapness.

For ease of reference, Table 5 tabulates the definitions of these empirical measures that we use as independent variables in regressions on the outcome of the auctions.

Panel A of Table 6 reports basic summary statistics of these five empirical measures (equally weighted) across both auctions and CUSIPs. We separately calculate the statistics for auction-CUSIP pairs that are included (but not necessarily purchased) in QE auctions (“All included”) and for auction-CUSIP pairs that are purchased (“All purchased”). The included bonds by the Fed are on average cheaper than the yield curve-implied value by 8.6 bps, with a standard deviation of 21 bps.[19] In comparison, the average cheapness measure of all purchased bonds by the Fed is 12.1 bps and has a standard deviation of about 23 bps. Because the average cheapness is higher on all purchased bonds than included bonds by about 3.5 bps, this evidence suggests that the Fed did purchase bonds that are cheaper.

On the other four measures, the included bonds and the purchased bonds are very similar. The average pre-auction volatilities of the two groups are 0.34% and 0.37%. The average specialness of the two groups is 2.6 and 2.5 bps. The average outstanding balances of both groups are about $30 billion. And the average bid-ask spreads of both groups are about 3.5 bps. (A side observation from Panel A of Table 6 is that the cheapness measure has some outliers. Our results are similar with or without winsorizing the cheapness measure.)

Panel B of Table 6 reports the correlation matrix of these five empirical measures, pooled across all auctions and CUSIPs, for both included and purchased bonds. While most correlations are small in magnitude, some are large, such as the the correlations between cheapness and volatility (positive), between cheapness and outstanding balance (negative), and between cheapness and bid-ask spread (positive). Thus, in the subsequent regression analysis, we include cheapness and other covariates both separately and jointly.

The STRIPS can also be reassembled or reconstituted in the right proportion to get back to the original form of the Treasury bond.

[19] It can appear puzzling that the average cheapness of all bonds is positive instead of close to zero. Our cheapness measure is not weighted by the inverse of duration as in Eq. (5). We compute the inverse duration-weighted fitting errors of our spline model as in Eq. (5) and find them to be tiny, only about one cent per $100 par value (or 1 bp) in magnitude.

Table 7
Cheapness of the bonds (not) purchased by the Fed.
This table reports the results of panel regression in Eq. (5). The indicator function 1(*Purchased by the Fed*) is equal to one if bond j is purchased in auction t and zero otherwise. The sample is 139 quantitative easing auctions of nominal Treasury securities executed from November 12, 2010 to September 9, 2011. *Cheapness* is in basis points (see Table 5). Robust t-statistics based on two-way clustered standard errors at auction and CUSIP levels are reported in parentheses. Significance levels are: *** for $p < 0.01$, ** for $p < 0.05$, and * for $p < 0.1$, where p is the p-value.

Independent variables	*Cheapness* (1)	*Cheapness* (2)	*Cheapness* (3)	*Cheapness* (4)
1(*Purchased by the Fed*)	3.925***	3.941***	6.471***	8.204***
	(11.305)	(7.459)	(6.424)	(3.890)
N	3432	3432	3432	3432
R^2	0.950	0.898	0.757	0.038
CUSIP fixed effects	Yes	Yes	No	No
Auction fixed effects	Yes	No	Yes	No

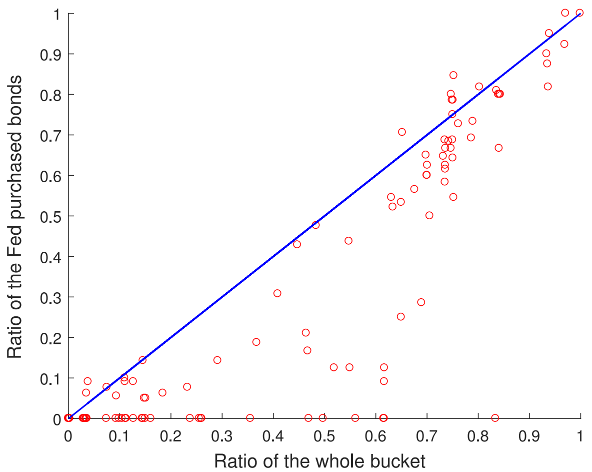

Fig. 3. The ratio of rich bonds.
The x-axis shows the ratio of the number of rich bonds in the auctioned maturity bucket, and the y-axis shows the ratio of the number of rich bonds among those purchased by the Fed in the corresponding auction.

5. The Fed's purchase quantities

In this section, we study the determinants of the purchase quantities by the Fed across auctions and CUSIPs. The prediction is that cheaper bonds are purchased more on average, controlling for the other covariates.

Fig. 3 provides a scatter plot of the ratio of the number of rich bonds (i.e., cheapness is negative based on our measure) in the auctioned maturity bucket (x-axis) against the ratio of the number of rich bonds among those purchased by the Fed in the corresponding auction (y-axis). If the Fed prefers cheaper (richer) bonds, the dots should lie mostly below (above) the 45-degree line. The Fed's purchases tilt toward cheaper bonds, as predicted. Among the 139 auctions, 40 of them have both rich ratios equal to zero, shown as 40 overlapping dots at (0, 0). For these auctions, the Fed buys only cheap bonds because all bonds in the maturity bucket are cheaper than model. In 31 auctions, the Fed buys only cheap bonds even though some rich bonds are available; these are the dots on the x-axis but not at (0, 0). In 68 auctions, the Fed buys some rich bonds. But, for most of these auctions, rich bonds represent a smaller fraction among the purchased bonds than among the included bonds in the auctions. Overall, only in 11 out of the 139 auctions is the rich ratio for the Fed's purchased bonds greater than the rich ratio for the auctioned bucket, which happens mostly when the auctioned maturity bucket is predominantly rich to start with. One possible reason for the Fed to purchase any rich bond is that rich bonds tend to be more liquid (see Table 6) and could be offered in larger amounts by dealers.[20] Cheap bonds can be particularly difficult for dealers to locate or source if most of the bonds in the auctioned maturity bucket are rich. Measurement error in our cheapness proxy is another possible reason that the Fed buys some rich bonds.

We now move on to regression analysis, starting with the following simple panel regression:

$$Cheapness_{tj} = \sum_j \alpha_j D_j + \sum_t \alpha_t D_t + \beta \cdot 1(Purchased\ by\ the\ Fed) + \epsilon_{tj}, \tag{8}$$

where 1(*Purchased by the Fed*) is equal to one if bond j is purchased by the Fed in auction t and zero otherwise and D_j and D_t are CUSIP and auction fixed effects, respectively. The standard errors are two-way clustered by auction and CUSIP (Petersen, 2009), which is the case for all other regressions.

Table 7 reports the results of regression in Eq. (8). Column 1 shows that a bond purchased by the Fed is on average 3.9 bps cheaper (relative to model) than a bond not purchased by the Fed, controlling for auction and CUSIP fixed effects. In other columns, when one or both fixed effects are dropped, the statistical and economical significance of the coefficient on cheapness is similar or larger.

Next, we run the following panel regression of the auction purchase amount q_{tj} (in billions of dollars) for bond j in auction t on empirical measures of bond cheapness, controlling for the other four covariates:

$$q_{tj} = \sum_j \alpha_j D_j + \sum_t \alpha_t D_t + \beta_1 \cdot Cheapness_{tj} + \beta_2 \cdot VOL_{tj} + \beta_3 \cdot Specialness_{tj} + \beta_4 \cdot OB_{tj} + \beta_5 \cdot Bid\text{-}Ask_{tj} + \epsilon_{tj}, \tag{9}$$

[20] The dealers' supply is a conjecture because we do not observe the rejected offers and, hence, cannot distinguish dealers' supply from the Fed's demand.

