Central bank, government interventions and financial markets

by

Yusuke Tsujimoto

A thesis submitted in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

in

Finance

Faculty of Business

University of Alberta

Abstract

This thesis comprises three essays examining the interactions between central bank and government interventions and financial markets. In Chapter 2, I compare corporate bond purchase programs by the Federal Reserve and the Bank of Japan during the COVID-19 crisis. Both programs shared one eligibility criterion: remaining maturity of five years or less. I demonstrate that Japanese firms, but not U.S. firms, *catered* to such a sudden maturity-specific demand shock by shifting the maturity of new bond issues, presumably due to the much larger purchase size of the Bank of Japan. Theoretically, the Japanese result aligns with the gap-filling theory of Greenwood et al. (2010), which predicts that corporations, as bond suppliers, fill supply-demand imbalances in specific maturity segments of the bond market.

Chapter 3 examines the microstructure of the reverse auctions the Fed holds for QE-driven purchases of Treasury bonds. Although the minimum tick size is set to be 1/256th (i.e., 0.390625 cents) per $100 par value, I document that primary dealers—the only direct participants—submit coarsely priced offers. Importantly, primary dealers with larger market shares price more finely, and my empirical analysis suggests that this coarse pricing originates from the information processing costs associated with increased pricing precision, in line with the theory of Grossman et al. (1997). I also document that the coarseness of prices is related to the *level* of prices (among *accepted* offers). Topmost dealers therefore play a special role in advancing market efficiency and promoting price competition.

In Chapter 4, we study how firms responded to the creation of the MSCI Empowering Women Index (WIN), an index for Japanese firms with superior gender diversity that Morgan Stanley

launched in cooperation with the Government Pension Investment Fund of Japan in 2017. Importantly, this index includes only firms in each industry's top 50% in women's workforce participation. This allows us to implement a difference-in-differences analysis based on firms around the inclusion threshold (treated firms) and those farther away from the threshold (control firms). We show that the treated firms improved gender diversity. This paper thus illustrates a capital market channel for inducing corporate social behavior changes.

Preface

Chapters 2 and 3 are based solely on my own work. Chapter 4 is a joint work with Professor Vikas Mehrotra at the University of Alberta, Professor Lukas Roth at the University of Alberta, and Professor Yupana Wiwattanakantang at the National University of Singapore. Professors Mehrotra, Roth and Wiwattanakantang identified the initial research idea and contributed to most of the writing. I contributed to the development and execution phases of the paper, including participating in framing the research question and methodology. In particular, I was primarily responsible for data curation, collaborated on the development and execution of empirical analysis, and made contributions to the writing and revisions as needed.

Acknowledgements

I would like to express my profound gratitude to Professor Vikas Mehrotra and Professor Randall Morck, my co-supervisors, for their caring, encouraging, and insightful guidance throughout my Ph.D. They have generously provided me with invaluable opportunities to grow academically and professionally, and I truly appreciate the trust they placed in me. My aspiration is to become a scholar who can pass on these virtues to my students.

I would like to thank Professor Mark Huson, my supervisory committee member, for his sincere and helpful advice. I have learned what it takes to write good papers and, more generally, to become a good academician. I also thank Professor Masahiro Watanabe, the Ph.D. in Finance program coordinator and the defense chair, for his guidance and support. I appreciate the entire Finance faculty of the Alberta School of Business for their valuable feedback. I would also like to note my thanks to Professor Lukas Roth and Professor Yupana Wiwattanakantang, my co-authors of Chapter 4, and to my arm's length examiners, Professor Sebastian Fossati and Professor Takeo Hoshi.

I am also grateful to my friends both within and outside the Alberta School of Business's Finance Ph.D. community. Despite my deep gratitude for their kindness, I choose not to take on the seemingly impossible task of listing them all here. I would like to acknowledge my long-time Finance Ph.D. colleagues, Rodrigo Cardenas, Joseph DeCoste, and Xiaowen Zhang, for making my Ph.D. life in Edmonton an enjoyable and positive one.

Lastly, but most importantly, I would like to thank my parents, Satoru Tsujimoto and Yumiko Tsujimoto. It was their unconditional love and support that made this pursuit of mine possible.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

Many "unconventional" monetary and financial policies have become common policy tools in developed economies (Kuttner, 2018). One representative example is central bank purchases of long-term government securities, commonly referred to as quantitative easing (QE). While the Bank of Japan's inception in 2001 led to widespread controversy, major central banks, most notably the Federal Reserve and the European Central Bank, eventually followed suit in the aftermath of the global financial crisis. Furthermore, confronted with the prolonged slow economic growth in the subsequent period, the Bank of Japan and the European Central Bank relied on a more aggressive policy—purchasing corporate bonds. This action is possibly more directly effective in lowering firms' cost of capital and thereby promoting investment. However, it comes with greater risk, as the central bank now takes on not just interest rate risk but also credit risk. Then, the COVID-19 crisis was a watershed moment. To navigate the economy through this turbulent time, policymakers worldwide relied on unconventional monetary policies more than ever. It was during this period that the Fed, with the backing of the Treasury Department, launched its first-ever corporate bond purchase programs.

Moreover, the *scope* of financial and monetary policy has been expanding. Policymakers worldwide are now more open to discussions on employing policy tools to address not only traditional goals (e.g., economic growth and financial stability) but also emerging broader objectives. A key driver of this trend is climate risk, which appears to be increasingly priced in (see Giglio et al. (2021) for a review). However, the debate can incorporate wider environmental, social, and governance (ESG) agendas. Governments' traditional approaches to tackling those issues are taxes, subsidies, and regulations. Nevertheless, governments can also induce changes in firm behavior through a *capital market* channel, by mobilizing funds under its control, such as sovereign wealth funds and public pension funds. More specifically, they can design a security purchase program that incentivizes private agents to make improvements toward those goals.

The design and evaluation of these new interventions require a careful, thorough examination of the responses from private agents in financial markets. On one hand, financial markets are dynamic and can quickly adapt to new environments through agents' competition to identify and exploit profit opportunities (Lo, 2019). On the other hand, history indicates that markets are subject to behavioral factors (Kindleberger, 1978; Shiller, 2000; Gennaioli and Shleifer, 2018), and even from a theoretical standpoint, market efficiency is bound to be limited if one acknowledges that information processing is costly (Grossman and Stiglitz, 1980).

At the same time, these new policy experiments provide empiricists with novel and fertile grounds for testing finance theories (Brunnermeier et al., 2021). For example, a large-scale asset purchase program by a central bank is essentially a demand shock to the target securities from the issuers' perspective. Finance researchers can rarely observe such a massive, sudden demand/supply shock to a specific segment of the market. This thesis, therefore, concerns the interaction between these new forms of central bank and government interventions and financial markets.

In Chapter 2, I study a situation in which a central bank becomes a significant buyer in a specific segment of the corporate bond market. During the COVID-19 pandemic, the Federal Reserve and the Bank of Japan announced massive purchases of corporate bonds maturing in five years or less. I demonstrate that firms in Japan, but not in the U.S., have catered to the maturity-specific demand shock by shifting the maturity of new bond issues, plausibly due to the much larger size of the Bank of Japan's actual purchases. I argue that the Japanese result is consistent with the "gap-filling" theory of Greenwood et al. (2010), which predicts that firms face a trade-off between the benefit of adjusting maturities to match the positive demand shock and the cost of deviating from their intrinsically optimal maturities. Thus, this essay has important policy implications for central banks becoming significant buyers in non-government debt markets.

Chapter 3 examines the reverse auctions that the Fed holds for implementing its QE-based Treasury bond purchases and studies how the Federal Reserve Bank of New York's primary dealers—the only direct participants—behave in this market. I document primary dealers' practice of coarse pricing; although the Fed explicitly sets the tick size of 1/256th in these reverse auctions, primary dealers' offer prices exhibit strong clustering on coarser grids. This coarse pricing has decreased over time but surged temporarily during March 2020. It grows when the offer-to-cover ratio is low, indicating a competitive constraint. Importantly, primary dealers with larger market shares price more finely, and laggard dealers employ coarse pricing especially when precise pricing is difficult. With regard to the Fed's purchase costs, coarsely priced offers have higher prices, conditional on winning the auction. However, I show that they likely have lower chances of winning as well. The results are consistent with the notion that dealers' coarse pricing results from the information cost associated with pricing precision, as proposed by Grossman et al. (1997), but not with dealer collusion. I thus illustrate that topmost dealers uniquely advance market efficiency in these multi-trillion-dollar operations of the Fed.

In Chapter 4, we study a novel financial market-based attempt by a government institution to influence firms toward the direction of its ESG targets. In 2017, the Government Pension Investment Fund of Japan, the world's largest pension fund, adopted the MSCI Empowering Women Index (WIN). To qualify for the prestigious index, firms must meet certain criteria for the advancement of women in their workforce. Inclusion in the WIN is structured loosely as a tournament—only the top 50% firms in each sector are included. Our analysis, using a difference-in-differences methodology, shows that firms around the inclusion threshold competing for index membership display significant improvements in women's participation in the workforce, including the C-suite, compared to firms farther away from the threshold. WIN firms also display an increase in paternity leaves suggesting a shift to a more women-friendly corporate culture. We also find that WIN firms gain institutional ownership and the change in corporate social behavior is not at the expense of operating profitability or valuation. Overall, our results document the social power of this novel policy tool, index creations.

Chapter 2

Do Firms Cater to Corporate QE? Evidence from the Bank of Japan's Massive Corporate Bond Purchases

"*[W]e remain watchful for the risk of creating distortions. One possibility could be that companies issue shorter maturities than they would otherwise to benefit from presence of the CCF [Corporate Credit Facility] backstop. So far, we see no evidence of this behavior.*"

—Daleep Singh, Speech at the Federal Reserve Bank of New York, 2020[1]

2.1 Introduction

In response to the novel coronavirus (COVID-19) pandemic, many central banks announced large-scale purchases of corporate bonds. An important feature of many of these programs is an eligibility criterion on (remaining) maturities. In the case of the U.S., with the backing of the Treasury Department, the Federal Reserve launched two facilities to purchase corporate bonds: the Primary

[1] Singh, Daleep (October 20, 2020). The Federal Reserve's Corporate Credit Facilities: Why, How, and For Whom. Speech, Federal Reserve Bank of New York. Retrieved from https://www.newyorkfed.org/newsevents/speeches/2020/sin201020.

Market Corporate Credit Facility (PMCCF) for primary market purchases and the Secondary Market Corporate Credit Facility (SMCCF) for secondary market purchases. The PMCCF targeted bonds maturing in four years or less, and individual bonds purchased by the SMCCF were required to have remaining maturities of five years or less. Similarly, the Bank of Japan (BOJ), which since the global financial crisis has purchased already-issued corporate bonds, dramatically scaled up purchase sizes, at the same time increasing the maximum eligible remaining maturity from three years to five years.

Under market segmentation, such a maturity-specific demand shock can influence the maturity choices of bond suppliers; firms may tilt the maturity of new bond issues toward the maturity segment experiencing such a demand shock to take advantage of favorable issuing conditions (Greenwood et al., 2010). I demonstrate that although this did not occur in the U.S., in line with the findings of Halling et al. (2020) and Boyarchenko et al. (2022) and the quoted speech of Daleep Singh (then Head of the Markets Group at the New York Fed), it did manifest in Japan. Figure 2.1 illustrates this stark difference between the two countries in the dynamics of the maturities of newly issued corporate bonds. The BOJ announced the most dramatic expansion of its corporate bond purchase program (hereafter, CBPP) on April 27, 2020. Figure 2.1 shows that in Japan the cumulative proportion of one- to five-year maturity bonds started to increase in May 2020 and reached a peak (62%) in July 2020, and at the same time, the proportion of maturities of (5,7] years, that is, a maturity category *slightly exceeding* the eligibility threshold, became almost negligible after April 2020. In contrast, no such changes occurred in the U.S., although the Fed announced the $750 billion corporate bond purchase plan on April 9 and the SMCCF started individual corporate bond purchases on June 16.

These Japanese results align with the notion that bond issuers face a trade-off when responding to demand shocks. By issuing a bond maturing in five years or less, they can benefit from favorable issuance conditions stemming from the positive demand shock. However, it is also costly for firms to deviate from their intrinsically optimal debt maturities, e.g., due to the maturity-matching principle (Myers, 1977), and this cost should increase as the deviation increases. This trade-off predicts that bond issuers whose target maturities are greater than, *but not too far from*, five years should respond the most to the positive demand shock by shortening the maturities of new bond issues. Consequently, there should be an increase in eligible maturities and a *disproportionate decrease* in maturities just to the right of the eligibility threshold.

To formally analyze the maturity decisions of Japanese firms, I employ the multinomial logit model. The dependent variable indicates the maturity bin of [1,3], (3,5], (5,7], (7,10], or >10 years. These maturity buckets are selected because original maturities of Japanese corporate bonds cluster on three, five, seven, and ten years. The analysis reveals that bond issuers significantly became less likely to choose the maturity bin just above the threshold, (5,7] years, as compared to *either* the neighboring shorter *or* longer maturity bin. The reduction is sizable; the baseline specification implies that the probability of selecting maturities of (5,7] years dropped from 17.9% to 6.8% (i.e., a 62% reduction).

Why are the U.S. and Japanese results so different? The most plausible reason is that the BOJ became a much more significant buyer in the target maturity segment than the Fed did. To begin with, the BOJ's purchase capacity was larger than the Fed's after taking into account the difference in the domestic corporate bond market size. More importantly, the *actual* purchase size of the Fed was much smaller than the initially announced purchase capacity; although the Fed allocated $750 billion to the two corporate bond purchase facilities in total, the PMCCF, while it was operative, did not purchase any bonds, and the SMCCF's actual purchases amounted to only $14 billion when the two facilities were concluded at the end of 2020. In contrast, the BOJ purchased ¥3 trillion (=$27.6 billion) of corporate bonds from March 23 to September 30, 2020. This is striking given that the Japanese corporate bond market size is only one-fifteenth that of the U.S. (Section 2.2.2). Furthermore, the BOJ's purchases were strictly confined to individual eligible bonds, i.e., investment-grade bonds maturing in from one to five years. In contrast, the SMCCF was authorized to buy individual bonds and bond ETFs. The latter option enabled the SMCCF to buy corporate bonds whose remaining maturities were above five years. This difference in purchase aggressiveness likely reflected different policy objectives.[2]

[2] U.S. corporate bond market liquidity dried up in the wake of the COVID-19 crisis (O'Hara and Zhou, 2021), and the main aim of the Fed's corporate bond market interventions was to address the dislocations in the market by providing *promises* of the market support (Boyarchenko et al., 2021; Haddad et al., 2023). Put differently, there was

Theoretically, firms with greater financial flexibility should be more responsive to a maturity-specific demand/supply shock thanks to their lower costs of deviating from the target maturity structure (Greenwood et al., 2010). This paper's analysis, however, does not reveal strong heterogeneity. There are three possible explanations. First and most importantly, as explained later, the protocol of the BOJ's reverse auctions implies that *higher-yield bonds*, and therefore bonds issued by *riskier firms*, were purchased first by the BOJ. That is to say, the BOJ-driven positive bond demand shock was directed more toward riskier firms. Second, the primary corporate bond market in Japan is practically restricted to only highly creditworthy issuers, and as a result, the cross-sectional variation is limited. For instance, *all* of the Japanese bonds included in my sample are investment grade and 95% of them are rated either AA or A. Lastly, firms might have different incentives during an acute crisis such as the COVID-19 crisis. Specifically, firms might care more about raising and stockpiling cash as opposed to keeping their maturity structures close to the target ones, and this might have been more so for firms with weaker balance sheets.

I focus my attention on changes in the maturities of new corporate bond issues in the vicinity of the eligibility threshold, and this is because sharp theoretical predictions can be made in this region.[3] The main concern in interpreting those changes as evidence of a *causal* effect of the BOJ's CBPP expansion is the possibility of confounding shocks influencing firms' maturity choices. However, it seems unlikely that other confounding factors, possibly related to COVID-19, completely brought about the substantial changes around the threshold. First, I am not aware of any other pandemic-related institutional arrangements that provided bond issuers with a *discontinuous* incentive shift around the maturity of five years.[4] Second, a simple shift toward shorter or longer maturities cannot explain the observed result—the decrease in maturities of (5,7] years is significant in comparison

no strong impetus for the Fed to proceed with further purchases once the market had recovered stability. On the other hand, the Japanese corporate bond market remained relatively stable even in early 2020, and the BOJ, like the ECB, has conducted corporate bond purchases as part of its quantitative easing policy (especially since the launch of the Quantitative and Qualitative Monetary Easing in 2013). The massive post-pandemic expansion of the BOJ's CBPP consequently had dual objectives: stabilizing the market and improving firms' funding conditions (Kuroda, 2020). The latter likely requires continual purchases.