Table 8
Regressions of purchase quantity.
This table reports results of panel regression in Eq. (9). The purchase quantity q_{tj} of bond j in auction t is in billions of dollars. Columns 1–5 report the results for all included bonds. Columns 6 and 7 report the results for all purchased bonds. The sample is 139 quantitative easing auctions of nominal Treasury securities executed from November 12, 2010 to September 9, 2011. The explanatory variables are *Cheapness* (in basis points), *VOL* (in percentage points), *Specialness* (in percentage points), *OB* (in \$10 billions), and *Bid-Ask* (in basis points), as defined in Table 5. Robust t-statistics based on two-way clustered standard errors at auction and CUSIP levels are reported in parentheses. Significance levels are: *** for $p < 0.01$, ** for $p < 0.05$, and * for $p < 0.1$, where p is the p-value.

Explanatory variables	q_{tj} (1)	q_{tj} (2)	q_{tj} (3)	q_{tj} (4)	q_{tj} (5)	q_{tj} (6)	q_{tj} (7)
Cheapness	0.012***		0.012***	0.006***	0.012***	0.013***	0.013***
	(4.439)		(4.179)	(4.206)	(4.895)	(3.423)	(3.333)
VOL		−2.222**	−2.137**	−0.874**	0.120		−3.600**
		(-2.320)	(-2.135)	(−2.393)	(1.199)		(−2.147)
Specialness		0.369	0.412	1.582*	0.151		0.594
		(0.637)	(0.688)	(1.859)	(0.249)		(0.519)
OB		0.087*	0.039	−0.001	−0.009		0.008
		(1.943)	(0.810)	(−0.030)	(−0.234)		(0.043)
Bid-Ask		−0.177***	−0.173***	−0.154***	−0.110***		−0.211***
		(−4.988)	(−4.945)	(−5.835)	(−4.500)		(−4.552)
N	3432	3432	3432	3432	3432	1952	1952
R^2	0.327	0.341	0.350	0.136	0.298	0.423	0.448
CUSIP fixed effects	Yes	Yes	Yes	No	Yes	Yes	Yes
Auction fixed effects	Yes	Yes	Yes	Yes	No	Yes	Yes

where D_j and D_t are CUSIP and auction fixed effects, respectively. As before, we use two-way clustered standard errors.

Table 8 reports the results of regression in Eq. (9). Columns 1–5 report the results with all included CUSIPs, and Columns 6 and 7 report the results with only purchased CUSIPs. As predicted, cheapness has a statistically and economically significant coefficient on purchase quantity, both as a single regressor in Column 1 and jointly with other covariates in Column 3. A coefficient of 0.012 implies that a bond that is 1 bp cheaper is purchased by \$12 million more. A one standard deviation increase in cheapness, which is about 23 bps (see Table 6), is associated with a higher purchase quantity of about \$276 million.

In Columns 4 and 5, we include one fixed effect at a time. With only the CUSIP fixed effects, the coefficient on cheapness remains 0.012 and the R^2 decreases only moderately. But with only the auction fixed effects, the coefficient on cheapness drops by about half (but still significant) and so does the R^2. This suggests that the explanatory power of the five measures is stronger for auction-to-auction variation in purchase quantities than for CUSIP-to-CUSIP variation.

Columns 6 and 7 repeat the regressions on the sample of bonds that are purchased by the Fed, with all five measures and both fixed effects. The effect of cheapness on purchase quantity is very similar to and slightly larger than that in the sample of included bonds; that is, a 1 bp increase of cheapness leading to a \$13 million increase of the purchase quantity.

Overall, we find strong evidence that the Fed purchases a larger quantity of bonds whose market prices are cheaper relative to the prices implied from a spline model. And, the economic significance of cheapness is large.

Finally, we briefly discuss the coefficients on the other four measures: volatility, specialness, outstanding balance, and bid-ask spread. Columns 3 and 7 show that, controlling for all else, the Fed buys a smaller quantity of bonds that have higher volatilities and wider bid-ask spreads, which tend to be associated with lower liquidity. This pattern can be interpreted in two ways. First, the measurement error in our cheapness proxy could be correlated with volatility and bid-ask spread, so that the Fed's true preference among various bonds seems to lean toward lower volatilities and narrower bid-ask spreads, controlling for our measure of cheapness. Second, that the Fed does not necessarily prefer more liquid bonds, but dealers have more difficulty obtaining bonds with higher volatilities and wider bid-ask spreads from long-term investors. Telling them apart would be possible if we were able to observe all of dealers' offers (not only accepted offers but also rejected ones), but such data are unavailable to us.

6. The Fed's costs

Following the literature (see, e.g., Cammack, 1991; Nyborg and Sundaresan, 1996; Han et al., 2007; Hortacsu and Kastl, 2012), we measure the Fed's cost by the auction price markup, namely, the difference between the price paid by the Fed and the secondary market price on the days the auctions are executed.[21] Let $p_{t,j,d,o}$ and $q_{t,j,d,o}$ be the oth winning offer price and par value from dealer d on CUSIP j in auction t, and recall that $P_{t,j}$ is the secondary market price of CUSIP j at the time auction t is closed. Then, the weighted-average price markup on bond j in auction t is

$$Markup_{tj} = \frac{\sum_{d,o} \left(p_{t,j,d,o} - P_{t,j}\right) \cdot q_{t,j,d,o}}{\sum_{d,o} q_{t,j,d,o}}. \tag{10}$$

[21] 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 correlation between the average price-based cost and the worst price-based cost is about 99%.

Table 9
Summary statistics of the Fed's cost across CUSIPs and auctions.

This table presents summary statistics of the Fed's cost (in cents per $100 par value) across CUSIPs and auctions from November 12, 2010 to September 9, 2011. The quantity-weighted mean, standard deviation, and *t*-statistic are in Panel A. The percentiles are in Panel B. 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 closest ask and the 3:30 p.m. ask. For auctions closed at 11:00 a.m., 11:30 a.m., and 12:05 p.m., the closest ask equals the 11:30 a.m. secondary market ask price. For auctions closed at 2:00 p.m., the closet ask equals the 2:15 p.m. ask price.

Panel A: Basic summary

Statistic	N	Quantity weighted mean	Quantity weighted standard deviation	t-statistic
Cost to closest ask	1953	0.71	11.22	2.80
Cost to 3:30 p.m. ask	1953	2.73	22.11	5.46

Panel B: Percentiles

	Minimum	1st	10th	50th	90th	99th	Maximum
Cost to closest ask	−61.77	−40.02	−9.48	0.42	13.95	34.04	60.99
Cost to 3:30 p.m. ask	−117.19	−62.94	−18.84	0.78	27.95	67.82	89.89

By definition, this markup measure is meaningful only for purchased bonds.

Which secondary market prices do we use for $P_{t,j}$? Ideally, we would want secondary market prices at the moment the auction is closed, but the NPQS data provide bid and ask prices only at 8:40 a.m., 11:30 a.m., 2:15 p.m., and 3:30 p.m. In our sample, 134 of the 139 auctions are closed at 11:00 a.m. Given the data limitation, we use the 11:30 a.m. ask price as $P_{t,j}$ for auctions closing at 11:00 a.m., 11:30 a.m., and 12:05 p.m. and use the 2:15 p.m. ask price for auctions closing at 2:00 p.m. We refer to these ask prices as the "closest ask" because they are as close to the auction closing times as possible in our data.

We denote by $Markup^{c}_{tj}$ the auction price markup based on the closest ask, calculated from Eq. (10). The time gap between auction closing and price quotes implies that, for the vast majority of the auctions, $Markup^{c}_{tj}$ measures the realized cost of the Fed 30 min after the auction, which can be interpreted as the profit of dealers by selling the bonds to the Fed in the auction and then covering their short positions in the secondary market 30 min later. In addition, we compute the price markup using the 3:30 p.m. ask price and denote it by $Markup^{3:30}_{tj}$. In a similar interpretation, $Markup^{3:30pm}_{tj}$ is the Fed's realized cost 4.5 h after the auction. Realized costs constructed this way are standard measures in market microstructure (see Hasbrouck, 2007).