[3] A possible alternative approach, a difference-in-differences design, is not chosen due to the absence of a credible control group; *almost all* bond issuers in Japan satisfied the CBPP's eligibility conditions throughout the sample period. Due to this methodological limitation, the possible effects of the BOJ's CBPP expansion on firm-level outcomes, while important, are beyond the scope of this paper.

[4] I rule out a possibility that shocks to *government* debt supply, which can affect firms' maturity choices to the extent that investors regard government and corporate bonds as substitutes, drove the result.

with *either* the shorter neighboring maturity bin ((3,5] years) *or* the longer bin ((7,10] years), consistent with my theoretical prediction.

The possibility of confounding factors is, nevertheless, still worrisome.[5] I therefore exploit the practice of firms issuing *multiple* bonds with *different* maturities at the same time (Helwege and Turner, 1999). These multiple-maturity bond issuances provide a cleaner testing ground for firms' catering behavior—I can analyze how individual firms select a *set* of different maturities, after controlling for the *range* of maturities they selected. Specifically, I examine one common type of multiple-maturity issuance—those including at least one maturity of five years or less and one maturity of ten years or more. This sample of issuers, by construction, must have had similar desired maturity *ranges*.[6] This paper shows that these issuances became much less likely to include a bond maturing in (5,7] years in the period following the BOJ's CBPP expansion. This result strongly bolsters the view that the maturity-specific demand shock disproportionately decreased the attractiveness of maturities that were just above the maturity eligibility threshold. It also implies that simultaneous issuances of multiple-maturity bonds were utilized by firms to cater to the BOJ-driven demand shock *without* much increasing their overall short-term debt reliance.

This paper concerns the effects of the BOJ's corporate bond purchases on the primary market, and therefore, the effects on the corporate bond secondary market are outside the main scope of this paper.[7] My main contribution is to document that a tight maturity eligibility criterion of central bank corporate bond purchases can lead to firms' gap-filling behavior. To the best of my knowledge, this unintended effect of central bank corporate bond purchases is new to the literature.[8] While

[5] One particular reason is a large decrease in very long maturities (>10 years) in the post-CBPP period (Figure 2.1). This result suggests that the pandemic, and ensuing uncertainty, did affect firms' maturity choices.

[6] The basic idea can be illustrated by the following example of SoftBank Corp. (9434). On March 12, 2020, the company issued four ¥10 billion-bonds maturing in 3, 5, 7, and 10 years. Then, on July 21, 2020, the company returned to the public debt market by issuing three new bonds maturing in 3, 5, and 10 years, which raised 10, 70, and 20 billion yen, respectively. Although both of the simultaneous issuances had identical maturity ranges and very similar average maturities, only the first issuance in the pre-crisis period included a seven-year bond.

[7] Gilchrist et al. (2020) show that the Fed's corporate bond purchase announcements disproportionately decreased secondary market yields of SMCCF-eligible bonds. In the case of the Japanese corporate bond market, Suganuma and Ueno (2018) study a period preceding the pandemic and find that the BOJ's corporate bond purchases decreased credit spreads.

[8] I am not aware of any academic studies examining changes in firms' maturity choices in response to the BOJ's CBPP. At the same time, I note that the possibility that the post-COVID-19 expansion of the BOJ's CBPP influenced Japanese firms' bond maturity choices has been suggested by the financial press. An article by Nikkei on October 16, 2020 reported that the issuance amount of corporate bonds maturing in six to seven years was much lower during

the Fed's corporate bond purchase programs in response to the COVID-19 crisis had tight maturity criteria, there is no evidence that they induced firms to shorten the maturity of newly issued bonds (Halling et al., 2020; Boyarchenko et al., 2020). The European Central Bank (ECB), another major central bank, has also conducted corporate bond purchases since the 2008–2009 financial crisis. In particular, the ECB initiated a large-scale corporate bond purchase program called the Corporate Sector Purchase Programme (CSPP) in 2016 and the Pandemic Emergency Purchase Programme (PEPP) in March 2020. However, the maturity eligibility criteria of the CSPP and PEPP are very lax: (remaining) maturities of less than 31 years. Therefore, this eligibility threshold itself should not be binding for most issuers (Mäkinen et al., 2020).

The debt maturity structure of industrial firms is relevant to policymakers. Firms shortening their debt maturities are also increasing their rollover risk and, consequently, their vulnerability to adverse capital supply shocks. Almeida et al. (2011) document that in the wake of the global financial crisis U.S. firms with greater rollover risk contracted investment more than other similar firms did. Similarly, Kalemli-Özcan et al. (2022) show that rollover risk deterred investment of European firms during the crisis. Moreover, these investment constraints attributable to rollover risk can dampen future growth in firms' productivity (Duval et al., 2020).

The BOJ's original motive for the strict maturity eligibility criterion was to limit its risk exposure (Bank of Japan, 2009), and in this sense, the distortionary effect on firms' bond maturity choice was *unintended*. It is, however, important to note that this does not mean that the effect was *unexpected*. Rather, the fact that the BOJ increased the maximum eligible maturity at the same time as increasing its purchase size indicates that the BOJ was well aware of this possible side effect. Therefore, one direct policy implication of this paper is that the BOJ's decision to relax the maturity eligibility criterion was appropriate—as otherwise, its massive purchases could have distorted firms' bond maturity choices toward the original shorter threshold (three years), which is likely less desirable during a time of crisis.

July–September 2020 relative to the same months in the previous year and explained that some firms shortened bond maturities to meet the BOJ's maturity criterion—Kohei Onishi (October 16, 2020) *Nichigin taishogai no shasai ni juyo toshika 6 nensai ni katsuro* (Demand for corporate bonds not targeted by the Bank of Japan increases—Investors moving to 6-year bonds). *Nikkei*.

Related literature: This paper draws from two strands of literature. First, it relates to the literature on central bank corporate bond purchase programs. Many papers have studied the effects of the ECB's CSPP. The CSPP's positive effects on firms' financing conditions are documented by Abidi and Miquel-Flores (2018), Grosse-Rueschkamp et al. (2019), Zaghini (2019), Pegoraro and Montagna (2022), Todorov (2020), Arce et al. (2021), and De Santis and Zaghini (2021). Furthermore, Pegoraro and Montagna (2022) show that to meet the eligibility criteria, CSPP-eligible issuers changed bond characteristics, such as exchange listing, use of central securities depository (CSD), seniority, collateral, and credit guarantee.[9] As aforementioned, the ECB's maturity eligibility criterion is quite lax—31 years or less. Pegoraro and Montagna (2022) and Galema and Lugo (2021) show that CSPP-eligible issuers substituted private debt with bonds and chose *longer* bond maturities. Those authors interpret that firms' motive to "lock in" the favorable but possibly transient issuing conditions drove this maturity shift. Demirgüç-Kunt et al. (2023) employ the event study approach to analyze the announcement effect of the ECB's PEPP on March 18, 2020.

The Fed's first-ever corporate bond purchases during the COVID-19 pandemic have also attracted much attention. First, many researchers investigate the effect of the Fed's interventions on the liquidity crisis in the corporate bond market in the wake of the pandemic (O'Hara and Zhou, 2021; Haddad et al., 2021; Nozawa and Qiu, 2021; Gilchrist et al., 2020; Kargar et al., 2021; Boyarchenko et al., 2022; Falato et al., 2021; Ma et al., 2022). Second, the impact of the Fed's corporate bond purchases on firms' financing decisions is studied by Acharya and Steffen (2020), Halling et al. (2020), Boyarchenko et al. (2022), Becker and Benmelech (2021), and Pettenuzzo et al. (2021). Notably, Halling et al. (2020, p. 503–504) find that the average maturity of newly issued corporate bonds became *longer* during the first months of the pandemic and state that this finding is "surprising" for two reasons: First, it is known that corporate bond maturities tend to be pro-cyclical (Erel et al., 2012; Chen et al., 2021). Second, the Fed's purchase programs mainly targeted relatively shorter-term bonds. A possible explanation offered by Halling et al. (2020) is that capital demand-side desire to reduce the immediate rollover risk might have played a special

[9]Notably, it is not just central banks that predominantly buy only a subset of corporate bonds satisfying certain conditions, and issuers have an incentive to cater to them as well. Dathan and Davydenko (2020) and Calomiris et al. (2022) show that firms modify bond features to cater to the demands from bond ETFs and major indexes, respectively.

role during the pandemic. Boyarchenko et al. (2022) also find no evidence that firms shortened the maturity of new bond issues due to the Fed's facilities. Lastly, Flanagan and Purnanandam (2020) analyze the target selection of the SMCCF's individual corporate bond purchases.

A couple of papers study the effects of the BOJ's CBPP using data preceding the COVID-19 crisis.[10] First, Suganuma and Ueno (2018) analyze the effects of the BOJ's government and corporate bond purchases on credit spreads using secondary market data up to 2016. They find that the BOJ's corporate bond purchases led to lower credit yields of eligible bonds. Second, Linh (2021) uses data up to 2018 and studies the effects of the BOJ's equity ETF purchases and corporate bond purchases on firms' capital structure decisions.[11] This paper differs from these two papers in terms of not only the BOJ's CBPP regime (i.e., pre-COVID-19 vs. post-COVID-19) but more importantly the research focus—I study the novel distortionary effect of the BOJ's massive purchases on firms' bond maturity choices.

Second, this paper extends the literature on the determinants of the maturity of corporate debt.[12] More specifically, this paper essentially tests the gap-filling theory of Greenwood et al. (2010). Finding an exogenous demand/supply shock to a specific maturity segment is difficult. Therefore, Greenwood et al. (2010), Badoer and James (2016), and Lugo and Piccillo (2019) rely on variations in *government bond supply*, based on the assumption that government and corporate bonds are (partial) substitutes for preferred-habitat investors (that is, investors predominantly investing in debt with specific maturities). On the other hand, I am aware of only two papers exploiting *corporate bond demand shocks*. First, Butler et al. (2023) document that the demand from insurance companies, which have a strong appetite for long-term bonds, influences firms' decision to supply long-term bonds. Second, Lugo (2021) focuses on money market mutual funds,

[10] The BOJ's equity ETF purchase program, which was also intended to boost firms' investment, is studied by Barbon and Gianinazzi (2019) and Charoenwong et al. (2021).

[11] Nevertheless, estimating the causal effects of the BOJ's CBPP on firm-level outcomes is challenging due to the lack of a clear control group. Linh (2021) defines the "treatment group" as firms that had outstanding bonds that were within the scope of the BOJ purchases (i.e., maturing in one to three years and not "fallen angels") and reports that these firms replaced bank debt with bonds more than other firms did. Such a difference-in-differences analysis seems to be unsuitable for my research setting because there is no a priori reason to expect that firms whose *existing* bonds are within the scope of the BOJ purchases have a stronger incentive to engage in the catering behavior.

[12] Classical explanations about firms' debt maturity choices are the so-called matching principle, i.e., matching the asset and liability maturities, (Myers, 1977), and signaling models under asymmetric information (Flannery, 1986; Diamond, 1991).

which are regulated to buy only very short-term (less than one year) bonds and shows that their demand affects firms' short-term debt issuances. I add to this literature by studying the effect of a large, well-confined maturity-specific bond demand shock in the medium range of the maturity spectrum. In addition, the demand shock in my study comes not from private investors but from a central bank. This paper thus informs policy about the influence of central banks becoming significant buyers in non-government debt markets.

2.2 Institutional background

2.2.1 The BOJ's corporate bond purchase program

When the COVID-19 pandemic emerged, the BOJ had already implemented corporate bond purchases as part of its long-lasting quantitative easing policy. Based on the terms announced in 2013, the bank was allowed to purchase up to ¥2.2 trillion of commercial paper (CP) and up to ¥3.2 trillion of corporate bonds. The key eligibility criteria for corporate bonds were credit ratings of BBB or above (i.e., investment grade) and residual maturities of one to three years. Also, eligible bonds must have been denominated in Japanese yen and governed by Japanese law. A brief history of the BOJ's corporate bond purchases is provided in Appendix 2.A.

During the COVID-19 pandemic, the BOJ's CBPP was expanded, both in terms of size and scope, through the following two key announcements: First, the BOJ announced on March 16, 2020 that the caps on its CP and corporate bond purchases would be increased by ¥1 trillion each (Bank of Japan, 2020b). Second, on April 27, 2020, the bank unveiled a plan to expand its corporate bond purchases to an unprecedented scale (Bank of Japan, 2020a). The size of the additional cap was increased from ¥1 trillion to ¥7.5 trillion for each of the CP and corporate bond purchases. As a result, the total cap on the BOJ's corporate bond purchases amounted to ¥10.7 trillion, an increase of 234% from the pre-pandemic level. There were also some modifications to the eligibility criteria. One of the most significant changes was the extension of the maximum eligible remaining maturity from three years to five years. Another noteworthy change was that the maximum bond purchase amount per issuer was increased from ¥100 billion to ¥300 billion. The April announcement said that the additional corporate debt purchases would continue until September 2020. Nevertheless, in May 2020, the date of termination was extended to March 2021 (Bank of Japan, 2020c). The purpose of the BOJ's CBPP during COVID-19 was twofold: to stabilize the financial market through liquidity provision and to make it easier for firms to raise capital (Kuroda, 2020). Table 2.1 summarizes major policies the BOJ announced on the two dates.[13]

The BOJ indeed accelerated corporate bond purchases following these announcements. The BOJ employed reverse auctions to purchase corporate bonds, and the auctions were conducted separately for bonds maturing in [1,3] and (3,5] years. Figure 2.2 shows that the monthly purchase amount of corporate bonds with remaining maturities of [1,3] years increased from ¥100 billion in February 2020 to ¥200 billion in March and further to ¥300 billion in the next month. Also, the BOJ began purchasing bonds maturing in three to five years in May and purchased ¥200 billion per month until the sample period ended.

Therefore, simply put, Figure 2.2 indicates that each of the maturity segments of [1,3] and (3,5] years experienced a ¥200 billion demand shock (on a monthly basis). This shock, however, could have led to a larger impact on the former segment. This is because the BOJ's previous purchases plausibly had already made the 'stock' of bonds with remaining maturities of [1,3] years scarce in the secondary market. Outcomes of the BOJ's reverse auctions, which are summarized in Table A2.1 in Appendix 2.B, align with this view. In the reverse auctions of bonds maturing in [1,3] years, the offer-to-cover ratio—the amount of (acceptable) offers divided by the purchase amount at the auction—became unusually low after the BOJ scaled up its purchases. For instance, in the May-8 auction, the first auction held after the BOJ's announcement on April 27, the offer-to-cover ratio was only marginally above one (1.05), with the lowest winning offer yield being −0.14%. This offer-to-cover ratio was much lower than those in the pre-COVID-19 period: in the auctions held on January 23 and February 20, 2020, the ratios were 3.15 and 2.54, respectively. In contrast, in

[13] Of course, measures taken by the BOJ were not limited to enhancing its corporate debt purchases. There were three other important actions: (a) loans to financial institutions through the Special Funds-Supplying Operations, (b) expanded Japanese government bond purchases, and (c) expanded equity ETF purchases. All these policies were also announced on March 16 and April 27, which were the dates of the bank's Policy Board meetings. For a summary of monetary and fiscal policy responses to the COVID-19 crisis in Japan, see Ogawa (2021).

the case of auctions of bonds maturing in (3,5] years, more offers were submitted relative to the purchase size; when it was held for the first time on May 20, 2020, the offer-to-cover ratio was 2.96. These facts indicate that for the BOJ demand shock to be 'absolved,' more bond issuances (or greater 'flows') were required in the maturity segment of [1,3] years than in that of (3,5] years.

2.2.2 Comparing the corporate bond purchase programs of the BOJ and the Fed

On March 23, 2020, the Fed announced the creation of its first-ever programs to purchase corporate bonds with financial support from the Treasury: the PMCCF to purchase newly issued corporate bonds and syndicated loans and the SMCCF to purchase already issued bonds and bond ETFs. The Fed also announced on April 9, 2020 that the combined purchase capacity of the PMCCF and the SMCCF was greatly expanded: from \$200 billion to \$750 billion. Importantly, these facilities' maturity eligibility criteria for individual bonds were close to that of the BOJ's CBPP. The SMCCF's criterion was five years or less remaining to maturity, whereas that of the PMCCF was slightly shorter—original maturities of four years or less.