Panel A of Table 9 presents summary statistics of the Fed's auction cost $Markup_{tj}$, in cents per $100 par value. The average cost over all purchase auctions, weighted by purchase quantity, is 0.71 cents measured relative to the closest ask and 2.73 cents relative to the 3:30 p.m. ask in the secondary Treasury market. To put the magnitude of $Markup_{tj}$ 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 (in Table 6, the mean bid-ask spread of 3.5 bps is equal weighted). Therefore, the average QE auction cost is below or comparable to the average bid-ask spread. These results suggest that the Fed suffers only moderate market-impact costs in purchasing the huge amount ($780 billion) of Treasury securities in our sample period. The aggregate dollar cost is far from trivial, $55 million at 0.71 cents per $100 par value and 213 million at 2.73 cents per $100 par value.

It is informative to compare the average cost in QE auctions with the cost of previous Treasury buyback operations, Treasury issuance auctions, and QE operations in other countries. Han et al. (2007) find that the Treasury's buyback operations from March 2000 to April 2002 incurred the average cost of 4.38 cents per $100 par value, relative to near-simultaneous market ask prices provided by GovPX. This is the only prior study we are aware of that focuses on the cost of buying back US Treasury securities. 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, relative to transaction prices in the when-issued market over a short time window around the auction time. Cammack (1991) and Nyborg and Sundaresan (1996) provide similar estimates in earlier samples.[22] In addition, Breedon and Turner (2016) find, in the European market, that the Bank of England's QE operations incur an average cost of 0.72 bps, in terms of yield, relative to end-of-day market yield.

Panel B of Table 9 reports various percentiles of the auction cost. Some outliers exist on both ends of the distribution. For example, a large difference is evident between the minimum and the 1% percentile, especially for $Markup^{3:30\,pm}_{tj}$, and between the maximum and the 99% percentile, especially for $Markup^{c}_{tj}$. Although we do not have definitive evidence as to why these large outliers oc-

[22] More recently, Fleming and Liu (2014) find that, right after issuance auctions, two-, five-, and ten-year notes have four-h returns between 1.49 and 12.51 bps, suggesting nontrivial costs of temporary price impact borne by the US Treasury. Lou et al. (2013) find that prices of US Treasury securities drop a few days prior to scheduled Treasury auctions and recover afterward.

Table 10
Panel regression of auction cost across CUSIPs and auctions.
This table reports results of panel regression in Eq. (11). The sample is the 139 quantitative easing auctions of nominal Treasury securities executed from November 12, 2010 to September 9, 2011. The auction price markup $Markup^c_{tj}$ is computed as the average, weighted by the amount of each accepted offer, of the differences between the accepted offer price and the corresponding closest ask in the secondary market. $Markup^{3:30}_{t,j}$ is measured relative to the 3:30 p.m. ask price. The explanatory variables are *Cheapness* (in basis points), *VOL* (in percentage points), *Specialness* (in percentage points), *OB* (in \$10 billions), and *Bid-Ask* (in basis points), as defined in Table 5. Robust t-statistics based on two-way clustered standard errors at auction and CUSIP levels are reported in parentheses. Significance levels are: *** for $p<0.01$, ** for $p<0.05$, and * for $p<0.1$, where p is the p-value.

Explanatory variables	$Markup^c_{tj}$ (1)	$Markup^c_{tj}$ (2)	$Markup^c_{tj}$ (3)	$Markup^c_{tj}$ (4)	$Markup^c_{tj}$ (5)	$Markup^{3:30}_{tj}$ (6)	$Markup^{3:30}_{tj}$ (7)
Cheapness	0.102**		0.115***	0.054***	0.275*	0.127**	0.142***
	(2.567)		(2.968)	(2.828)	(1.936)	(2.335)	(2.700)
VOL		2.944	4.122	−3.885	15.376		−4.545
		(0.480)	(0.676)	(−1.206)	(1.426)		(−0.368)
Specialness		−8.948**	−8.639**	6.124*	22.192		−6.178*
		(-2.289)	(−2.204)	(1.656)	(1.251)		(−1.699)
OB		−2.658***	−3.281***	0.003	−5.085**		−4.577***
		(−3.606)	(-3.935)	(0.016)	(−2.039)		(−6.697)
Bid-Ask		−0.628***	−0.649***	−0.224*	−0.746		−0.496*
		(−4.181)	(−4.347)	(−1.784)	(−0.763)		(−1.764)
N	1934	1934	1934	1934	1934	1934	1934
R^2	0.963	0.964	0.965	0.952	0.179	0.977	0.977
CUSIP fixed effects	Yes	Yes	Yes	No	Yes	Yes	Yes
Auction fixed effects	Yes	Yes	Yes	Yes	No	Yes	Yes

cur, conversations with the Fed suggest that the very left tail of the cost distribution could be due to very low priced offers from anxious sellers (dealers or their customers), and the very right tail of the cost distribution could be a result of monopolistic dealers controlling a large fraction of the supply of particular CUSIPs. To make sure that our results are not unduely affected by outliers, in conducting the regressions below, we trim observations for which the measured auction cost is smaller than the 0.5% percentile or larger than the 99.5% percentile across auction-CUSIP pairs (truncating more leads to similar or even stronger results). Because there are about 1953 auction-CUSIP pairs with the cost measure, we only lose 20 data points.

We now investigate the economic determinants of the auction cost. Our main hypothesis is that cheaper bonds incur a higher cost to the Fed, controlling for other covariates. To test it, we run the panel regression of the form:

$$\begin{aligned} Markup_{tj} = & \sum_j \alpha_j D_j + \sum_t \alpha_t D_t + \beta_1 \cdot Cheapness_{tj} \\ & + \beta_2 \cdot VOL_{tj} + \beta_3 \cdot Specialness_{tj} + \beta_4 \cdot OB_{tj} \\ & + \beta_5 \cdot Bid\text{-}Ask_{tj} + \epsilon_{tj}. \end{aligned} \tag{11}$$

This regression is run on the cost to closest ask and the cost to the 3:30 pm ask. As before, all standard errors are two-way clustered by auction and by CUSIP (Petersen, 2009).[23]

Columns 1–5 of Table 10 report the panel regression of the auction cost relative to the closest ask, $Markup^c_{tj}$, in various specifications. We find that the Fed pays a higher cost on cheaper bonds, controlling for the other four covariates and the two fixed effects. In Column 3, the main specification, the coefficient of 0.115 on cheapness means that, controlling for all else, a 23 bps increase in cheapness (roughly one standard deviation of cheapness across auctions and CUSIPs in our sample according to Table 6) increases the auction cost by about 2.6 cents ($= 23 \times 0.115$) per \$100 par value. This magnitude is economically large, given the low average cost in QE auctions.

Columns 4 and 5 report the results from regressions with one fixed effect at a time. Cheapness remains significant in both specifications. The regression R^2 reveals that the variation of auction costs is primarily driven by its time series variation than cross-section variation, which seems intuitive. Because the total size of each QE auction is relatively inflexible, the Fed has limited ability to smooth its purchases over time. But the Fed has substantial flexibility to choose among different CUSIPs, and this endogenous choice tends to reduce the cost variation across CUSIPs.

Columns 6 and 7 report the panel regression of the auction cost relative to the 3:30 p.m. ask. In Column 7, the coefficient for cheapness is 0.142. So, the same 23 bps increase in cheapness (one standard deviation) increases the auction cost by about 3.3 cents per \$100 par value. Not only is the average auction cost larger if measured at the end of day, but the dependence of auction cost on cheapness is also stronger at the end of day.

In Columns 3 and 7, controlling for all else, the Fed pays a higher cost on bonds that have a lower specialness, a lower outstanding balance, and a lower bid-ask spread. Again, these patterns could be a result of the correlation between the measurement error of the cheapness measures and other covariates.

7. Heterogeneity in dealers' volume and profitability

The implementation cost of the Federal Reserve is the profit of primary dealers. Taking advantage of dealers'

[23] In unreported results, we also include the auction number in the regression as an explanatory variable, to control for the possibility that dealers learn about the bond values or the Fed's internal benchmark prices over time (Milgrom and Weber, 1982; Ashenfelter, 1989; Han et al., 2007). We do not find any time trend in the data.

Table 11
Dealer profitability.
This table summarizes the dealers' profitability across the 139 quantitative easing auctions of nominal Treasury securities from November 12, 2010 to September 9, 2011. We compute, for each dealer, the profit margin as the average (weighted by the amount of each accepted offer) of the differences between the offer price and the corresponding closest ask for the bond, 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. The dealers are ranked according to their aggregate profits. We report the rankings, dealer identities, aggregate profits in millions of dollars, profit margins in cents per $100 par value, aggregate offer amount in billions of dollars of par value sold to the Fed, total number of auctions, number of winning offers per auction, and total number of CUSIPs sold.

Rank by aggregate profit	Dealer identity (1)	Aggregate profit (millions of dollars) (2)	Profit margin (cents per $100) (3)	Aggregate offer amount (billions of dollars) (4)	Total number of auctions (5)	Number of winning offers per auction (6)	Total number of winning CUSIPs (7)
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 Paribas	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	BAML	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 Capital	−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

identities in our data set, we now study the granular heterogeneity of profitability across primary dealers and the likely determinants of the heterogeneity.

For each dealer d, we compute his aggregate profit as

$$Aggregate\ Profit_d = \sum_{t,j,o} \left(p_{t,j,d,o} - P_{i,j}\right) \cdot q_{t,j,d,o} \tag{12}$$

and profit margin as

$$Margin_d = \frac{\sum_{t,j,o} \left(p_{t,j,d,o} - P_{i,j}\right) \cdot q_{t,j,d,o}}{\sum_{t,j,o} q_{t,j,d,o}}, \tag{13}$$

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

Table 11 summarizes the dealer-by-dealer profitability across the 139 QE auctions in our sample period. The data are organized by the rankings according to dealers' aggregate profits. The dealers' identities are provided in Column 1. Columns 2, 3, and 4 provide, respectively, the aggregate profit in millions of dollars, the profit margin in cents per $100 par value, and the aggregate accepted offer amount in billion dollars of par value. In addition, we also report the total number of auctions in Column 5, the number of winning offers per auction in Column 6, and the total number of CUSIPs sold in Column 7.

From columns 2–4, we observe a striking concentration of aggregate profits and aggregate amounts of accepted offers among dealers. The top five dealers (Morgan Stanley, 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, measured relative to the closest ask. These five dealers also account for about half of the $776.6 billion total purchase amount. The bottom five dealers (RBS, Credit Suisse, UBS, Barclays Capital, and Citigroup) account for about 40% of the $776.6 billion total purchase amount, but their profit margins are much lower. The bottom five dealers all have negative profit margins, averaging −0.44 cents per $100 par value, and the top five dealers all have positive profit margins, averaging 1.67 cents per $100 par value.

The concentration of profits for the top few dealers in QE auctions is similar to that in other over-the-counter markets, including municipal bonds, corporate bonds, and asset-backed securities (see Li and Schurhoff, 2014; Di-Maggio et al., 2017; Hollifield et al., 2017).

At least two potential reasons exist for the strong heterogeneity of dealers' profitability. The first is that the top dealers could be more capable of finding bonds that the Fed wishes to purchase. The second is that, conditional on selling the same bond, the top five dealers could be able to charge a higher price markup than other dealers.

To investigate the first potential explanation, we run the following panel regression of sales quantity at the auction-CUSIP-dealer level:

$$q_{tjd} = \sum_j \alpha_j D_j + \sum_t \alpha_t D_t + \beta_1 \cdot Cheapness_{tj} + \beta_2 \cdot VOL_{tj} + \beta_3 \cdot Specialness_{tj} + \beta_4 \cdot OB_{tj} + \beta_5 \cdot Bid\text{-}Ask_{tj}$$

Table 12
Panel regression of markup and quantity at auction-CUSIP-dealer level.

This table reports results of panel regressions in Eqs. (14) and (15), for each bond purchased by each dealer in each of the 139 quantitative easing auctions of nominal Treasury securities executed from November 12, 2010 to September 9, 2011. The explanatory variables are *Cheapness* (in basis points), *VOL* (in percentage points), *Specialness* (in percentage points), *OB* (in \$10 billions), and *Bid-Ask* (in basis points), as defined in Table 5. The dummy variable 1(*Top 5*) is equal to one if the dealer is among the top five in Table 11 and zero otherwise. Robust *t*-statistics based on two-way clustered standard errors at auction and CUSIP levels are reported in parentheses. Significance levels are: *** for $p<0.01$, ** for $p<0.05$, and * for $p<0.1$, where p is the p-value.

	q_{tjd}				$Markup^c_{tjd}$			
Variable	All (1)	All (2)	Top five (3)	Others (4)	All (5)	All (6)	Top five (7)	Others (8)
Cheapness	0.0002	0.0012**	0.0026***	0.0003	0.0705*	0.0608*	0.0412	0.0673*
	(0.3315)	(2.2748)	(2.7224)	(0.5685)	(1.9130)	(1.6750)	(1.0277)	(1.7701)
VOL	−0.2552**	−0.6185**	−0.9997**	−0.2539**	−3.7973	−0.2697	3.4750	−3.8353
	(−2.3829)	(−2.4436)	(−2.2583)	(−2.3762)	(−0.8056)	(−0.0569)	(0.5277)	(−0.8170)
Specialness	−0.1699	−0.0028	0.1172	−0.1699	−5.1007*	−4.1191	−3.6786	−5.0964*
	(−1.3029)	(−0.0147)	(0.3269)	(−1.3075)	(−1.6878)	(−1.4044)	(−0.9622)	(−1.6972)
OB	−0.0003	0.0042	0.0272	−0.0012	−1.9063***	−1.7785***	−1.2930***	−1.8871***
	(−0.0259)	(0.2350)	(0.7382)	(−0.0952)	(−3.1124)	(−3.8824)	(−2.5989)	(−3.0486)
Bid-Ask	−0.0123**	−0.0221***	−0.0362***	−0.0123**	−0.7309***	−0.6599***	−0.5558***	−0.7297***
	(−2.4668)	(−4.3756)	(−3.9614)	(−2.4914)	(−4.1953)	(−4.2222)	(−2.9325)	(−4.2138)
1(*Top 5*)	0.2649***				−3.7919**			
	(2.6544)				(−2.1792)			
1(*Top 5*)· *Cheapness*	0.0023***				−0.0303			
	(2.9827)				(−1.0020)			
1(*Top 5*)· *VOL*	−0.7543*				7.3785			
	(−1.9221)				(1.3045)			
1(*Top 5*)· *Specialness*	0.2876				1.4202			
	(0.9046)				(0.4223)			
1(*Top 5*)· *OB*	0.0287				0.6080			
	(1.0153)				(0.8339)			
1(*Top 5*)· *Bid-Ask*	−0.0241***				0.1764			
	(−2.7251)				(1.0193)			
N	7397	7397	3330	4067	7323	7323	3285	4038
R^2	0.2038	0.1310	0.2125	0.1655	0.9514	0.9486	0.9524	0.9506
CUSIP fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Auction fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
1(*Top 5*)·CUSIP fixed effects	Yes	No	No	No	Yes	No	No	No
1(*Top 5*)·Auction fixed effects	Yes	No	No	No	Yes	No	No	No

$$+1(d\in Top\ 5)\cdot\Bigg[\sum_j \alpha'_j D_j+\sum_t \alpha'_t D_t$$
$$+\beta'_1\cdot Cheapness_{tj}+\beta'_2\cdot VOL_{tj}+\beta'_3\cdot Specialness_{tj}$$
$$+\beta'_4\cdot OB_{tj}+\beta'_5\cdot Bid\text{-}Ask_{tj}\Bigg]+\epsilon_{tjd}, \tag{14}$$

where q_{tjd} is dealer d's sales quantity of bond j in auction t. The dummy variable $1(d\in Top\,5$ is equal to one if dealer d is among the top five dealers in Table 11 and zero otherwise. The first two lines of Eq. (14) are very similar to Eq. (9), but the last two lines of Eq. (14) capture the additional behavior of the top five dealers relative to the other dealers.