Figure 2.3 compares both the purchase capacities and actual purchase amounts (up to September 2020) of the Fed and the BOJ. To gauge their economic magnitudes, I also deflate those amounts by their domestic corporate bond market sizes at the beginning of 2020, which are obtained from the Japan Securities Dealers Association for Japan[14] and SEC (2020) for the U.S.[15] The top graphs show that the BOJ's purchase cap was even larger as a proportion of the domestic corporate bond market size than the massive \$750 billion purchase limit of the Fed. The BOJ's value (15.4%) consists of the existing purchase cap (4.6%) and the additional one (10.8%), and the latter alone is larger than the value of the Fed (7.8%). A more striking difference is found in the comparison of the *actual* purchase amounts in Figure 2.3. While the BOJ's actual purchase amount relative to the domestic corporate bond market size was 4.3%, that of the Fed was merely 0.13%.

[14] I use the total outstanding amount of "corporate straight bonds" (*futsu shasai*) as a proxy for the domestic bond market size. The amount was ¥69.2 trillion (equivalent to \$637 billion) as of December 2019.

[15] To calculate the relative sizes of the Fed's purchase capacity and actual purchase amount, I rely on SEC (2020), which estimates the size of the outstanding corporate bonds in the U.S., issued by U.S. and foreign firms, to be \$9.1 trillion as of December 31, 2019.

2.3 Theoretical framework

The gap-filling theory of Greenwood et al. (2010) suggests that the BOJ's CBPP expansion could have encouraged firms to "fill the gap" in the targeted maturity segment to absorb the positive demand shock. The driver of the gap-filling theory is the existence of a group of investors who predominantly invest in debt with specific maturities, i.e., preferred-habitat investors (Modigliani and Sutch, 1966, 1967; Vayanos and Vila, 2021). If arbitrage capital is limited and if firms cannot change the maturity of new debt issues, a change in preferred-habitat investors' demand leads to a violation of the expectations hypothesis. In reality, however, firms can change the maturity of new debt issues. Greenwood et al. (2010) therefore posit that firms have an incentive to alter the debt maturity toward a maturity segment in which the debt supply is low relative to the demand from preferred-habitat investors and that firms do so as long as the benefit outweighs the cost.

The demand shock examined by this paper is confined to a specific maturity segment, and such a shock can lead to sharp changes in firms' maturity choices in the vicinity of the threshold. This is because firms face a trade-off in choosing bond maturity: by selecting a maturity that is highly demanded by (preferred-habitat) investors, a firm can possibly issue the bond at favorable terms, while doing so can be costly in the sense that the firm's maturity structure is moved away from the target one (Greenwood et al., 2010). Importantly, the cost for the firm is expected to increase as the size of the deviation from its target maturity increases. Thus, in this paper's setting, where demand for bonds maturing in five years or less suddenly increased, the following predictions can be made: the firms whose target maturities are five years or less keep choosing their target maturities, those whose target maturities are only moderately above five years might find it optimal to shorten the maturity of new bond issues (to five years), and the firms with target maturities well above five years are unlikely to shorten the maturity to meet the demand shock because it would be too costly for them to do so. As a result, if we compare the distributions of the maturities of newly issued bonds in the periods before and after the BOJ's CBPP expansion on April 27, 2020, there should be an increase in the share of maturities eligible to be bought by the BOJ and a *disproportionate* decrease in the share of maturities slightly exceeding five years.[16]

The above reasoning implies that all other things being equal, the increase in the issuance of eligible maturities should predominantly occur at the maturity that just meets the eligibility criterion, namely, five years. Nevertheless, this argument does not take into account that the BOJ held separate reverse auctions of bonds with remaining maturities of [1,3] and (3,5] years, and the demand shock to the shorter maturity segment was likely larger in effect due to the BOJ's preceding purchases (Section 2.2.1). Consequently, the demand shock looks more like two separate shocks that occurred in the two maturity segments, and each shock could have encouraged firms to shift the maturity toward the segment.

2.4 Data

The BOJ's corporate bond purchases did not explicitly exclude specific firms or sectors.[17] It is important to note the BOJ reserved the right to reject offers at its discretion. The BOJ has not disclosed the identities of either the purchased or the rejected corporate bonds. However, media reports have suggested that the BOJ has indeed excluded offers for certain (eligible) corporate bonds and the auction participants have been aware of that (Nikkei, 2021).

Japanese corporate bond issuance data were obtained from the Japan Securities Dealers Association (JSDA).[18] I retain bond code 40 ("corporate straight bonds" with fixed interest rates). This group includes callable bonds and they are excluded. I also removed permanent bonds, se-

[16]Notably, Calomiris et al. (2022) document discontinuous changes in issuance amounts, instead of maturities, around a threshold that they focus on and provide an explanation that somewhat resembles that of this paper. They analyze U.S. dollar-denominated corporate bonds issued by emerging country firms and document a disproportionate increase in clustering of issuance amounts exactly at $500 million in the period following the global financial crisis. Calomiris et al. (2022) argue that the driving force is a post-crisis increase in the demand for emerging country corporate bonds that can be included in major indexes. Importantly, a key index eligibility criterion of major bond indexes is the minimum issuance size of $500 million. Thus, according to the authors, firms whose target issuance sizes are below the threshold face a trade-off. By "stretching" to $500 million, they can issue bonds at more favorable conditions. At the same time, raising more than the target amount is costly because individual firms have limited investment opportunities whose (risk-adjusted) expected returns are at least as good as the interest rate. As a result, only emerging country issuers whose target sizes are moderately below $500 million would change their issuance amount to $500 million after the demand shock.

[17]The only exception is that the BOJ committed not to buy bonds issued by financial institutions holding current accounts at the BOJ and their parent holding companies (i.e., "counterpart financial institutions").

[18]`https://www.jsda.or.jp/en/statistics/bonds/index.html`

cured (asset-backed) bonds, equity-like bonds (e.g., bonds with an optional interest deferral clause), and investment corporation bonds.[19] The data include information on issuance amounts, coupon rates, yields, maturities, and credit ratings, among others. Maturities were rounded to the nearest half-year. Offering spreads were calculated by subtracting the yields of the closest-maturity Japanese government bonds (JGBs) from offering yields.[20] This bond issuance data were merged with firm-level data obtained from the Worldscope.

Sample issuers: I first exclude firms under government control such as Japan Tobacco Inc. (JT) and companies in the Japan Railway (JR) group as did Tanaka (2014). Then, my main analysis considers public companies in non-financial and non-utility industries. The inclusion of financial firms and utilities, however, does not qualitatively change the main result of this paper (Table A2.4 in Appendix 2.D).

Sample period: My sample period consists of three sub-periods. As the BOJ announced expansions of its CBPP on March 16 and April 27, 2020, I refer to the period between the two announcements as *expansion1*, and the period starting from the date of the second announcement *expansion2*. The sample period ends on September 30, 2020. One full year leading up to the first announcement (March 16, 2019–March 15, 2020) is called *pre-expansion*. The empirical results remain similar, however, even if the pre-expansion period includes only six months before March 16, 2020.

One salient feature of the Japanese corporate bond market is illustrated in Panel A of Table 2.2: the primary corporate bond market in Japan is practically limited to highly creditworthy issuers.[21] All of the sample bonds are investment grade, and the vast majority of them (94.8%) are rated A or

[19] The BOJ set no explicit restrictions on non-straight bonds. The bank, however, has discretion in selecting which eligible bonds to buy, and it *might* take special clauses attached to bonds into account. This issue remains unclear due to non-identification of bonds purchased by the BOJ.

[20] Data on JGB yields were obtained from the website of the Ministry of Finance: `https://www.mof.go.jp/english/jgbs/reference/interest_rate/index.htm`

[21] The Japanese corporate bond market is rated by four major credit rating agencies: two global institutions, Standard & Poor's (S&P) and Moody's, and two domestic ones, Rating and Investment Information (R&I) and Japan Credit Rating Agency (JCR). Most bond issuers obtain credit ratings only from the domestic agencies. In the interest of comparability, the highest ratings are employed for bonds rated by multiple agencies. This is because the global credit agencies tend to provide lower ratings (Packer, 2002; Han et al., 2012) and because bigger, and therefore likely safer, firms are apt to be rated by the global agencies in addition to the domestic ones. Consequently, the use of the lowest (or even average) ratings could deflate the creditworthiness of relatively safer issuers.

above.[22] Notably, this result is not driven by less creditworthy firms' greater difficulty in accessing the public debt market during the COVID-19 crisis. Bonds rated AA or A account for 93.3% even during *pre-expansion*. Panel B of Table 2.2 shows the distribution of maturities by credit rating.[23]

The summary statistics of basic bond- and firm-level variables are provided in Table 2.4. Of the 499 newly issued bonds, I was able to merge 490 with all of the firm-level variables from Worldscope that will be included in the subsequent multivariate analysis. Variable definitions are provided in Table 2.3. Bond- and issuer-characteristics in the three sub-periods are compared in Table 2.5. Most notably, it shows that the average log maturity is longer (shorter) in the expansion1 (expansion2) period than in the pre-expansion period. The average maturities, however, can be greatly influenced by long-term bond issues, and my main interest lies in changes in maturities in the neighborhood of the BOJ's eligibility threshold.

2.5 Results

Preliminary observations

Figure 2.4 shows the number of newly issued bonds and their total proceeds. It displays data from not only Japan but also the U.S. for reference. Interestingly, while in the U.S. corporate bond issuances surged in March 2020 and remained high in the following two months, a significant increase in corporate bond issuances in Japan occurred only in June and July 2020.

Two striking observations can be made from Figure 2.5, which displays the distributions of bond maturities for the three sub-periods. First, corporate bond maturities in Japan exhibit clustering on some integer years, namely, 3, 5, 7, and 10, across the sub-periods, suggesting that these maturities

[22] The first-ever speculative-grade (BB) bond was issued by Aiful in June 2019; however, this bond is not included in the sample as the issuer is a financial firm. For background information, see Finbarr Flynn and Komaki Ito (2019, May 28) Japan's First Junk Bond Would Break Barriers in Wary Market. *Bloomberg.* Retrieved from `https://www.bloomberg.com/news/articles/2019-05-28/japan-s-first-junk-bond-would-break-down-barriers-in-wary-market`.

[23] Although both AA- and A-rated bonds are dispersed across a wide maturity spectrum, BBB-rated bonds exhibit a high concentration (46%) in the maturity bin of (3,5] years. This observed pattern is largely consistent with the finding of Guedes and Opler (1996) in the U.S. that less creditworthy firms are more likely to issue medium-term debt.

conform to a norm (Weld et al., 2009). Second, the distributions seem to differ by sub-period. As the interpretation of the results during the expansion1 period requires caution for the reasons discussed later, the focus is on comparing the pre-expansion and expansion2 periods. One salient difference is the decrease in the share of the maturity of seven years, which was a previously popular maturity marginally exceeding the maximum (remaining) maturity of the BOJ purchases.

In contrast, there are no signs that U.S. issuers exhibited a similar catering behavior. First, Figure 2.1 shows that the cumulative proportion of one- to five-year maturity bonds remained constant at around 20% throughout the period examined. Second, it appears that the proportion of maturities of (5,7] years *increased* following the pandemic. This result holds true even for U.S. issues rated at least A, whose bonds were eligible for the SMCCF purchases (Figure A2.6 in Appendix 2.D). Collectively, these results are consistent with Boyarchenko et al. (2022), who conclude that the Fed's facilities did not induce U.S. firms to shorten the maturities of new bond issues.

In the following analysis of Japanese data, maturities are grouped into five maturity bins: (1,3] years, (3,5] years, (7,10] years, and greater than 10 years. The selection of these maturity buckets is motivated by the following factors: First, it is natural to have maturity bins each of which embraces only (at most) one of the popular maturity years, i.e., 3, 5, 7, and 10 years. Second, because the maximum maturity for the BOJ's purchases was originally three years and then increased to five years, the tests of firms' maturity choices around the maturity eligibility criteria require that these years are used as cut-off values.

Multinomial logit analysis of firms' bond maturity choices

To formally study Japanese firms' maturity choices of new bond issues, I adopt the multinomial logit model (MNLM), following previous studies on this topic (Guedes and Opler,1996; Badoer and James, 2016). To estimate changes in firms' bond maturity choices following the BOJ's CBPP expansions, my MNLM includes two indicator variables, $Expansion1$ and $Expansion2$, which take a value of one if the bond was issued during the expansion1 and expansion2 periods, respectively, and zero otherwise. Issuer control variables and industry fixed effects are also included to control for compositional changes in issuers. The baseline model includes a dummy variable for BBB-rated bonds, the natural logarithm of total assets, net book leverage, profitability, and asset tangibility. It should be noted that some of the variables lead to a separation problem, under which the maximum likelihood estimates do not exist (Albert and Anderson, 1984).[24] I therefore use the bias-corrected multinomial logit model of Kosmidis and Firth (2011), which can handle the problem of separation.[25]

Note that the coefficients of my MNLM correspond to the effects on the log odds ratio of choosing a particular maturity bin relative to the reference maturity bin. In my baseline analysis the maturity bin of (7,10] years is set as the reference category since it is the most popular maturity category. Note that while the choice of the reference category determines the parameterization of regression coefficients, it does *not* affect the estimated marginal effects.

I emphasize that the results of $Expansion1$ should be taken with great caution, and my interest is in those of $Expansion2$. To begin with, the expansion1 period was quite short, spanning only 41 days. Moreover, it appears that the April-27 announcement of the CBPP's massive expansion had been anticipated by market participants. The announcement was made immediately after a pre-scheduled Policy Board meeting on that day, and some of the topics to be discussed had already been reported by newspapers. First, a Reuters exclusive on April 14 said that the meeting would discuss options to expand its CP and corporate bond purchases.[26] Second, on April 23, Nikkei, Japan's leading business newspaper, reported that the BOJ was planning to expand both the size and scope of the purchases.[27] These anticipations are important because the vast majority of bond

[24] Complete separation occurs when the outcome variable is perfectly predicted by a set of explanatory variables. A similar situation but with some overlaps in prediction is called quasi-complete separation. For a more detailed explanation, see Albert and Anderson (1984). In my sample, one instance of a separation problem is that no health care firms (ICB industry code 20) chose the maturity bin of >10 years.

[25] The MNLM of Kosmidis and Firth (2011) is based on the penalized likelihood method of Firth (1993), which has been shown to be effective at solving the problem of separation (Heinze and Schemper, 2002).

[26] Leika Kihara and Takahiko Wada (2020, April 14) Exclusive: BOJ considering steps to ease corporate funding strains in April - sources. *Reuters.* Retrieved from `https://www.reuters.com/article/uk-health-coronavirus-japan-boj-exclusiv/exclusive-boj-considering-steps-to-ease-corporate-funding-strains-in-april-sources-idUKKCN21W0IY?`.

[27] *Kokusai kounyuu seigen naku, nichigin giron he CP shasai kounyuu baizou* (The BOJ to discuss removing its JGB purchase limit and doubling its CP and corporate bond purchases) (in Japanese) (2020, April 23) *Nikkei.* Retrieved from `https://www.nikkei.com/article/DGXMZO58430050T20C20A4MM8000`.

issuances during *expansion1* took place later in the period; 21 out of the 23 bonds (91.3%) issued during this period were issued on or after April 16.[28] For these reasons, my main attention is directed to comparing the pre-expansion and expansion2 periods.

The MNLM results are presented in Table 2.6. Panel A shows the baseline result where the reference category is the maturity bucket of (7,10] years, while alternative reference categories are used in Panel C. To gauge the economic significance, Panel B computes the average marginal effects (AMEs) of $Expansion1$ and $Expansion2$. These marginal effects are measured as the changes in the predicted probabilities by changing the values of the dummy variables from zero to one.[29]

The estimated coefficients of $Expansion2$ strongly support the view that, during *expansion2*, some firms shifted maturities in the vicinity of the eligibility threshold (five years). In Panel A of Table 2.6, the significant and negative coefficient of $Expansion2$ for the maturity bin of (5,7] years means that the maturity bin became less likely to be selected as compared to the reference maturity bin of (7,10] years when I compare the pre-expansion and expansion2 periods, holding issuer characteristics constant. Likewise, Column 2 of Panel C indicates that firms also became less likely to choose maturities of (5,7] years in comparison to maturities of (3,5] years. Moreover, Panel C shows that these results hold even when the maturity bin of (5,7] years is compared with the shorter maturity bins combined (from 1 to 5 years, Column 3) or the longer maturity bins combined (greater than 7 years, Column 4). Therefore, the decrease in the probability of choosing the maturity bin of (5,7] years is specific to this maturity region, rather than reflecting a shift toward either shorter or longer maturities. This result is consistent with the theoretical prediction that firms whose target maturities are only moderately above the eligibility threshold (five years) are the most likely to shorten maturities to meet the criterion.