Analogously, to investigate the second potential explanation of dealers' heterogeneity, we run the panel regression of price markups at the auction-CUSIP-dealer level:

$$Markup^c_{tjd}=\sum_j \alpha_j D_j+\sum_t \alpha_t D_t+\beta_1\cdot Cheapness_{tj}$$
$$+\beta_2\cdot VOL_{tj}+\beta_3\cdot Specialness_{tj}+\beta_4\cdot OB_{tj}$$
$$+\beta_5\cdot Bid\text{-}Ask_{tj}+1(d\in Top\ 5)\cdot\Bigg[\sum_j \alpha'_j D_j$$
$$+\sum_t \alpha'_t D_t+\beta'_1\cdot Cheapness_{tj}+\beta'_2\cdot VOL_{tj}$$
$$+\beta'_3\cdot Specialness_{tj}+\beta'_4\cdot OB_{tj}$$
$$+\beta'_5\cdot Bid\text{-}Ask_{tj}\Bigg]+\epsilon_{tjd}, \tag{15}$$

where $Markup^c_{tjd}$ is dealer d's profit margin of selling bond j in auction t, relative to the closest ask.

Table 12 reports the results of regressions in Eqs. (14) and (15), together with results from one pooled specification and two separate regressions for the two subgroups of dealers, all without interaction terms involving $1(d\in Top\,5)$.

Column 1 shows the result of regression in Eq. (14). The coefficient of 0.0002 on cheapness suggests that a 1 bp increase in bond cheapness increases a non-top five dealer's sales quantity by 0.2 million, although this increase is statistically insignificant. The significantly positive coefficient of 0.0023 on 1(*Top 5*)·*Cheapness* indicates that a top five dealer's sales quantities go up by additional 2.3 million (hence, 2.5 million total) if bond cheapness increases by 1 bp. Columns 3 and 4, which report separate regressions for the two groups of dealers, further confirm this large difference. Column 2 shows that the dealers' average sensitivity of sales quantity to cheapness is 0.0012, or 1.2 million per dealer if cheapness increases by 1 bp.

The coefficients on other covariates are consistent with the auction-CUSIP-level results in Table 8. That is, dealers' sales quantities tend to be larger if a bond is less volatile or has a lower bid-ask spread. In addition, relative to non-top five dealers, the top five dealers' sales quantities are four times as sensitive to volatility and three times as sensitive to bid-ask spread.

Column 5 shows the result of regression in Eq. (15). Interestingly, per 1 bp increase in cheapness, the non-top five dealers seem to extract a higher price markup per unit of bond sold, of 0.07 cent per \$100 par value, relative to the top five dealers, whose corresponding slope is 0.04 cent per \$100 par value. Columns 7 and 8 again confirm this difference. The coefficients on other covariates are consistent with the auction-CUSIP-level results in Table 10. As in the case of cheapness, relative to top five dealers, non-top five dealers' price markups are 1.4 time as sensitive to specialness, 1.5 times as sensitive to outstanding balance, and 1.3 times as sensitive to bid-ask spread.

Overall, the evidence from Table 12 reveals that the top five dealers make higher profits mainly because they are more capable of finding and selling to the Fed bonds that are cheaper (than model), are less volatile, and have lower bid-ask spreads. In the data, these bond characteristics are strongly associated with higher purchase quantities by the Fed. The other dealers seem more strategic in charging the Fed slightly higher prices, but their price advantage is moderate and not large enough to offset their quantity disadvantage.

8. Conclusion

Since the global financial crisis, central banks around the world have implemented quantitative easing in a large scale. This paper provides the first empirical analysis of QE auctions in the United States, which the Fed uses to purchase US Treasury securities. Relative to Treasury issuance auctions, a unique feature of QE auctions is that the Fed uses a relative value approach, based on an internal yield curve model, to evaluate dealers' offers on multiple, substitutable bonds. By purchasing undervalued bonds, the Fed harvests gains from trade, but this preference could be partly inferred from public price data. This trade-off suggests that bonds that appear cheaper should be purchased more by the Fed but also incur the Fed a higher cost.

Using a proprietary, transaction-level data set that contains outcomes of QE auctions from November 2010 to September 2011, we find that bond cheapness is a significant predictor of the Fed's purchase quantity and cost, controlling for standard covariates. A one standard deviation increase in cheapness increases purchase quantity by 276 million and costs by 2.6 cents per \$100 par value. Moreover, dealers show remarkable heterogeneity in their transaction volume and aggregate profits. More granular panel regressions reveal that the top five dealers have a large comparative advantage in finding Treasury bonds that the Fed seems to prefer, whereas the other dealers have a small comparative advantage in charging slightly higher prices.

A potential direction of future research is to apply a similarly granular analysis on the implementation mechanism of other major central banks. For example, the Bank of England also uses reverse auctions, but it evaluates offers on different gilts based purely on secondary market prices. By not using an internal model, the Bank of England's mechanism seems more transparent and less discretionary than the Fed's. In implementing the European Central Bank's QE program, most national central banks in the Eurozone conduct purchases bilaterally, which have less transparency but more discretion. (A few national central banks in the Eurozone use auctions, however.) The wide heterogeneity of implementation mechanisms among major central banks can provide more insights about the costs and benefits of each mechanism and raise the policy question of whether a different mechanism should be tried on a pilot program and finally adopted. Given that QE operations tend to be large and could become a standard tool in monetary policy, any improvement in the implementation mechanism, in terms of faster fill of targeted purchase amount or lower cost, could bring substantial efficiency gains.

Appendix. Robustness

In this Appendix, we conduct robustness checks on our main empirical findings of the determinants of purchase quantities and costs in QE auctions.

From Eqs. (5) and (6), our baseline bond cheapness measure is $Cheapness_{tj} = P_{tj}(\theta_t)/P_{tj} - 1,$ where $P_{tj}(\theta_t)$ is the theoretical price of the bond based on the cubic spline model estimated using all available bond prices and P_{tj} is the market price. We consider seven perturbations of the baseline model.

(1) Instead of using a penalty coefficient of $\lambda = 0$ in Eq. (5), we use $\lambda = 0.01$.
(2) Instead of using the maturity knots of 2, 5, 10, 20, and 30 years, we use 1, 2, 3, 5, 7, 10, 20, and 30 years. This setup is labeled "Cubic18" as there are 18 free parameters in fitting the yield curve.
(3) Instead of using the bond prices to construct cheapness, we use the simple yield difference, $y_{tj} - y_{tj}(\theta_t),$ where $y_{tj}(\theta_t)$ is the model-implied yield and y_{tj} is the market yield.
(4) Instead of using the midpoint of bid and ask prices in fitting the model, we use the ask price.
(5) Instead of excluding all Treasury securities with maturities less than one year, we exclude those with maturities less than one month.
(6) Instead of using the 8:40 a.m. prices on the auction day to fit the yield curve, we use the 3:30 p.m. prices on the day before the auction.
(7) Instead of using the cubic spline, we use Svensson (1994) model of the yield curve. Svensson (1994) model assumes the following specification of the instantaneous forward rate:

$$f(m) = \beta_0 + \beta_1 \exp\left(-\frac{m}{\tau_1}\right) + \beta_2 \frac{m}{\tau_1} \exp\left(-\frac{m}{\tau_1}\right) + \beta_3 \frac{m}{\tau_2} \exp\left(-\frac{m}{\tau_2}\right), \tag{A1}$$

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 parame-

Table A1
Alternative measures of cheapness: purchase quantity.