Panel B documents that the changes from *pre-expansion* to *expansion2* are also economically significant; the probability of a bond issuer choosing maturities of (5,7] years was lower by 11.2

[28] In the U.S., the SEC requires bond issuers to file a registration statement and then a pricing document. Butler et al. (2023) report that the median number of days between the two filings is six for the 100 bonds randomly selected from their sample. Such an exercise is not possible for Japanese bond issuers as they are required to file only one document that includes full issuance information.

[29] Because they can be viewed as dummy variables for a categorical variable with three levels (i.e., *pre-expansion*, *expansion1*, and *expansion2*), the marginal effects of *Expansion1* (*Expansion2*) are computed by using only bonds issued during the pre-expansion and expansion1 (expansion2) periods, as recommended by Bartus (2005).

percentage points in *expansion2* than in *pre-expansion*, holding issuer characteristics fixed. At the same time, the AMEs of $Expansion2$ for the shorter maturity bins, which were eligible for the BOJ purchases, are positive and significant. To put these marginal effects into perspective, Figure 2.6 displays the predicted probabilities derived from the MNLM of Table 2.6 for the three sub-periods. It shows that the 11.2 percentage-point decrease of the maturity bin of (5,7] years translates into a 62% decrease relative to the predicted probability of this maturity segment for the pre-expansion period (17.9%).

However, the changes in firms' bond maturity choices from *pre-expansion* to *expansion2* were not limited to the maturity segments close to the eligibility threshold (five years). First, an increase in the eligible maturities occurred not only in (3,5] years but also in [1,3] years. In Section 2.3, it is argued that the BOJ's CBPP expansion was like two separate shocks to bonds with remaining maturities of [1,3] years and (3,5] years, with the shock to the former segment possibly being greater. The MNLM result can be seen as consistent with the view that firms catered to each of the demand shocks. It should be noted, however, that this result could simply reflect the tendency of newly issued debt to have shorter maturities during a downturn (Erel et al., 2012). It is worth repeating that possible confounding effects of COVID-19-driven shocks are the main reason why I focus on changes in firms' maturity choices around the BOJ's eligibility threshold for identification.

Second, the longest maturity segment (>10 years) became less popular in the expansion2 period. The discussion in Section 2.3 suggests that this decrease was *not* directly driven by firms' incentives to cater to the BOJ-driven demand shock. There are two possible alternative explanations for this result. First, as previously noted, the maturities of newly issued debt tend to be shorter during a 'typical' crisis time, and Erel et al. (2012) provide two possible explanations. First, higher adverse selection costs can encourage issuers to issue less information-sensitive securities. Second, capital suppliers might become more risk averse during a crisis (Caballero and Krishnamurthy, 2008; Vayanos, 2004), leading to a decrease in the demand for riskier securities. Given that very long-term bonds are the most information-sensitive and the riskiest, these capital demand- and supply-side factors might have particularly discouraged firms from issuing very long-term bonds during the COVID-19 crisis in Japan.

The second possible explanation for the decrease in the issuance of very long-term bonds is the rise in very long-term yields in the expansion2 period. Figure 2.7 documents that during *expansion2*, yields of JGBs with very long maturities (20 and 30 years) kept increasing and that these movements were in contrast to those of shorter-term rates, which remained stable during the period.[30] This distinct rise in the very long end of the term structure in Japan might have discouraged firms from choosing those segments when issuing a bond.[31]

The results of $Expansion2$ are robust to various changes in the model settings. First, controlling for possible seasonality in firms' bond maturity choices does not alter the results. Murfin and Petersen (2016) document seasonal variations in interest rates and loan volumes using U.S. data and note that one way to control for possible seasonality is by adding month-by-year fixed effects. I therefore estimated the baseline MNLM with month-by-year fixed effects on an extended sample covering four years from October 2016 to September 2020. The result remains similar (Table A2.3 in Appendix 2.D). Second, using an alternative pre-expansion period of six months before the BOJ's first announcement did not qualitatively change the MNLM result. Lastly, the $Expansion2$ results remain qualitatively similar even when the control variables and industry fixed effects are dropped.

Cross-sectional differences

Greenwood et al. (2010) predict that firms with greater financial flexibility respond to demand/supply shocks more strongly. The reason is that deviating from the intrinsically optimal maturity is expected to be less costly for them. The existing literature on this subject largely supports this prediction. Greenwood et al. (2010) state that their empirical analysis supports this prediction "largely, though not entirely." Badoer and James (2016) report from their analysis of very long-term debt issuance that only firms rated A or above exhibited gap-filling behavior.

[30] Figure 2.7 also suggests that a possible yield curve change was *not* a driver of the main finding of this paper—a disproportionate decrease in issuances of corporate bonds maturing in (5,7] years. Throughout the sample period, the yields for five- and seven-year JGBs moved almost in tandem.

[31] Under the expectations hypothesis firms cannot benefit from this kind of timing behavior. However, it might result from managers' desire to inflate near-term earnings or their speculation on future interest rate changes (Faulkender, 2005; Chernenko and Faulkender, 2011; Greenwood et al., 2010). Indeed, Graham and Harvey (2001, pp. 224–225) document that more than a third of CFOs answered that issuing "short-term when short-term interest rates are low compared to long-term rates" is "important" or "very important." Also, Guedes and Opler (1996) find that firms tend to choose shorter maturities when the yield curve is steeper.

I therefore examine the cross-sectional heterogeneity in the changes in bond maturities from *pre-expansion* to *expansion2*. This is done by using the MNLM to compute the AMEs of *Expansion2* for sub-samples divided by proxies of financial strength.[32] The results reported in Table 2.7 show, however, no strong cross-sectional differences. In Panel A, bonds are divided into AA-rated and A/BBB-rated, with the vast majority (92%) of the latter category being rated A. The signs of the AMEs of *Expansion2* are the same for the two groups, and none of the differences are statistically significant at the conventional levels. Similar analyses using market capitalization (Panel B) and net book leverage (Panel C) yield no clear differences. In Appendix 2.D, I also divide firms based on the ratio of bank debt to total debt (obtained from the Capital IQ Capital Structure database) and repeat the exercise. The underlying idea is that a close relationship with a bank might affect the firm's financial flexibility. However, this analysis does not reveal statistically significant cross-sectional differences, either (Table A2.5).

Therefore, in contrast to previous studies, I do not find strong cross-sectional differences in firms' catering behavior. Of course, the unique nature of my sample period, the COVID-19 crisis, might have influenced firms' incentives. There are, however, two other possible reasons for the null result in Japan. First, the cross-sectional variation in the financial flexibility of Japanese bond issuers is quite limited. Specifically, as already mentioned, access to the Japanese primary bond market has been virtually restricted only to highly creditworthy firms. I had to use a cutoff point of AA versus A/BBB in the cross-sectional analysis because 94.7% of the sample bonds are rated AA or A (and none of them are rated below BBB). In contrast, Badoer and James (2016) divide their sample bonds into those rated AAA/AA/A and those rated BBB or below. Also, in one of their cross-sectional analysis, Greenwood et al. (2010) compares dividend payers (financially strong issuers) and non-payers (financially weak issuers) and find that the former more strongly exhibits gap-filling behavior. In my sample, nearly all (99.6%) bonds were issued by dividend payers.

Second, the BOJ's auction protocol prioritizes corporate bonds with higher yields, i.e., bonds

[32] More specifically, this analysis is performed as follows: First, for each measure of financial strength, I create a dummy variable separating the sample. I then estimate the MNLM where the dependent variable is the same as before but the independent variables are *Expansion1*, *Expansion2*, the dummy variable about financial strength, and the interaction term of *Expansion2* and the dummy. Finally, I compute the AMEs of *Expansion2* separately for each subset of firms and test the differences in these AMEs between the two groups.

issued by *riskier firms*. This is because offers with the highest yields, *irrespective of the offered bonds' market yields*, are preferred and accepted first in the BOJ's reverse auctions of corporate bonds.[33] The preferential treatment could have more strongly encouraged riskier firms to engage in the catering behavior.

2.5.1 Discussions

Government debt supply

The literature has established that firms tend to avoid maturity segments in which government debt supply is abundant (Greenwood et al., 2010; Badoer and James, 2016; Lugo and Piccillo, 2019). Therefore, changes in the government debt supply might have contributed to the disproportionate changes in firms' maturity choices around the eligibility threshold of five years. Following the pandemic, the JGB supply fluctuated due to two (countervailing) factors: a supply *increase* through accelerated debt issuances by the Ministry of Finance (MOF) and a supply *decrease* through the BOJ's concurrent scaled-up quantitative easing operations. Nevertheless, Appendix 2.C shows that government debt supply did *not* change in the direction that would account for my main finding—the disproportionate reduction in issuances of (5,7]-year corporate bonds.

"BOJ-trade" demand

This paper has characterized the expansion of the BOJ's CBPP as a corporate bond demand shock. However, the BOJ's purchases of *already-isssued* bonds may also have impacted the primary market through a seemingly unique channel—by encouraging the so-called "BOJ trade." This section first details this trading strategy and then clarifies that this channel is essentially a byproduct of the BOJ's massive purchases, rather than a distinct phenomenon.

In short, the BOJ trade aims to profit from the BOJ's quantitative easing operations. According

[33] A more natural way to compare offers of different securities would be to compare each offer's difference between the offer yield and the market yield. The government bond reverse auctions of the BOJ and other central banks, such as the Fed (Song and Zhu, 2018) and the BoE (Breedon, 2018), essentially take this approach to compare offers of different government bonds. The BoE's reverse auctions of corporate bonds also follow this approach (Boneva et al., 2022). See Appendix 2.B for further details.

to the financial press, the accepted offer yields in the BOJ's reverse auctions of JGBs and corporate bonds tended to be lower than the at-issuance and secondary market yields of the associated securities. The BOJ trade tries to exploit this difference by predicting securities that could later be sold to the BOJ and purchasing them in the primary or secondary market (Tomisawa and Hazama, 2019). Media reports suggest that the massive expansion of the BOJ's CBPP provided a great opportunity for the BOJ trade, and thus escalated private investors' demand for eligible corporate bonds.[34] Appendix 2.B also presents some empirical evidence of the BOJ trade opportunity.

The above discussion implies that bond issuers could have catered to private investors' BOJ-trade-driven demand by choosing a maturity of five years or less. This paper is unable to isolate this specific influence. Nevertheless, the existence of the BOJ trade opportunity can be viewed as an unavoidable byproduct of the BOJ purchases given their size; it is plausible that such gigantic purchases inevitably cause price impact, no matter how implemented. Notably, corporate bond markets, particularly those outside the U.S., are characterized by infrequent trading and low liquidity. Therefore, the massive demand shock from the BOJ is also the root cause of this BOJ-trade demand.

2.5.2 Simultaneous issuances of bonds with varying maturities

One important feature of corporate bond issuances is that firms often issue bonds of *different* maturities at the *same* time (Houston and Venkataraman, 1994; Helwege and Turner, 1999). While my preceding analyses have ignored this feature, in this section bonds issued on the same day are grouped together. Furthermore, I utilize these multiple-maturity issuances to conduct a cleaner test of firms' catering to the BOJ-driven demand shock.

[34]For instance, in July 2020, an article by Nikkei stated that an anonymous fund manager told the newspaper that her/his fund had been deciding which corporate bonds to purchase based on whether they could be later bought by the BOJ rather than credit ratings. Satoshi Matsui, Shugo Yamada, and Taichiro Sunaga (2020, July 23) *Nichigin kanwa 3kagetsu shasai CP shijou ni onkei senmei* (Easing of the Bank of Japan for 3 months—clear benefits to the corporate bond/CP market) (in Japanese). *Nikkei*. Retrieved from `https://www.nikkei.com/article/DGXLASFL22I19_S0A720C2000000/`.

Proportion of multiple-maturity bond issuances

Interestingly, multiple-maturity bond issuances became more popular in the post-pandemic period. Panel A of Table 2.8 shows that while in *pre-expansion* around half (50.5%) of bond issuance events were multiple-maturity issuances, this percentage increased to 70.5% in *expansion2*. Fisher's exact test of the distributions being different is statistically significant (p-value < 0.01).

There are two reasons for the increase in the proportion of multiple-maturity bond issuances. On one hand, the increase appears to align with the findings of Choi et al. (2018, 2021), who show that when rollover risk increases, firms increase the dispersion of their debt maturities. Based on this *maturity dispersion view*, the COVID-19 pandemic encouraged firms to choose multiple-maturity bond issuances because of the heightened economic uncertainty (and the resulting rollover risk increase). On the other hand, it is also possible that the BOJ's corporate bond purchases contributed to this increase. This is because multiple-maturity bond issuances can be a means to *partially* cater to the demand shock: by issuing a bond whose maturity is short enough for the BOJ purchases and another bond with a longer maturity, the firm was able to partially cater to the demand shock without much increasing its reliance on relatively short-term debt. I call this a *partial catering view*.

Maturity *compositions* of multiple-maturity bond issuances have changed in manners consistent with the partial catering view, although this finding does not rule out the maturity dispersion view. Panel B of Table 2.8 shows that the increase in multiple-maturity bond issuances was concentrated on those including maturities that meet the BOJ's eligibility criterion. Conditional on being a multiple-maturity issuance, the probability of including a bond maturing in (3,5] years ([1,3] years) increased from 59.1% (18.3%) in *pre-expansion* to 92.7% (47.3%) in *expansion2*. At the same time, that of including a bond maturing in (5,7] years decreased from 44.1% to 16.4%. Furthermore, there was a significant increase in the proportion of multiple-maturity issuances including an eligible maturity bond *and* a bond maturing in ten years or more. The proportion increased from 50.5% in *pre-expansion* to 74.5% in *expansion2* (the last row of Panel B). Arguably, this type of multiple-maturity bond issuance particularly fits the partial catering view.

Analysis of changes in maturity compositions of multiple-maturity bond issuances

Panel C of Table 2.8 shows maturity compositions of one common type of multiple-maturity issuances: those including maturities of [1,5] and ≥10 years. It is again shown that individual firms altered maturities in the manner of partially catering to the BOJ's massive corporate bond purchases. For instance, in *pre-expansion*, there were no simultaneous issuances of bonds maturing in 3, 5, and 10 years, and there was only one simultaneous issuance of bonds maturing in 3, 7, and 10 years. In *expansion2*, while the former increased to 16, the latter decreased to zero. These contrasting changes can hardly be explained by heightened uncertainty or the consequent rise in rollover risk; those two types of multiple-maturity issuances consist of the same number of different maturities (three), cover the same maturity range (from three to ten years), and have very similar average maturities (6 years for the former and 6.7 years for the latter). In contrast, these changes are consistent with the notion that the BOJ's massive purchases made the maturity of seven years disproportionately less appealing for bond issuers.

Therefore, these issuances offer a unique opportunity to examine how bond issuers with *similar desired maturity ranges* select a *set of different maturities*. As such, this sample enables a cleaner test of firms' catering to the BOJ-driven bond demand shock. I thus perform a more formal test using this sub-sample (i.e., multiple-maturity issuances including maturities of [1,5] and ≥10 years). The logit model is employed to estimate the change in the probability of including a maturity of seven years—the only maturity chosen within the maturity bin of (5,7] years in Panel C of Table 2.8—after controlling for issuer- and issuance-characteristics. To be specific, the dependent variable takes the value of one if the simultaneous bond issuance includes a bond with a seven-year maturity, and zero otherwise. The control variables for issuer-characteristics are the same as those used for the MNLM. The following variables are also included to control for issuance-characteristics: the natural log of proceeds-weighted average maturity and a dummy variable that takes a value of one if the multiple-maturity issuance includes three different maturities or more. The bias-corrected logit model of Kosmidis and Firth (2009) is employed as it can obtain maximum likelihood estimates even when quasi-complete or complete separation is present.

Table 2.9 confirms that, during the expansion2 period, the sample multiple-maturity bond issuers were less likely to include a bond with a seven-year maturity than in the pre-expansion period. In Column 1 the coefficient of *Expansion*2 is negative and significant at the 1% level, and the economic significance is large; the AME of *Expansion*2 is estimated to be −22.4 percentage points, which represents the difference in the predicted probability of including the maturity of seven years between the pre-expansion period (41.7%) and the expansion2 period (19.4%). In addition, this result is robust: Columns 2 and 3 show that the coefficient of *Expansion*2 increases in size when more independent variables are added to control for issuance- and issuer-characteristics. In Columns 4–6, I add a further sample restriction that the sample multiply-maturity issuances must consist of three maturities. This additional restriction results in the sample maturity compositions of only two types: maturity combinations of 3, 5, and 10 years and 3, 7, and 10 years. Clearly, these issuers must have had very similar preferences for the overall maturity structures of their new issues. Although this restriction nearly halves the sample size, the coefficients of *Expansion*2 are also negative and remain significant at the 1% level. The magnitude of the effect seems greater for this sub-sample. In Column 4, the estimated AME is −58.0 percentage points as the predicted probability decreases from 81.3% to 23.3% when moving from the pre-expansion period to the expansion2 period.