This table reports results of panel regressions of the auction purchase quantity q_{tj} (in billions of dollars), for each included bond of the 139 quantitative easing auctions of nominal Treasury securities executed from November 12, 2010 to September 9, 2011. The explanatory variables are *Cheapness* (in basis points), *VOL* (in percentage points), *Specialness* (in percentage points), *OB* (in \$10 billions), and *Bid-Ask* (in basis points), as defined in Table 5. Different versions of cheapness are fitting the spline curve model with the roughness penalty coefficient $\lambda = 0.01$; fitting the spline curve model with knot points of 1, 2, 3, 5, 7, 10, 20, and 30 years ("Cubic18"); the difference between the yield to maturity of the market and model prices ("Yield-based"); fitting the curve using ask prices ("Ask based"); fitting the curve using the bonds with maturities greater than one month ("Maturity $\geq$ 1 month"); fitting the curve using the bond price quotes one day before the auction ("Previous day"); and fitting the curve using the Svensson (1994) model ("Svensson"). Robust *t*-statistics based on two-way clustered standard errors at the auction and CUSIP levels are reported in parentheses. Significance levels are: *** for $p < 0.01$, ** for $p < 0.05$, and * for $p < 0.1$, where p is the p-value.

Explanatory variables	$\lambda = 0.01$, q_{tj} (1)	Cubic18, q_{tj} (2)	Yield-based, q_{tj} (3)	Ask-based, q_{tj} (4)	Maturity $\geq$ 1m, q_{tj} (5)	Previous day, q_{tj} (6)	Svensson, q_{tj} (7)
Cheapness	0.012***	0.015***	0.076***	0.012***	0.011***	0.011***	0.009***
	(4.174)	(4.343)	(4.742)	(4.176)	(4.029)	(3.830)	(4.025)
VOL	−2.137**	−2.004*	−2.262**	−2.148**	−2.130**	−2.087**	−2.771***
	(−2.135)	(−1.955)	(−2.287)	(−2.146)	(−2.124)	(−2.105)	(−2.862)
Specialness	0.412	0.441	0.429	0.412	0.414	0.414	0.437
	(0.688)	(0.725)	(0.725)	(0.687)	(0.689)	(0.698)	(0.724)
OB	0.039	0.018	0.056	0.039	0.041	0.044	0.046
	(0.810)	(0.342)	(1.203)	(0.802)	(0.860)	(0.915)	(1.055)
Bid-Ask	−0.173***	−0.174***	−0.173***	−0.168***	−0.174***	−0.173***	−0.184***
	(−4.944)	(−4.990)	(−5.021)	(−4.861)	(−4.949)	(−4.931)	(−5.133)
N	3432	3432	3432	3432	3432	3432	3429
R^2	0.350	0.354	0.356	0.350	0.350	0.348	0.349
CUSIP fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Auction fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table A2
Alternative measures of cheapness: auction cost.

This table reports results of panel regressions of the auction price markup, for each purchased bond in each of the 139 quantitative easing auctions of nominal Treasury securities executed from November 12, 2010 to September 9, 2011. The auction price markup $Markup^c_{t,j}$ is computed as the average, weighted by the amount of each accepted offer, of the differences between the offer price and the corresponding closest ask of the bond. The explanatory variables are *Cheapness* (in basis points), *VOL* (in percentage points), *Specialness* (in percentage points), *OB* (in \$10 billions), and *Bid-Ask* (in basis points), as defined in Table 5. Different versions of cheapness are fitting the spline curve model with the roughness penalty coefficient $\lambda = 0.01$; fitting the spline curve model with knot points of 1, 2, 3, 5, 7, 10, 20, and 30 years ("Cubic18"); the difference between the yield to maturity of the market and model prices ("Yield-based"); fitting the curve using ask prices ("Ask based"); fitting the curve using the bonds with maturities greater than one month ("Maturity $\geq$ 1 month"); fitting the curve using the bond price quotes one day before the auction ("Previous day"); and fitting the curve using the Svensson (1994) model ("Svensson"). Robust *t*-statistics based on two-way clustered standard errors at the auction and CUSIP levels are reported in parentheses. Significance levels are: *** for $p < 0.01$, ** for $p < 0.05$, and * for $p < 0.1$, where p is the p-value.

Explanatory variables	$\lambda = 0.01$, $Markup^c_{tj}$ (1)	Cubic18, $Markup^c_{tj}$ (2)	Yield-based, $Markup^c_{tj}$ (3)	Ask-based, $Markup^c_{tj}$ (4)	Maturity $\geq$ 1m, $Markup^c_{tj}$ (5)	Previous day, $Markup^c_{tj}$ (6)	Svensson, $Markup^c_{tj}$ (7)
Cheapness	0.115***	0.147***	0.799***	0.114***	0.119***	0.114***	0.057**
	(2.971)	(3.770)	(4.105)	(2.964)	(3.263)	(2.843)	(2.008)
VOL	4.122	4.439	2.724	4.002	4.158	4.862	−0.138
	(0.676)	(0.742)	(0.451)	(0.660)	(0.684)	(0.778)	(−0.021)
Specialness	−8.641**	−8.383**	−8.286**	−8.648**	−8.490**	−8.868**	−8.283**
	(−2.204)	(−2.227)	(−2.114)	(−2.206)	(−2.183)	(−2.287)	(−2.155)
OB	−3.285***	−3.405***	−2.893***	−3.290***	−3.338***	−3.193***	−2.314***
	(−3.934)	(−4.427)	(−3.759)	(−3.928)	(−4.054)	(−3.986)	(−2.932)
Bid-Ask	−0.648***	−0.633***	−0.663***	−0.592***	−0.652***	−0.647***	−0.721***
	(−4.345)	(−4.179)	(−4.533)	(−4.036)	(−4.394)	(−4.361)	(−4.477)
N	1934	1934	1934	1934	1934	1934	1933
R^2	0.965	0.966	0.966	0.965	0.965	0.965	0.964
CUSIP fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Auction fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes

ters 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 et al. (2007), and Hu et al. (2013) for details on the interpretations of these parameters in terms of the yield curve.

These perturbations are included one at a time, while keeping the other model components identical to the baseline specification.

Tables A1 and A2 report results of our main regressions in Eqs. (9) and (11) using these seven alternative measures of bond cheapness. In Table A1, on quantity regressions, the coefficients on the six alternative price-based cheapness measures are between 0.009 (Svensson) and 0.015 (Cubic18), very similar to 0.012 in the baseline specification. The yield-based cheapness measure has a coefficient of 0.076, which is consistent. Because the proportional price difference is roughly the yield difference mul-

Table A3
Winsorize cheapness.
This table reports panel regressions of the auction purchase quantity q_{tj} for each included bond (in Columns 1 and 3) and panel regressions of the auction price markup for each purchased bond (in Columns 2 and 4) in each of the 139 quantitative easing auctions of nominal Treasury securities executed from November 12, 2010 to September 9, 2011. The auction price markup $Markup^c_{tj}$ is computed as the average, weighted by the amount of each accepted offer, of the differences between the offer price and the corresponding closest ask of the bond. The explanatory variables are *Cheapness* (in basis points), *VOL* (in percentage points), *Specialness* (in percentage points), *OB* (in $10 billions), and *Bid-Ask* (in basis points), as defined in Table 5. Columns 1 and 2 report regressions for the sample winsorized at the top and bottom 0.5% of cheapness, and Columns 3 and 4 for the sample winsorized at the top and bottom 1% of cheapness. Robust *t*-statistics based on two-way clustered standard errors at auction and CUSIP levels 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.

	Winsorize cheapness: 1%		Winsorize cheapness: 2%	
Explanatory variables	q_{tj} (1)	$Markup^c_{tj}$ (2)	q_{tj} (3)	$Markup^c_{tj}$ (4)
Cheapness	0.013***	0.117***	0.013***	0.120***
	(4.223)	(2.953)	(4.223)	(2.961)
VOL	−2.150**	4.045	−2.150**	3.941
	(−2.140)	(0.663)	(−2.140)	(0.649)
Specialness	0.418	−8.658**	0.418	−8.672**
	(0.705)	(−2.200)	(0.705)	(−2.199)
OB	0.054	−3.301***	0.054	−3.318***
	(1.217)	(-3.946)	(1.217)	(−3.939)
Bid-Ask	−0.173***	−0.650***	−0.173***	−0.651***
	(−4.936)	(−4.367)	(−4.936)	(−4.395)
N	3432	1934	3432	1934
R^2	0.351	0.965	0.351	0.965
CUSIP fixed effects	Yes	Yes	Yes	Yes
Auction fixed effects	Yes	Yes	Yes	Yes

Table A4
Treasury notes versus bonds.
This table reports results of panel regressions of the purchase quantity q_{tj} and panel regressions of the auction price markup, for each purchased bond in each of the 139 quantitative easing auctions of nominal Treasury securities executed from November 12, 2010 to September 9, 2011. Regressions are run on two subsamples split by time to maturity or original maturity, with the cutoff point of ten years. The auction price markup $Markup^c_{tj}$ is computed as the average, weighted by the amount of each accepted offer, of the differences between the offer price and the corresponding closest ask of the bond. The explanatory variables are *Cheapness* (in basis points), *VOL* (in percentage points), *Specialness* (in percentage points), *OB* (in $10 billions), and *Bid-Ask* (in basis points), as defined in Table 5. Robust *t*-statistics based on two-way clustered standard errors at auction and CUSIP levels are reported in parentheses. Significance levels are: *** for $p < 0.01$, ** for $p < 0.05$, and * for $p < 0.1$, where p is the p-value.

	Time to maturity				Original maturity			
	> 10 years		≤ 10 years		> 10 years		≤ 10 years	
Variable	q_{tj}	$Markup^c_{tj}$	q_{tj}	$Markup^c_{tj}$	q_{tj}	$Markup^c_{tj}$	q_{tj}	$Markup^c_{tj}$
Cheapness	0.017***	0.128**	0.002	0.094**	0.024***	0.220***	0.002**	0.074*
	(4.675)	(2.484)	(1.172)	(2.431)	(5.068)	(5.441)	(2.042)	(1.769)
VOL	−5.565***	−6.608	−0.052	14.024	−7.444***	−0.428	0.046	8.316
	(−3.045)	(−1.035)	(−0.373)	(1.582)	(−3.350)	(−0.060)	(0.323)	(0.876)
Specialness	0.351	−8.548**	0.503	10.022	0.052	−3.772	0.191	−11.466
	(0.535)	(−2.164)	(0.799)	(1.044)	(0.064)	(−1.535)	(0.875)	(−1.021)
OB	0.033	28.991	−0.024	−2.577***	0.009	−211.159	−0.036	−2.950***
	(0.501)	(1.051)	(−0.500)	(-5.875)	(0.124)	(−0.879)	(−0.724)	(−3.432)
Bid-Ask	−0.234***	−0.585***	−0.027	−1.107***	−0.292***	−0.637***	−0.027	−0.870***
	(−5.247)	(−3.161)	(−1.204)	(−3.669)	(−5.642)	(−4.511)	(−1.472)	(−2.905)
N	2972	1632	460	302	2575	1387	857	547
R^2	0.362	0.956	0.533	0.980	0.375	0.970	0.561	0.974
CUSIP fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Auction fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

tiplied by the bond duration, we expect the coefficient on the yield-based cheapness measure to be larger by a factor of about five or six (the average maturity of bonds in QE auctions), relative to the coefficients on price-based cheapness measures.

In Table A2, on auction cost regressions, the coefficients on the six alternative price-based cheapness measures are between 0.057 (Svensson) and 0.119 (maturity ≥ 1 month), similar to 0.115 in the baseline specification. The fact that the Svensson-based measure has a visibly smaller coefficient suggests that the spline-based model is probably closer to the Fed's own model. Again, the yield-based cheapness measure has a coefficient of 0.799, consistent with the multiplication by duration.

Table A3 reports the results from regressions in Eqs. (9) and (11), but with cheapness winsorized at the top and bottom 0.5% (Columns 1 and 2) and winsorized at the top and bottom 1% (Columns 3 and 4). The results are very similar to and even slightly stronger than the baseline regressions.

Table A4 reports the quantity and cost regressions by splitting the sample of Treasury securities into two halves, first by time to maturity and then by original maturity. The cutoff maturity is ten years. The positive coefficients of purchase quantity and cost on cheapness are robust for the subsamples, although the point estimates for longer-maturity bonds are larger than those for shorter-maturity ones.