2.5.3 Offering yields and spreads

Having documented changes in firms' bond maturity choices, I now move on to examine issuing costs. Table 2.4 reveals two notable features. First, the average offering spread is *higher* than the average offering yield, reflecting that yields of JGBs with maturities of up to 10 years were negative for most of the sample period (see Figure 2.7). Due to this complication arising from negative interest rates, the subsequent analysis mainly focuses on offering yields. Second, offering yields (and spreads) tend to be very low, with the median offering yield being merely 28 basis points. The financial press suggests that these extremely low offering yields were, in part, the result of BOJ-trade-oriented investors, for whom at-issuance yields were of second-order importance (Tomisawa and Hazama, 2019).

To analyze the relationship between maturity and the pricing of newly issued bonds, I regress offering yields and spreads on the following variables: the natural logarithm of proceeds, the issuer control variables (*Ln(total assets)*, *Net book leverage*, *Profitability*, and *Asset tangibility*), as well as fixed effects for maturity bin, credit rating, and industry. These models are estimated separately for the pre-expansion and expansion2 periods to investigate whether the coefficients of the maturity bin dummies differ between the two periods.

Table 2.10 presents the regression results of offering yields. It implies that the relationship between offering yields and maturities changed only weakly at best. First, although all the maturity bin dummies have greater coefficient values in the expansion2 period (i.e., the curve became steeper), the difference is statistically significant only for the bin of (7,10] years. Second, there is no sign of a discontinuous change around the eligible maturity threshold. No statistically significant changes are found when a similar analysis is performed on offering spreads (Table A2.6 in Appendix 2.D).

Yet, cautions should be taken when interpreting these results of credit yield and spread curves. First, they do not necessarily represent the causal effect of maturities on credit yields/spreads due to the self-selected nature of bond maturities (Helwege and Turner, 1999). Second, this paper has documented that Japanese bond issuers became much less likely to choose the maturity bin of (5,7] year, and this in turn means that the sample size of these bonds is very limited for the expansion2 period. As a result, in the expansion2 period, the coefficients of the maturity bin of (5,7] years cannot be precisely estimated.

There is also a theoretical reason why one does *not* necessarily expect the BOJ's purchases to have a strong effect on the term structure of offering yields/spreads—firms as a whole can elastically respond to demand-driven yield differences across the maturity spectrum. Note that this is exactly the main finding of this paper. Theoretically, in the most extreme case of a frictionless capital market where the maturity choice is irrelevant, firms do this until all the differences in expected returns are eliminated (Greenwood et al., 2010). As a result, Greenwood et al. (2010, p. 997) state: "[W]e should stress that our model's implications for returns are neither as fundamentals nor as robust as its implications for quantities." Therefore, it is possible that the increase in the supply of bonds maturing in five years or less offset the disproportionate effect that the positive demand shock could have had on the pricing of newly issued bonds in this maturity segment if firms had not altered bond maturities.

2.6 Conclusions

This paper demonstrates that the BOJ's post-pandemic massive corporate bond purchases led firms to shorten the maturities of newly issued bonds to meet the maturity eligibility criterion. Given that the COVID-19 shock itself should have influenced firms' debt maturity choices and there is no clear 'control' group, this paper focuses on maturity distribution changes in the vicinity of the eligibility threshold (five years). Consistent with theoretical predictions, the proportions of eligible, shorter-term maturity segments increased, while that of the maturity segment *just above* the threshold, (5,7] years, decreased relative to the neighboring maturity bins. In addition, I take advantage of multiple-maturity bond issuances to document cleaner evidence that individual issuers indeed shifted maturities at the intensive margin to cater to the BOJ-driven bond demand shock.

In contrast, no evidence suggests that U.S. firms shortened bond maturities to meet the maturity eligibility criteria of the Fed's facilities. Why did the catering occur only in Japan? I argue that this likely reflects that the BOJ's *actual* purchase amount was much larger as a proportion of the target market size. There were, however, at least two other noteworthy institutional differences. First, while the BOJ's program targeted only already issued bonds, the Fed created the PMCCF and the SMCCF to purchase newly issued and already issued bonds, respectively. (Note, however, that while the PMCCF was operative, it did not conclude any purchases.) Second, the purchase methods were different. While the BOJ purchased corporate bonds through reverse auctions, the SMCCF directly purchased bonds from the secondary market with help from BlackRock. Nevertheless, it seems unlikely that these institutional differences, rather than the massive difference in actual purchase size, led to the very different results.

As corporate bond purchases are increasingly becoming a common policy tool of central banks, this paper's finding is relevant to policymakers not only in Japan but also in many other economies. The BOJ and the ECB have conducted corporate bond purchases since the 2008–2009 financial crisis, and many central banks have employed this tool in response to the pandemic. Also, persistent low interest rates in many high-income countries, which Blanchard and Summers (2020) call "Japanification," suggest the increasing importance of unconventional monetary policies. The Japanese evidence illustrates a potential trade-off for central banks purchasing non-government debt—although a stringent maturity eligibility criterion can limit the central bank's risk exposure, it may also generate a distortionary effect on debt issuers' maturity choices.

2.7 Figures

Figure 2.1: Maturities of newly issued corporate bonds in Japan and the U.S.

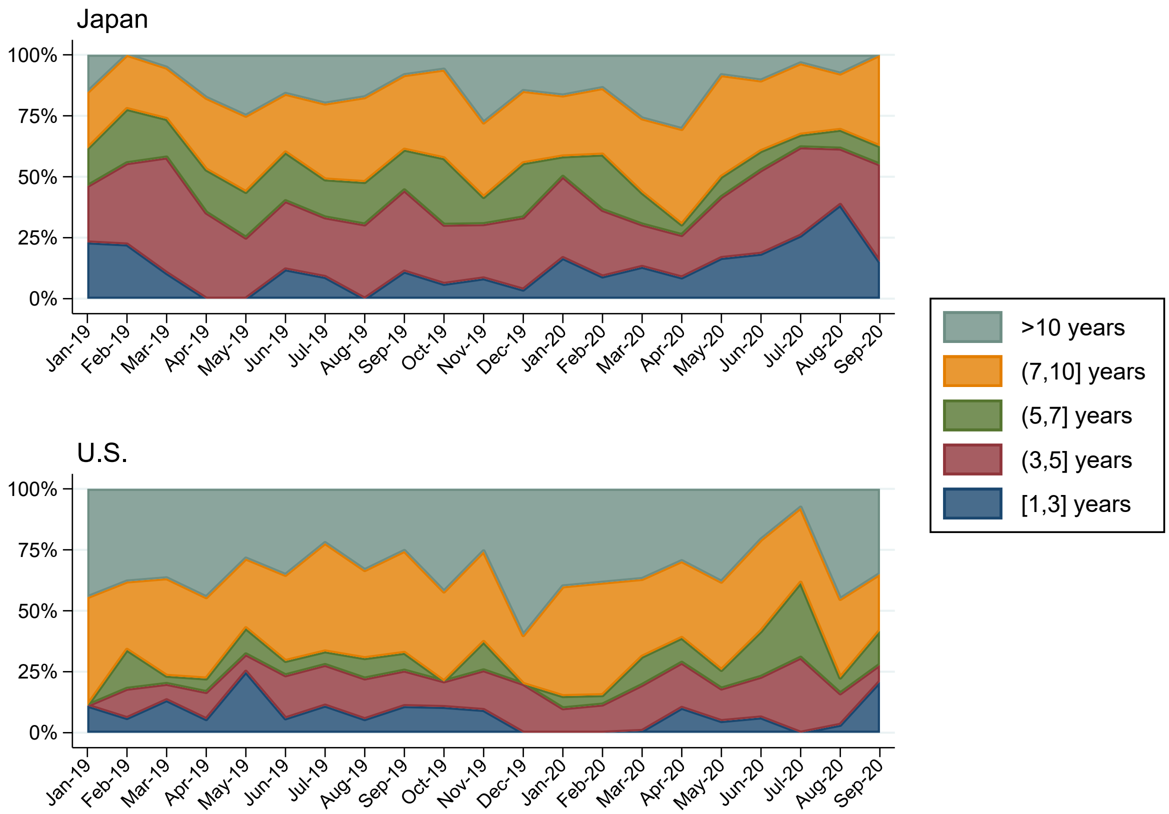

This figure shows the proportions of maturities of newly issued corporate bonds in Japan and the U.S. The sample firms are public, non-financial, and non-utility firms. See Section 2.4 for the data sources.

Figure 2.2: The BOJ's new corporate bond purchase amounts

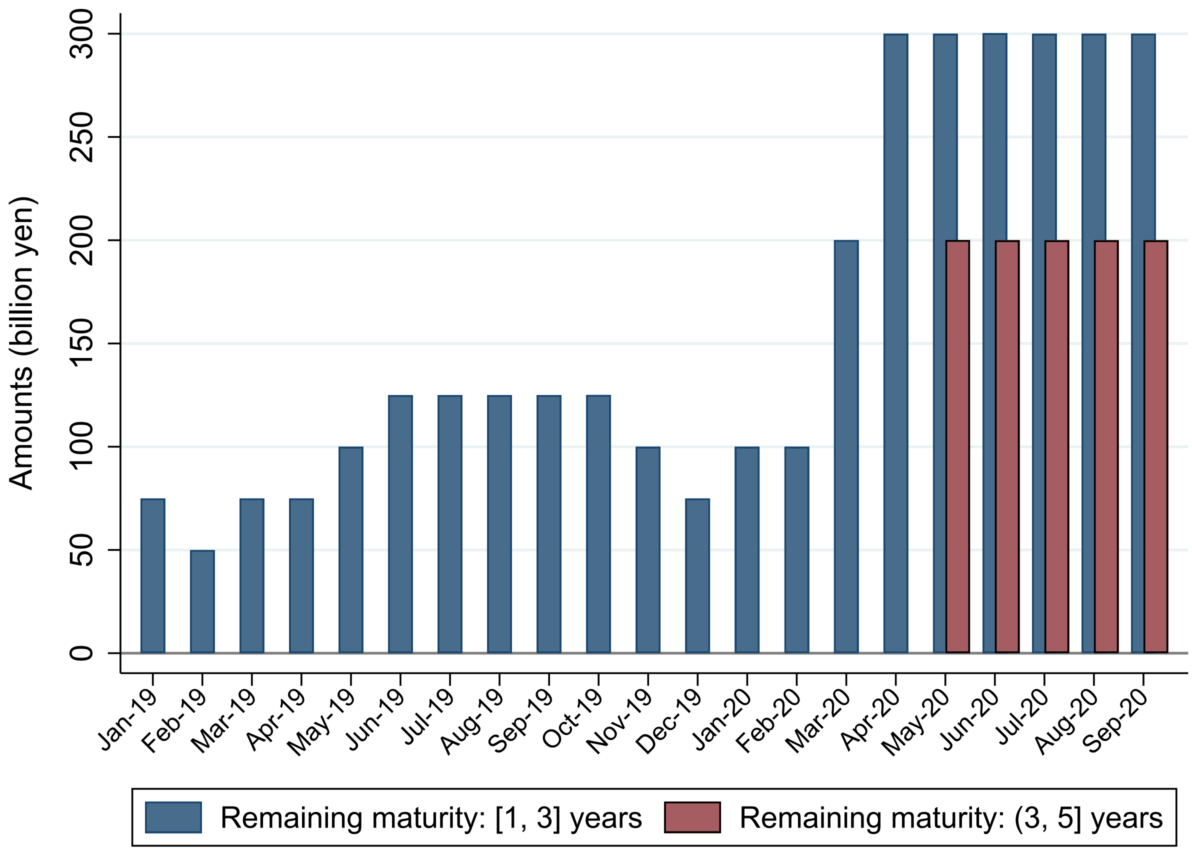

This figure shows the amounts of the BOJ's new corporate bond purchases (in billions of yen) on a monthly basis from January 2019 to September 2020. The data were obtained from the BOJ's website regarding its market operations (https://www.boj.or.jp/en/statistics/boj/fm/ope/index.htm/).

Figure 2.3: Comparison of the corporate bond purchase sizes of the BOJ and the Fed

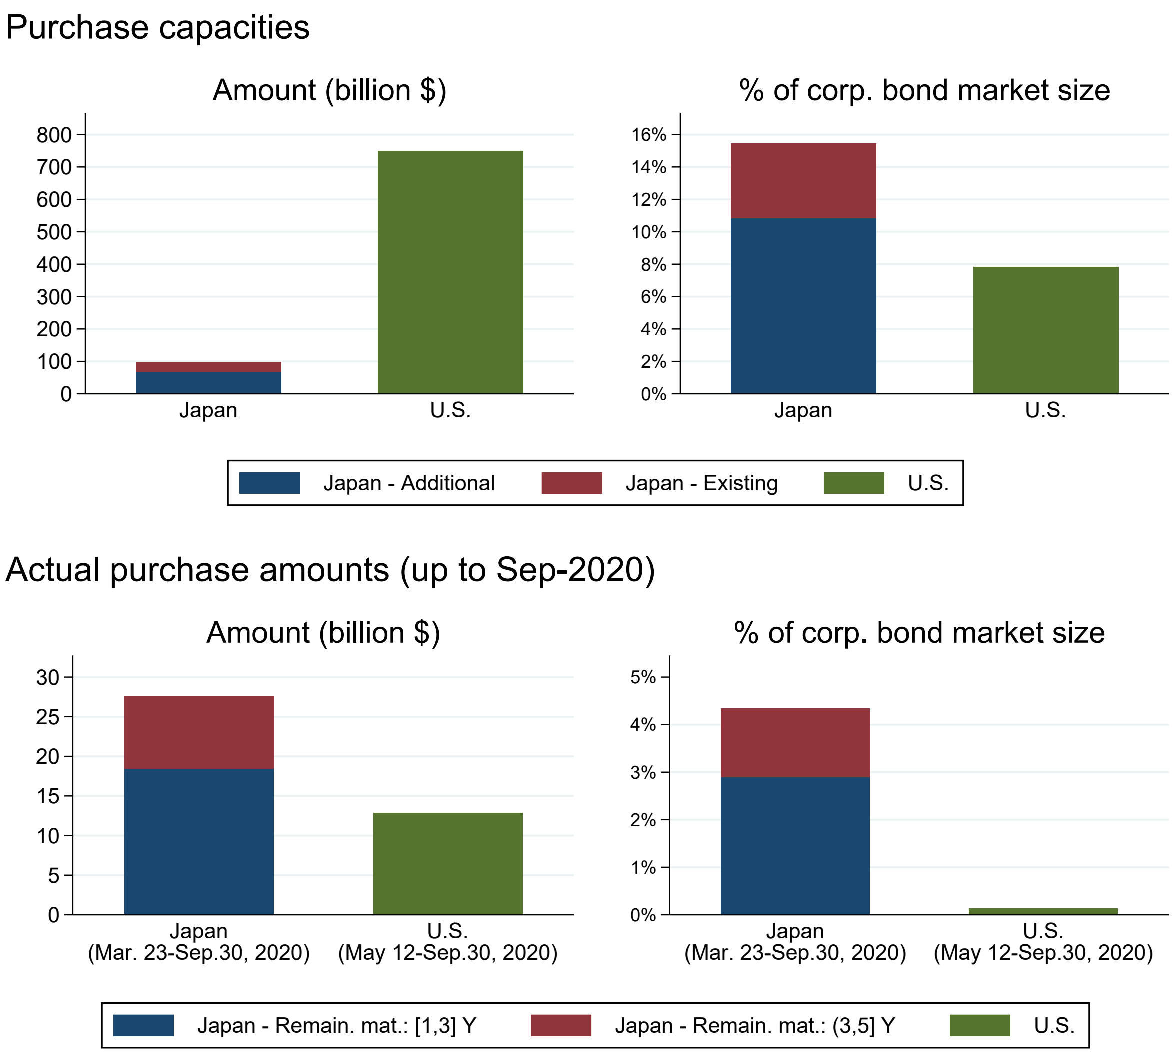

The top left graph shows the corporate bond purchase capacities of the BOJ and the Fed (i.e., the PMCCF and the SMCCF) during the COVID-19 pandemic. The BOJ's purchase capacity consists of the preexisting amount (Japan - Existing) and the additional amount announced on April 27, 2020 (Japan - Additional). In the top right graph, the purchase capacities are normalized by the domestic corporate bond market sizes, whose data come from the Japan Securities Dealers Association (JSDA) for Japan and SEC (2020) for the U.S. The bottom graphs show the actual purchase amounts up to September 2020. The BOJ's purchase amount is based on the reverse auctions conducted from March 23 to September 23, 2020. The actual purchase amount of the Fed is based on the SMCCF's total outstanding amount as of September 30, 2020. (The PMCCF did not conclude any purchases.) The data source of the actual purchase size of the BOJ is the same as that of Figure 2.2. The SMCCF purchase amounts were obtained from the Fed's public releases available at `https://www.federalreserve.gov/publications/2020-reports-to-congress-in-response-to-covid-19.htm`. Again, these amounts are normalized by the domestic corporate bond market sizes in the bottom right graph.