References

Ashenfelter, O., 1989. How auctions work for wine and art. J. Econ. Perspect. 3, 23–36.
Ausubel, L., Cramton, P., Pycia, M., Rostek, M., Weretka, M., 2014. Demand reduction and inefficiency in multiunit auctions. Rev. Econ. Stud. 81, 1366–1400.
Back, K., Zender, J., 1993. Auctions of divisible goods: on the rationale for the treasury experiment. Rev. Financ. Stud. 6, 733–764.
Bauer, M., Rudebusch, G., 2014. The signaling channel for federal reserve bond purchases. Int. J. Central Bank. 10, 233–289.
Bikhchandani, S., Huang, C., 1989. Auctions with resale markets: a model of treasury bill auctions. Rev. Financ. Stud. 2, 311–340.
Bikhchandani, S., Huang, C., 1993. The economics of treasury securities markets. J. Econ. Perspect. 7, 117–134.
Boyarchenko, N., Lucca, D.O., Veldkamp, L., 2015. Taking Orders and Taking Notes: Dealer Information Sharing in Financial Markets. Staff report. Federal Reserve Bank of New York.
Breedon, F., Turner, P., 2016. On the Transactions Costs of Quantitative Easing. Bank for International Settlements, Basel, Switzerland. Unpublished working paper.
Cammack, E., 1991. Evidence on bidding strategies and the information in treasury bill auctions. J. Political Econ. 99, 100–130.
Chatterjee, A., Jarrow, R., 1998. Market manipulation, price bubbles, and a model of the US treasury securities auction market. J. Financ. Quant. Anal. 33, 255–290.
Christensen, J., Rudebusch, G., 2012. The response of interest rates to US and UK quantitative easing. Econ. J. 122, 385–414.
D'Amico, S., English, W., Lopez-Salido, D., Nelson, E., 2012. The federal reserve's large-scale asset purchase programs: rationale and effects. Econ. J. 122, 415–446.
D'Amico, S., Fan, R., Kitsul, Y., 2013. The Scarcity Value of Treasury Collateral: Repo Market Effects of Security-specific Supply and Demand Factors. Federal Reserve Bank of Chicago, IL. Unpublished working paper.
D'Amico, S., King, T., 2013. Flow and stock effects of large-scale treasury purchases: evidence on the importance of local supply. J. Financ. Econ. 108, 425–448.
Di-Maggio, M., Kermani, A., Song, Z., 2017. The value of trading relationships in turbulent times. J. Financ. Econ. 124, 266–284.
Duarte, J., Longstaff, F., Yu, F., 2007. Risk and return in fixed income arbitrage: nickels in front of a steamroller? Rev. Financ. Stud. 20, 769–811.
Duffie, D., 1996. Special repo rates. J. Financ. 51, 493–526.
Engle, R., Fleming, M., Ghysels, E., Nguyen, G., 2012. Liquidity, Volatility, and Flights to Safety in the US Treasury Market: Evidence from a New Class of Dynamic Order Book Models. Staff report. Federal Reserve Bank of New York.
Eser, F., Schwaab, B., 2016. Evaluating the impact of unconventional monetary policy measures: empirical evidence from the ECB's securities markets programme. J. Financ. Econ. 119, 147–167.
Fisher, M., Nychka, D., Zervos, D., 1995. Fitting the Term Structure of Interest Rates with Smoothing Splines. Federal Reserve System, Washington, DC. Unpublished working paper.
Fleming, M., Liu, W., 2014. Intraday Pricing and Liquidity Effects of US Treasury Auctions. Staff report. Federal Reserve Bank of New York.
Fleming, M., Mizrach, B., 2009. The Microstructure of a US Treasury ECN: The Brokertec Platform. Staff report. Federal Reserve Bank of New York.
Gagnon, J., Raskin, M., Remanche, J., Sack, B., 2011. The financial market effects of the Federal Reserve's large-scale asset purchases. Int. J. Central Bank. 7, 3–44.
Goldreich, D., 2007. Underpricing in discriminatory and uniform-price treasury auctions. J. Financ. Quant. Anal. 42, 443–466.
Gordy, M., 1999. Hedging winner's curse with multiple bids: evidence from the Portuguese treasury bill auction. Rev. Econ. Stat. 81, 448–465.
Goswami, G., Noe, T., Rebello, M., 1996. Collusion in uniform-price auctions: experimental evidence and implications for treasury auctions. Rev. Financ. Stud. 9, 757–785.
Gurkaynak, R., Sack, B., Wright, J., 2007. The US treasury yield curve: 1961 to the present. J. Monet. Econ. 54, 2291–2304.
Hamilton, J.D.D., Wu, J.C., 2012. The effectiveness of alternative monetary policy tools in a zero lower bound environment. J. Money Credit Bank. 44, 3–46.
Han, B., Longstaff, F., Merrill, C., 2007. The US treasury buyback auctions: the cost of retiring illiquid bonds. J. Financ. 62, 2673–2693.
Hancock, D., Passmore, W., 2011. Did the Federal Reserve's MBS purchase program lower mortgage rates? J. Monet. Econ. 58, 498–514.
Hasbrouck, J., 2007. Empirical Market Microstructure. Oxford University Press, New York.
Hollifield, B., Neklyudov, A., Spatt, C., 2017. Bid-ask spreads, trading networks, and the pricing of securitizations. Rev. Financ. Stud. 30, 3048–3085.
Hortacsu, A., Kastl, J., 2012. Informational advantage: a study of Canadian treasury auctions. Econometrica 80, 2511–2542.
Hortacsu, A., Kastl, J., Zhang, A., 2017. Bid Shading and Bidder Surplus in the U.S. Treasury Auction System. National Bureau of Economic Research, Cambridge, MA. Unpublished working paper.
Hortacsu, A., McAdams, D., 2010. Mechanism choice and strategic bidding in divisible good auctions: an empirical analysis of the Turkish treasury auction market. J. Political Econ. 118, 833–865.
Hu, X., Pan, J., Wang, J., 2013. Noise as information for illiquidity. J. Financ. 68, 2341–2382.
Jordan, B., Jordan, S., 1997. Special repo rates: an empirical analysis. J. Financ. 52, 2051–2072.
Kastl, J., 2011. Discrete bids and empirical inference in divisible good auctions. Rev. Econ. Stud. 78, 978–1014.
Keloharju, M., Nyborg, K., Rydqvist, K., 2005. Strategic behavior and underpricing in uniform price auctions: evidence from Finnish treasury auctions. J. Financ. 60, 1865–1902.
Kremer, I., Nyborg, K., 2004. Underpricing and market power in uniform-price auctions. Rev. Financ. Stud. 17, 849–877.
Krishnamurthy, A., 2002. The bond/old-bond spread. J. Financ. Econ. 66, 463–506.
Krishnamurthy, A., Vissing-Jorgensen, A., 2011. The Effects of Quantitative Easing on Interest Rates. National Bureau of Economic Research, Cambridge, MA. Unpublished working paper.
Li, C., Wei, M., 2013. Term structure modeling with supply factors and the Federal Reserve's large-scale asset purchase programs. Int. J. Central Bank. 9, 3–39.
Li, D., Schurhoff, N., 2014. Dealer Networks. Swiss Finance Institute, Geneva, Switzerland. Unpublished working paper.
Lou, D., Yan, H., Zhang, J., 2013. Anticipated and repeated shocks in liquid markets. Rev. Financ. Stud. 18, 981–1020.
McCulloch, H.J., 1971. Measuring the term structure of interest rates. J. Bus. 44, 19–31.
McCulloch, H.J., 1975. The tax-adjusted yield curve. J. Financ. 30, 811–830.
Meaning, J., Zhu, F., 2011. The Impact of Recent Central Bank Asset Purchase Programmes. Bank for International Settlements, Basel, Switzerland.
Milgrom, P., Weber, R., 1982. A theory of auctions and competitive bidding. Econometrica 50, 1089–1122.
Nyborg, K., Rydqvist, K., Sundaresan, S., 2002. Bidder behavior in multiunit auctions: evidence from Swedish treasury auctions. J. Political Econ. 110, 394–424.
Nyborg, K., Sundaresan, S., 1996. Discriminatory versus uniform treasury auctions: evidence from when-issued transactions. J. Financ. Econ. 42, 63–104.
Pasquariello, P., Roush, J., Vega, C., 2014. Government Intervention and Strategic Trading in the U.S. Treasury Market. University of Michigan, MI Federal Reserve System, Washington, DC. Unpublished working paper.
Petersen, M., 2009. Estimating standard errors in finance panel data sets: comparing approaches. Rev. Financ. Stud. 22, 435–480.
Piazzesi, M., 2010. Affine term structure models. In: Ait-Sahalia, Y., Hansen, L.P. (Eds.), Handbook of Financial Econometrics: Tools and Techniques. Elsevier, Amsterdam, pp. 691–766.

Potter, S., 2013. The implementation of current asset purchases. In: Remarks at the Annual Meeting with Primary Dealers. Federal Reserve Bank of New York.
Sack, B., 2011. Implementing the federal reserve's asset purchase program. In: Remarks at Global Interdependence Center Central Banking Series Event. Federal Reserve Bank of Philadelphia, PA.
Schlepper, K., Hofer, H., Riordan, R., Schrimpf, A., 2017. Scarcity Effects of QE: A Transaction-level Analysis in the Bund Market. Bank for International Settlements, Basel, Switzerland. Unpublished working paper.
Simon, D., 1994. Markups, quantity risk, and bidding strategies at treasury coupon auctions. J. Financ. Econ. 35, 43–62.
Stroebel, J., Taylor, J., 2012. Estimated impact of the Federal Reserve's mortgage-backed securities purchase program. Int. J. Central Bank. 8, 1–42.
Svensson, L., 1994. Estimating and Interpreting Forward Rates: Sweden 1992–1994. National Bureau of Economic Research, Cambridge, MA. Working paper.
Swanson, E.T., 2011. Let's Twist Again: A High-frequency Event-study Analysis of Operation Twist and its Implications for QE2. Federal Reserve Bank of San Francisco, CA. Unpublished working paper.
Umlauf, S., 1993. An empirical study of the Mexican treasury bill auction. J. Financ. Econ. 33, 313–340.
Vayanos, D., Weill, P.O., 2008. A search-based theory of the on-the-run phenomenon. J. Financ. 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? Econ. J. 122, 447–466.