Figure 2.4: Corporate bond issues in Japan and the U.S.

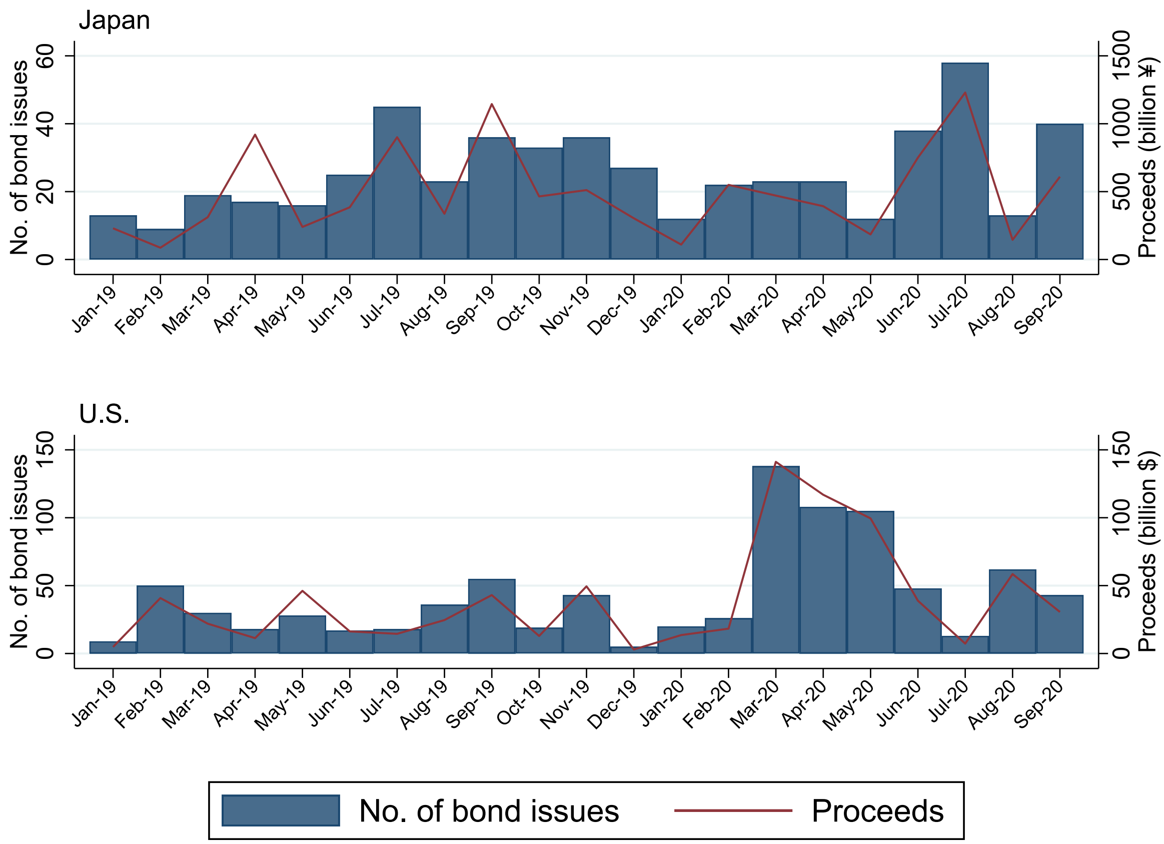

This figure shows the number of newly issued bonds and their aggregate proceeds. The sample firms are public, non-financial, and non-utility firms. The Japanese data were obtained from the JSDA and Datastream. U.S. bond issuance data were collected from the SDC Platinum. To obtain a sample similar to that used by Halling et al. (2020), I kept only public firms in non-financial and non-utility industries that could be matched to the Compustat.

Figure 2.5: Maturity distribution

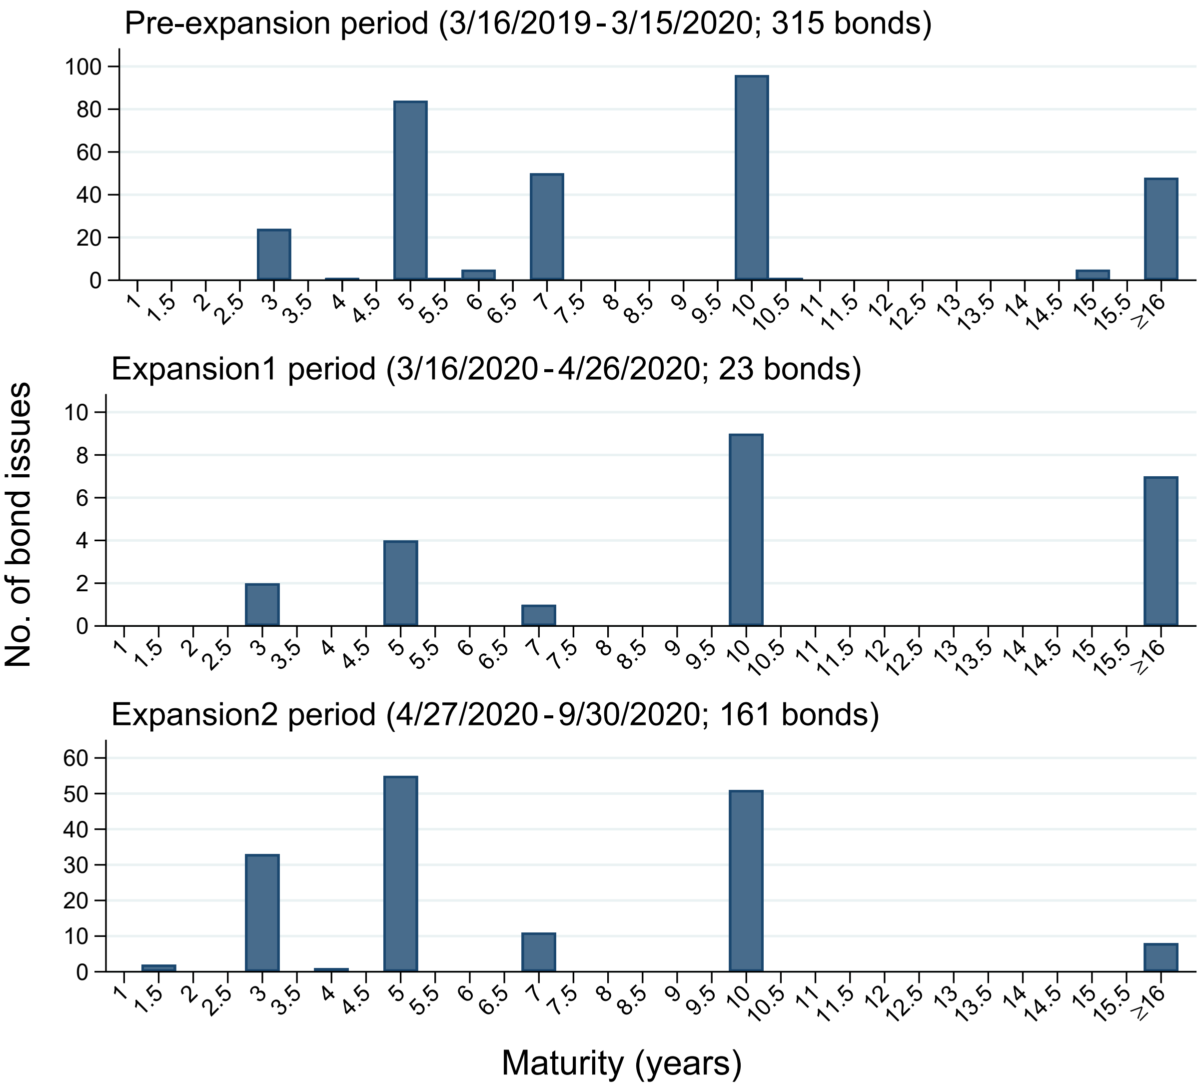

This figure shows the distributions of corporate straight bond issues by maturity. The issuers are public firms in non-financial and non-utility industries in Japan. The data were obtained from the JSDA.

Figure 2.6: Predicted probabilities based on the baseline multionomial logit model

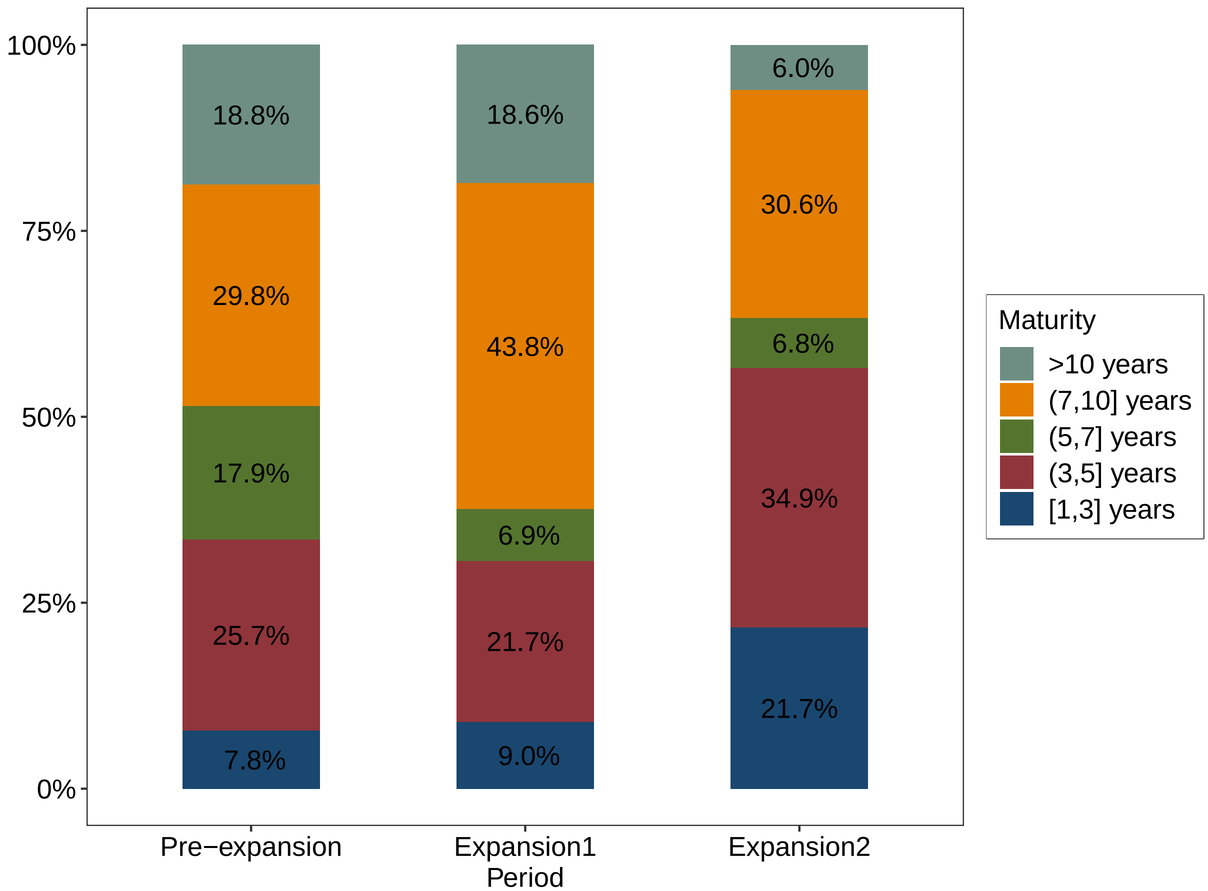

The estimated predicted probabilities of the maturity bins for the pre-expansion, expansion1, and expansion2 periods are displayed. They are derived from the multinomial logit model in Table 2.6.

Figure 2.7: Dynamics of the term structures of Japanese government bonds

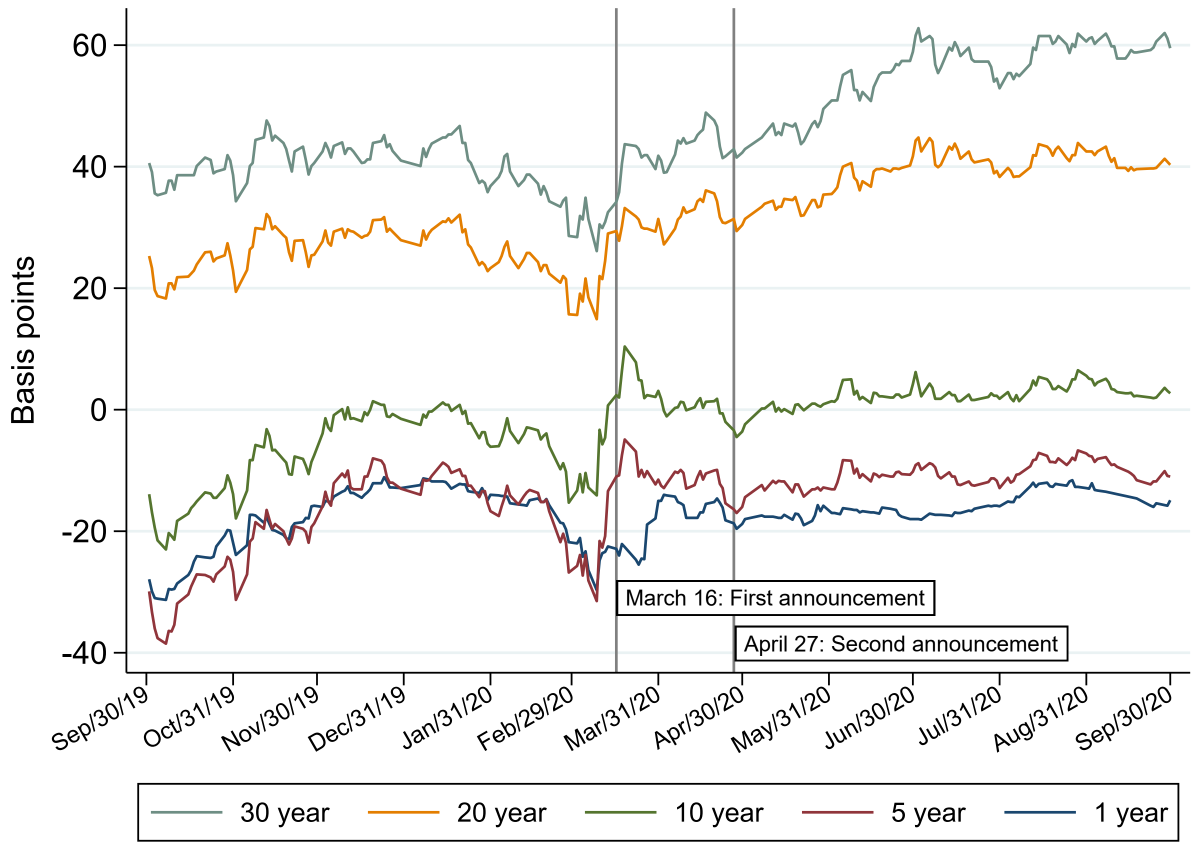

This figure plots the yields of Japanese government bonds (JGBs) with constant maturities of 1, 5, 7, 10, 20, and 30 years for the one-year period leading up to September 30, 2020. The JGB yield data were obtained from the Ministry of Finance's website at `https://www.mof.go.jp/english/jgbs/reference/interest_rate/index.htm`.

2.8 Tables

Table 2.1: Key contents of the BOJ's announcements on March 16 and April 27, 2020

	March 16 announcement: "Enhancement of Monetary Easing in Light of the Impact of the Outbreak of the Novel Coronavirus (COVID-19)"	**April 27 announcement: "Enhancement of Monetary Easing"**
CP and corporate bond purchases	• The caps for CP and corporate bond purchases were increased by one trillion yen each. Consequently, the cap for CP was increased from ¥2.2 trillion to ¥3.2 trillion and that for corporate bonds was increased from ¥3.2 trillion to ¥4.2 trillion. • The additional purchases of one trillion yen each would be conducted until September 2020. • The key eligibility criteria remained the same. Eligible bonds were still required to have credit ratings of BBB or above (i.e., investment grades) and one- to three-year remaining maturities.	• The amounts of additional purchases of CP and corporate bond purchases were increased from one trillion yen each to ¥7.5 trillion each. As a result, the total cap became around ¥20 trillion. • The termination date was moved from September 2020 to March 2021. • While the credit rating eligibility criterion did not change, the maximum remaining maturity was extended from three years to five years. • The maximum purchase amount per issuer was increased from ¥100 billion to ¥500 billion for CP and from ¥100 billion to ¥300 billion for corporate bonds. Similarly, the BOJ's maximum ownership share per issuer was increased from 25% to 50% for CP and from 25% to 30% for corporate bonds.
Equity ETF purchases	The cap was increased from ¥6 trillion to ¥12 trillion.	The cap remained at ¥12 trillion.
"Special funds-supplying operations to facilitate corporate financing regarding the novel coronavirus (COVID-19)"	The BOJ launched the special funds-supplying operations, through which the BOJ would provide financial institutions with up to one-year loans at zero rate. The loans could be made against a variety of corporate debt, which totaled ¥8 trillion as of February 29, 2020.	Most notably, the eligible collaterals were expanded to a wider variety of private debt, which totaled ¥23 trillion as of March 31, 2020.
JGB purchases	The BOJ's JGB purchases would be maintained. The amount of JGBs held by the BOJ was planned to be increased by around ¥80 trillion per year.	The cap for JGB purchases was removed.

This information was obtained from the BOJ's website: https://www.boj.or.jp/en/mopo/mpmdeci/state_2020/index.htm/.

Table 2.2: Bond issuances and maturity by credit rating

Panel A: Credit rating distribution

	Pre-expansion period (3/16/19–3/15/20)	Expansion1 period (3/16/20–4/26/20)	Expansion2 period (4/27/20–9/30/20)	Total
AA	99	11	54	164
A	195	12	102	309
BBB	21	0	5	26
Total	315	23	161	499

Panel B: Maturity distribution by credit rating

	AA	A	BBB	Total
[1,3] years	28	30	3	61
(3,5] years	40	93	12	145
(5,7] years	20	45	3	68
(7,10] years	50	99	7	156
>10 years	26	42	1	69
Total	164	309	26	499

This table summarizes corporate straight bond issuances by credit rating at issuance for the period from March 16, 2019 to September 30, 2020. The issuers are public firms in non-financial and non-utility industries in Japan. When a bond is rated by more than one credit rating agency, the highest rating is assigned.

Table 2.3: Definition of variables

Variable	Description
Individual bond-level variables (Data source: JSDA)	
Maturity	Time to maturity (years)
Proceeds	Bond issue proceeds (billions of yen)
Offering yield	Offering yield (basis points)
Offering spread	Offering spread (basis points) measured as the difference between the offering yield and the yield of the same-maturity JGBs
Credit rating: BBB	Dummy variable that takes a value of one if the bond is rated BBB
Issuance event-level variables (Data source: JSDA)	
N. of different maturities	Number of different maturities included
Ln(sum of proceeds)	Natural logarithm of the sum of proceeds of the simultaneous bond issuance
Ln(average proceeds)	Natural logarithm of the average of proceeds of the simultaneous bond issuance
Including [1,3] Y	A dummy variable that takes a value of one if the simultaneous bond issuance includes a maturity in [1,3] years
Including (3,5] Y	A dummy variable that takes a value of one if the simultaneous bond issuance includes a maturity in (3,5] years
Including (5,7] Y	A dummy variable that takes a value of one if the simultaneous bond issuance includes a maturity in (5,7] years
Including (7,10] Y	A dummy variable that takes a value of one if the simultaneous bond issuance includes a maturity in (7,10] years
Including >10 Y	A dummy variable that takes a value of one if the simultaneous bond issuance includes a maturity greater than 10 years
Ln(weighted average maturity)	Natural logarithm of the proceeds-weighted average maturity
Weighted SD of maturities	Proceeds-weighted (population) standard deviation of maturities
Incl. [1,5] Y and ≥10 Y	A dummy variable that takes a value of one if the simultaneous bond issuance includes a maturity of [1,5] years and a maturity of 10 years or above
Firm-level variables (Data source: Datastream/Worldscope)	
Ln(total assets)	Natural logarithm of total assets (WC02999)
Net book leverage	Total debt (WC03255) minus cash and short-term investments (WC02001) divided by total assets
Profitability	EBITDA (WC18198) divided by total assets
Asset tangibility	Net PPE (WC02501) divided by total assets
Other variables	
Expansion1	Dummy variable that takes a value of one if the bond was offered between March 16 and April 26, 2020
Expansion2	Dummy variable that takes a value of one if the bond was offered on April 27, 2020 or later

This table lists the definitions and data sources of variables.

Table 2.4: Descriptive statistics

	Mean	S.d.	Min.	Median	Max.	N
Ln(maturity)	1.994	0.552	1.099	1.946	2.996	499
Ln(proceeds)	23.367	0.622	22.333	23.026	24.972	499
Offering yield (bps)	33.446	23.766	1.000	28.000	103.000	499
Offering spread (bps)	41.874	14.409	16.400	41.200	85.100	499
Credit rating: BBB	0.052	0.222	0.000	0.000	1.000	499
Ln(total assets)	27.944	1.158	25.638	28.030	30.449	490
Net book leverage	0.201	0.196	-0.196	0.206	0.554	490
Profitability	0.089	0.036	0.025	0.084	0.182	490
Asset tangibility	0.370	0.205	0.034	0.331	0.799	490
Expansion1	0.046	0.210	0.000	0.000	1.000	499
Expansion2	0.323	0.468	0.000	0.000	1.000	499

Variable definitions are provided in Table 2.3. All the continuous variables are winsorized at the 2.5% and 97.5% levels.

Table 2.5: Comparison by period

	Pre-expansion (N=311)	Expansion1 (N=23)	Expansion2 (N=156)	Expansion1 minus Pre-expansion	Expansion2 minus Pre-expansion
Ln(maturity)	2.073	2.273	1.815	0.200*	-0.258***
	(0.537)	(0.614)	(0.524)	[1.705]	[-4.938]
Ln(proceeds)	23.307	23.384	23.424	0.076	0.117**
	(0.636)	(0.581)	(0.539)	[0.558]	[1.966]
Offering yield (bps)	32.810	53.783	31.154	20.973***	-1.656
	(23.725)	(18.248)	(22.939)	[4.147]	[-0.719]
Offering spread (bps)	44.554	47.909	34.762	3.355	-9.792***
	(13.413)	(5.962)	(14.005)	[1.189]	[-7.332]
Credit rating: BBB	0.068	0.000	0.032	-0.068	-0.035
	(0.251)	(0.000)	(0.177)	[-1.287]	[-1.578]
Ln(total assets)	27.917	28.462	27.923	0.546**	0.006
	(1.200)	(0.789)	(1.104)	[2.145]	[0.056]
Net book leverage	0.199	0.268	0.196	0.069	-0.003
	(0.199)	(0.158)	(0.196)	[1.623]	[-0.155]
Profitability	0.093	0.070	0.086	-0.023***	-0.007*
	(0.037)	(0.021)	(0.034)	[-2.887]	[-1.929]
Asset tangibility	0.365	0.466	0.367	0.100**	0.001
	(0.205)	(0.213)	(0.202)	[2.256]	[0.071]

This table reports the means and the differences in means of bond- and issuer-characteristics by sub-period. The sample period consists of the following three sub-periods: the pre-expansion period (from March 16, 2019 to March 15, 2020), the expansion1 period (from March 16 to April 26, 2020), and the expansion2 period (from April 27 to September 30, 2020). The first three columns report the means accompanied by the standard deviations in parentheses. In the remaining two columns, bonds issued in the expansion1 and expansion2 periods are compared with those issued in the pre-expansion period, and t-statistics of the differences are reported in square brackets. Variable definitions are provided in Table 2.3. All the continuous variables are winsorized at the 2.5% and 97.5% levels. ***, **, and * indicate that the difference is significant at 1%, 5%, and 10% levels, respectively.

Table 2.6: Multinomial logit regressions of corporate bond maturity choices

Panel A: Coefficients (reference category = (7,10] years)

	[1,3] years	(3,5] years	(5,7] years	(7,10] years	>10 years
Expansion1	-0.265	-0.564	-1.366		-0.406
	(0.833)	(0.646)	(0.945)		(0.596)
Expansion2	1.042***	0.312	-1.012***		-1.395***
	(0.337)	(0.259)	(0.387)		(0.431)
Credit rating: BBB	0.542	0.663	0.166		-1.311
	(0.769)	(0.537)	(0.736)		(1.060)
Ln(total assets)	0.223	-0.093	0.089		0.365**
	(0.161)	(0.119)	(0.148)		(0.168)
Net book leverage	2.060*	0.011	-0.981		-0.089
	(1.218)	(0.876)	(1.105)		(1.245)
Profitability	0.122	-1.105	-4.092		-13.288**
	(5.882)	(4.273)	(5.292)		(6.650)
Asset tangibility	-0.381	-0.875	1.184		4.115***
	(1.007)	(0.808)	(1.026)		(1.059)
Industry FE	✓	✓	✓		✓
N	490				
McFadden's pseudo R^2	0.106				

Panel B: Average marginal effects

	[1,3] years	(3,5] years	(5,7] years	(7,10] years	>10 years
Expansion1	0.012	-0.040	-0.110*	0.140	-0.002
	(0.060)	(0.098)	(0.060)	(0.112)	(0.068)
Expansion2	0.139***	0.092**	-0.112***	0.009	-0.127***
	(0.036)	(0.046)	(0.030)	(0.046)	(0.027)

Panel C: Alternative reference categories

	[1,3] Y vs. Ref. cat. = (3,5] Y (1)	(5,7] Y vs. Ref. cat. = (3,5] Y (2)	(5,7] Y vs. Ref. cat. = [1,5] Y (3)	(5,7] Y vs. Ref. cat. = >7 Y (4)
Expansion1	0.299	-0.802	-0.825	-1.216
	(0.893)	(1.013)	(0.962)	(0.904)
Expansion2	0.729**	-1.324***	-1.536***	-0.676*
	(0.335)	(0.388)	(0.371)	(0.375)
Control variables	✓	✓	✓	✓
Industry FE	✓	✓	✓	✓
N	490		490	490
McFadden's pseudo R^2	0.105		0.105	0.078

This table reports the results of the multinomial logit regressions using the bias-correction methods of Kosmidis and Firth (2011). An R package `brglm2` (Kosmidis, 2020) is used for the estimation. The dependent variable is the bond maturity category. In Panel A, the reference category is the maturity bin of (7,10] years. The derived average marginal effects of *Expansion*1 and *Expansion*2 are reported in Panel B, together with standard errors obtained using the delta method. Panel C reports the MNLM results with alternative reference categories. In the first two columns of Panel C, the reference category is the maturity bin of (3,5] years. In Column 3 (4) of Panel C, the maturity bins of [1,3] years and (3,5] years ((7,10] years and >10 years) are grouped together and used as the reference category. The sample period starts on March 16, 2019, one year before the BOJ's first CBPP expansion announcement. It ends on September 30, 2020. Variable definitions are provided in Table 2.3. All the continuous variables are winsorized at the 2.5% and 97.5% levels. Standard errors are reported in parentheses below the estimated coefficients. ***, **, and * indicate that the difference is significant at 1%, 5%, and 10% levels, respectively.

Table 2.7: Effect heterogeneity: Comparing the average marginal effects of Expansion2

	Panel A: Credit rating (N=490)			**Panel B: Market capitalization** (N=483)			**Panel C: Net book leverage** (N=490)		
	AA (1)	A/BBB (2)	Diff. (1) - (2)	$\geq$ median (3)	$<$ median (4)	Diff. (3) - (4)	$\geq$ median (5)	$<$ median (6)	Diff. (5) - (6)
[1,3] Y	0.122*	0.144***	-0.022	0.128**	0.146***	-0.017	0.180***	0.092*	0.088
	(0.070)	(0.041)	(0.081)	(0.052)	(0.049)	(0.072)	(0.052)	(0.049)	(0.071)
(3,5] Y	0.081	0.084	-0.003	0.142**	0.025	0.117	0.076	0.096	-0.020
	(0.078)	(0.056)	(0.096)	(0.064)	(0.066)	(0.091)	(0.061)	(0.068)	(0.090)
(5,7] Y	-0.062	-0.136***	0.074	-0.130***	-0.092**	-0.038	-0.119***	-0.103**	-0.016
	(0.053)	(0.035)	(0.063)	(0.042)	(0.041)	(0.059)	(0.036)	(0.046)	(0.058)
(7,10] Y	0.002	0.021	-0.019	0.026	0.004	0.022	0.033	-0.001	0.034
	(0.080)	(0.055)	(0.097)	(0.063)	(0.066)	(0.091)	(0.063)	(0.065)	(0.090)
>10 Y	-0.143***	-0.114***	-0.029	-0.166***	-0.082**	-0.084	-0.170***	-0.084***	-0.086
	(0.053)	(0.032)	(0.061)	(0.041)	(0.037)	(0.055)	(0.049)	(0.021)	(0.052)

The average marginal effects (AMEs) of *Expansion2* are reported together with standard errors in parentheses. The AMEs are based on the MNLMs where the dependent variable is the categorical variable for the maturity bins and the independent variables are *Expansion1*, *Expansion2*, an indicator variable dividing the sample, and the interaction between *Expansion2* and the indicator variable. In Panel A, the indicator variable takes a value of one if the bond is rated AA. In Panel B (C), the indicator variable takes a value of one if the market capitalization (net book leverage) is equal to or greater than the median at the beginning of the sample period. The AMEs are calculated separately based on the indicator variable values, and the differences in AMEs are also tested. ***, **, and * indicate that the difference is significant at 1%, 5%, and 10% levels, respectively.

Table 2.8: Maturity compositions

Panel A: No. of different maturities	Pre-expansion (3/16/19–3/15/20)	Expansion1 (3/16/20–4/26/20)	Expansion2 (4/27/20–9/30/20)	Total
Single maturity	91	4	23	118
Two maturities	61	5	29	95
Three maturities	28	3	25	56
Four maturities	3	0	0	3
Five maturities	1	0	1	2
Total	184	12	78	274
Multiple maturities prop.	0.505	0.667	0.705	
Diff. vs. Pre: p-value (Fisher's exact test)		0.375	0.004	

Panel B: Characteristics of multiple-maturity issuances

	Pre-expansion (N=93)	Expansion1 (N=8)		Expansion2 (N=55)	
	Mean	Mean	Diff. vs. Pre (p-value)	Mean	Diff. vs. Pre (p-value)
N. of different maturities	2.398	2.375	-0.023 (0.92)	2.509	0.111 (0.28)
Ln(sum of proceeds)	24.217	24.326	0.109 (0.66)	24.353	0.136 (0.22)
Ln(average proceeds)	23.375	23.481	0.106 (0.62)	23.464	0.089 (0.34)
Including [1,3] Y	0.183	0.250	0.067 [0.64]	0.473	0.290 [0.00]
Including (3,5] Y	0.591	0.500	-0.091 [0.72]	0.927	0.336 [0.00]
Including (5,7] Y	0.441	0.125	-0.316 [0.13]	0.164	-0.277 [0.00]
Including (7,10] Y	0.753	1.000	0.247 [0.19]	0.800	0.047 [0.55]
Including >10 Y	0.344	0.500	0.156 [0.45]	0.109	-0.235 [0.00]
Ln(weighted average maturity)	2.158	2.134	-0.024 (0.88)	1.868	-0.290 (0.00)
Weighted SD of maturities	3.181	3.605	0.424 (0.55)	2.549	-0.632 (0.03)
Incl. [1,5] Y and >10 Y	0.505	0.625	0.120 [0.72]	0.745	0.240 [0.01]

Panel C: Maturity compositions of issuances including maturities of [1,5] years and ≥10 years

Maturities (years)	Pre-expansion	Expansion1	Expansion2	Total
5, 10	18	2	15	35
5, 7, 10	14	0	5	19
3, 5, 10	0	1	16	17
5, 10, >10	3	1	2	6
5, >10	3	0	0	3
3, 10	2	0	0	2
3, 10, >10	0	1	1	2
3, 5, 7, 10	2	0	0	2
1.5, 3, 5, 7, 10	0	0	1	1
3, >10	1	0	0	1
3, 5, 7, 10, >10	1	0	0	1
3, 7, 10	1	0	0	1
3, 7, 10, >10	1	0	0	1
4, 10	0	0	1	1
5, 7, >10	1	0	0	1
Total	47	5	41	93

Bonds that have different maturities but were issued by the same company on the same date are treated as part of one unique bond issuance event. Panel A addresses all the sample bond issuances, whereas Panel B considers only multiple-maturity issuances. In Panel B, characteristics of multiple-maturity bond issuances in the expansion1 and expansion2 periods are compared and tested against those of the pre-expansion period. Variable definitions are provided in Table 2.3. For non-binary variables, p-values obtained from t-tests are reported in round brackets. For binary variables, p-values obtained from two-sided Fisher's exact tests are reported in square brackets. All the continuous variables are winsorized at the 2.5% and 97.5% levels. Panel C further limits the sample to multiple-maturity issuances including bonds maturing in five years or less and in 10 years or more.

Table 2.9: Skipping a maturity of seven years: Analyzing multiple-maturity issuances including bonds maturing in [1,5] years and ≥10 years

Dependent variable: Indicator variable of the issuance including a seven-year bond						
Sample:	All			Issues with three maturities		
	(1)	(2)	(3)	(4)	(5)	(6)
Expansion1	-1.666***	-2.004**	-1.957***	-3.473***	-3.475***	-3.514***
	(0.307)	(0.782)	(0.719)	(0.411)	(0.345)	(1.021)
Expansion2	-1.092***	-1.261***	-1.449***	-2.662***	-2.681***	-3.121***
	(0.381)	(0.341)	(0.321)	(0.538)	(0.554)	(0.518)
Ln(weighted average maturity)		0.387	0.325		-0.120	-1.361
		(0.716)	(0.693)		(1.181)	(1.193)
N. of maturities: three or more		2.041***	2.371***			
		(0.313)	(0.341)			
Credit rating: BBB			-0.048			
			(0.579)			
Ln(total assets)			-0.154			-0.580*
			(0.143)			(0.350)
Net book leverage			-2.805**			-4.900**
			(1.154)			(1.959)
Profitability			-8.062			-17.906*
			(5.031)			(9.290)
Asset tangibility			-0.497			0.884
			(0.890)			(1.248)
Industry FE		✓	✓			✓
Observations	93	93	91	46	46	46
McFadden's pseudo R^2	0.026	0.125	0.132	0.128	0.128	0.210
Mean of dep. var.	0.280	0.280	0.264	0.457	0.457	0.457
Average marginal effects						
Expansion2	-0.224***	-0.219***	-0.242***	-0.580***	-0.583***	-0.564***
	(0.082)	(0.058)	(0.053)	(0.088)	(0.092)	(0.086)

In Panel A, the sample includes only multiple-maturity issuances including bonds maturing in [1,5] years and ≥10 years. In Panel B, the sample issuances are further restricted to those whose number of different maturities is three. The dependent variable is an indicator variable that takes a value of one if the issuance includes a bond maturing in (5,7] years. Variable definitions are provided in Table 2.3. The bias-corrected logit models of Kosmidis and Firth (2009) are estimated. Heteroskedasticity-robust standard errors are reported in parentheses. The AMEs of *Expansion*1 are not reported because Stata's `margins` command returns implausibly low standard errors for them in some columns, presumably due to the rarity of the observations falling into this category. All the continuous variables are winsorized at the 2.5% and 97.5% levels. ***, **, and * indicate significance at 1%, 5%, and 10% levels.

Table 2.10: Changes in the determinants of offering yields

	Pre-expansion period	Expansion2 period	Expansion2 vs. Pre-expansion	
			Diff. in coeff.	χ^2 stat. [p-value]
Maturity: (3,5] Y	8.40***	12.90***	4.50	1.07 [0.30]
	(2.01)	(4.16)		
Maturity: (5,7] Y	19.29***	24.72***	5.42	1.00 [0.32]
	(2.53)	(5.19)		
Maturity: (7,10] Y	24.61***	33.38***	8.77**	4.49 [0.03]
	(2.03)	(3.89)		
Maturity: >10 Y	64.34***	66.01***	1.67	0.13 [0.72]
	(2.73)	(4.19)		
Credit rating: A	10.15***	11.45***	1.30	0.10 [0.75]
	(1.74)	(4.01)		
Credit rating: BBB	24.54***	31.13**	6.59	0.24 [0.63]
	(3.22)	(14.10)		
Ln(proceeds)	1.87	-1.55	-3.41	1.19 [0.28]
	(1.62)	(2.90)		
Ln(total assets)	0.25	0.73	0.47	0.04 [0.84]
	(0.63)	(2.47)		
Net book leverage	4.52	23.73**	19.21*	3.15 [0.08]
	(3.16)	(11.11)		
Profitability	-45.62**	-57.83	-12.21	0.03 [0.86]
	(17.67)	(71.89)		
Asset tangibility	-1.92	-18.87***	-16.95**	5.12 [0.02]
	(3.66)	(7.06)		
Industry FE	✓	✓		
N	311	156		
Adjusted R^2	0.822	0.594		
Mean of dep. var.	32.81	31.15		
S.D. of dep. var.	23.72	22.94		

The dependent variable is the offering yield. Variable definitions are provided in Table 2.3. All the continuous variables are winsorized at the 2.5% and 97.5% levels. In the first two columns, heteroskedasticity-robust standard errors are reported in parentheses. The differences in coefficients are tested by the "stacking" method of Weesie et al. (2000). ***, **, and * indicate that the difference is significant at 1%, 5%, and 10% levels, respectively.

Appendices

Appendix 2.A: A brief history of the BOJ's corporate bond purchases

While the BOJ pioneered quantitative easing in March 2001 as an alternative monetary policy tool under close-to-zero interest rates, the bank did not take a further step of purchasing corporate bonds until the global financial crisis occurred. The possibility of purchasing short-term corporate debt was first mentioned in an announcement in December 2008.[35] The bank then began purchasing CP in the next month,[36] and in March 2009, corporate bonds under the following two key eligibility conditions: remaining maturities of one year or less and credit ratings of A or above.[37] The BOJ ended the program at the end of the year, stating that "issuing conditions in the CP and corporate bond markets have been improving markedly."[38]

The Comprehensive Monetary Easing (CME) program announced in October 2010 marked the resumption of the bank's corporate bond purchases. One of the pillars of the CME was the creation of the Asset Purchase Program, through which the BOJ could purchase a wide range of assets such as government bonds, CP, corporate bonds, and equity ETFs. Corporate bonds rated as investment grade and maturing from one to two years were eligible.[39]

The CME was then supplanted by a new program, the Quantitative and Qualitative Monetary Easing (QQME), which was unveiled on April 4, 2013. The newly appointed BOJ Governor Haruhiko Kuroda explained that the QQME was based on his belief that the BOJ "should do whatever is necessary to overcome deflation" and " should make all-out efforts to utilize every possible resource bestowed upon the Bank."[40] As part of the change, the BOJ was enabled to

[35] Bank of Japan (2008, December 19) On Monetary Policy Decisions. Retrieved from `https://www.boj.or.jp/en/announcements/release_2008/k081219.pdf`.

[36] Bank of Japan (2009, January 22) Outright Purchases of Corporate Financing Instruments. Retrieved from `https://www.boj.or.jp/en/announcements/release_2009/un0901b.pdf`.

[37] Bank of Japan (2009, February 19) Statement on Monetary Policy. Retrieved from `https://www.boj.or.jp/en/announcements/release_2009/k090219.pdf`.

[38] Bank of Japan (2009, October 30) Statement on Monetary Policy. Retrieved from `https://www.boj.or.jp/en/announcements/release_2009/k091030.pdf`.

[39] Bank of Japan (2010, October 28) Statement on Monetary Policy. Retrieved from `https://www.boj.or.jp/en/announcements/release_2010/k101028.pdf`.

[40] Haruhiko Kuroda (2013, April 12) Quantitative and Qualitative Monetary Easing. *Speech at a Meeting Held by the Yomiuri International Economic Society in Tokyo.* Retrieved from `https://www.boj.or.jp/en/announcements/`

buy investment-grade bonds with remaining maturities of from one to three years. This regime continued until the BOJ's first response to the COVID-19 pandemic on March 16, 2020 (Section 2.2.1).

Appendix 2.B: Institutional details of the BOJ's corporate bond purchases

Central banks' methods of large-scale government and corporate bond purchases

There are two major forms through which central banks conduct large-scale purchases of already issued government/corporate bonds: reverse auctions and bilateral purchases in the secondary market. A typical reverse auction goes as follows: First, the central bank announces the expected purchase amount and target securities in advance. Then, on the auction day, auction participants submit offers of price-quantity pairs of auction-eligible securities that they intend to sell. Lastly, the central bank decides which offers to accept based on its algorithm that may or may not be publicly disclosed. Reverse auctions are commonly used as a means of repurchasing government bonds. They have been used in QE operations of the BOJ and other major central banks, such as the Fed and the BoE. Alternatively, a central bank may choose to purchase securities directly from the secondary market. The Eurosystem has mainly employed this approach to purchase bonds issued by Eurozone countries.

The BOJ has also used reverse auctions to purchase (already issued) corporate bonds, as did the BoE in its Corporate Bond Purchase Scheme (CBPS), which was initiated in 2016. On the other hand, when the Fed started its first-ever corporate bond purchases in response to the pandemic, it hired BlackRock to purchase corporate bonds on its behalf in both the primary and secondary markets.

Auction design of the BOJ's corporate bond purchases

The BOJ's corporate bond reverse auctions are structured in the following manner. First, normally, on the last or the second last business day of a month, the bank discloses information regarding

press/koen_2013/data/ko130412a1.pdf.

planned auctions for the subsequent two months. During the sample period, the auctions were held once a month, except for April 2020, when two auctions were held. As already mentioned, the bank announced on April 27, 2020 that the maximum eligible remaining maturity was extended from three years to five years. Following the announcement, the BOJ started to hold separate auctions of bonds with remaining maturities of [1,3] years and for those with remaining maturities of (3,5] years. Each has been held once a month since May 2020. The pre-auction announcements include the auction dates and each auction's expected purchase amount.

Then, on the day of the auction, participants submit multiple offers, each of which is a pair of yield (price) and quantity for a specific eligible corporate bond. The auction participants are financial firms pre-screened by the BOJ. As of October 2020, there were 36 participating financial firms.[41] Each auction accepts all eligible corporate bonds, as long as they satisfy the auction's target remaining maturity range (i.e., [1,3] or (3,5] years). Note that the auctions are price-discriminatory ones, meaning that winning offers are executed at their own offer yields (prices).

Finally, the BOJ determines winning offers. The multi-good nature of these auctions means that the BOJ needs an algorithm to compare offers of different bonds. One natural method would be to rank them based on their yield "concessions" (Breedon, 2018), which are differences between the offer yields and the market yields of the associated securities. The government bond reverse auctions organized by the BOJ and the BoE employ this approach to compare offers of different (eligible) bonds. In the case of the Fed's QE reverse auctions, the offer yields are compared with the yields implied by its (confidential) yield curve model (Song and Zhu, 2018). The BoE's CBPS takes a somewhat similar approach, in the sense that it sets a "reserve spread" for each eligible corporate bond and the BoE considers the differences between offer and reserve spreads in determining which offers to accept (Boneva et al., 2022).

Nevertheless, the BOJ states that it compares only unadjusted offer yields in its reverse auctions of corporate bonds. Specifically, the BOJ states that it "accepts bids by starting with the highest desired yield and continuing down so that the amount purchased per single issuer's CP and corporate bonds remains within the unused purchase value," while "the Bank reserves the right to

[41] https://www.boj.or.jp/mopo/measures/select/opelist01.pdf

reject all or some of bids which a counterparty submits when deemed appropriate."[42] Therefore, taken at face value, the BOJ's protocol implies that bonds issued by relatively riskier firms, to the extent that they inherently bear higher yields, are preferred among those issued by eligible firms. This is a unique (and often overlooked) feature of the BOJ's corporate bond auctions.

Auction concessions and the "BOJ trade"

It has been argued that the BOJ's massive purchases of JGBs and corporate bonds have unintentionally distorted the Japanese financial market. At the core of this criticism is the existence of the so-called "BOJ trade." According to the financial press, this trading strategy aims to profit from yield concessions in the BOJ's reverse auctions, by purchasing assets targeted by BOJ's purchase operations in the primary or secondary market to sell them to the BOJ later. In the case of the BoE's reverse auctions of government bonds, significant yield concessions are documented by Breedon (2018).[43] It has been claimed that the BOJ trade has played an important role in the JGB market[44] and the corporate bond market (Tomisawa and Hazama, 2019).

To grasp the nature of the BOJ's corporate bond auctions, Table A2.1 summarizes the key information released by the BOJ regarding its reverse auctions of corporate bonds from January to September 2020. Notably, following the first CBPP expansion announcement on March 16, 2020, the BOJ held two auctions of bonds with remaining maturities of [1,3] years in April 2020. Then, as a result of the second expansion announcement on April 27, the auctions of bonds with remaining maturities of (3,5] years started in the next month.

Although only pre-screened financial firms can join the reverse auctions, other investors can sell bonds to the central bank indirectly through the auction participants. Using transaction-level data, Boneva et al. (2022) show that when the BoE's CBPS purchases corporate bonds, there are large increases in the net purchases by dealers, most of whom can participate in the reverse auctions, and in the net selling by other active bond investors such as insurance companies. Boneva et al. (2022,

[42] `https://www.boj.or.jp/en/mopo/measures/mkt_ope/ope_s/opetori19.htm/`

[43] I am not aware of any study that systematically analyzes yield concessions in the BOJ's reverse auctions of JGBs.

[44] Eleanor Warnock (2014, April 14) Massive BOJ-Buying Silences JGB Market. *The Wall Street Journal.* Retrieved from `https://www.wsj.com/articles/BL-MBB-19545`.

p. 10) conclude, "The aggregate quantity of bonds sold by insurance companies and asset managers during the purchase period suggests that these investors were the ultimate sellers of around half of the bonds bought by the BoE, with the remainder coming from dealers balance sheets."

Because the BOJ has not disclosed the identities of the purchased corporate bonds,[45] it is impossible to precisely measure yield concessions, and therefore the profit opportunity of the BOJ trade. Nevertheless, the available data strongly suggest the existence of yield concessions. First, I compare the average winning offer yields (Column (b)) and the average market yields of corporate bonds with eligible credit ratings and remaining maturities (Column (g)). The third last column shows that the average winning offer yields were always lower than the average market yields of eligible corporate bonds. Second, the average winning offer yields are further compared with the average market yields of AAA-AA rated bonds with eligible remaining maturities (Column (h)). Note that it is *unlikely* that this set of bonds provides appropriate benchmark market yields for bonds purchased by the BOJ. This is because the auction mechanism implies that higher-yield (i.e., riskier) bonds are purchased first, as mentioned earlier. The reason why AAA-AA rated bonds are examined is to provide reasonable lower bounds of corresponding market yields. The last column of Panel A shows that in the case of [1,3]-year bond auctions the average winning offer yields were even lower than the average market yields for the least risky bonds.

In addition, Table A2.1 indicates strikingly high yield concessions for reverse auctions of bonds maturing in [1,3] years after the BOJ's CBPP expansions. The average (lowest) winning offer yields had been *negative* since June (May) 2020. In June 2020, presumably as a result of the extremely low yields, the BOJ started to publish the lower limits that it set on offer yields for this shorter maturity category. The financial press explained that as the lower limit in June 2020 (-0.14%) was on par with the market yield of JGBs with comparable remaining maturities, the BOJ intended to signal not to buy corporate bonds at yields lower than the comparable JGB yields.[46]

[45] The BoE also does not publicly disclose information regarding individual corporate bonds that it has purchased in its reverse auctions (Boneva et al., 2022). It, however, discloses its sector-level corporate bond holding data, while the BOJ does not do so.

[46] *Nichigin, shasai ope de hatsu no kagen rimawari settei* (The BOJ set the lower limit yield in its corporate bond operations for the first time) (in Japanese) (2020, June 9) *Nikkei.* Retrieved from `https://www.nikkei.com/article/DGXMZO60155990Z00C20A6EN2000/`.

Table A2.1: Corporate bond operations of the BOJ, January–September 2020

Panel A: Remaining maturities: [1,3] years

									Eligible bonds's econdary-market yields			
									AAA-BBB bonds		AAA-AA bonds	
		(a)	(b)	(c)		(d)	(e)	(f)	(g)		(h)	
Auction date	Settle. date	Planned purchase amount (bill. ¥)	Offer amount (bill. ¥)	Purchase amount (bill. ¥)	Offer-to-cover ratio: (b) / (c)	Average winning yield (%)	Lowest winning yield (%)	Lower limit on yields (%)	Average yield (%)	Diff: (d) - (g)	Average yield (%)	Diff: (d) - (h)
1/23/2020	1/29/2020	100.0	315.1	100.0	3.15	0.064	0.055		0.142	-0.078	0.087	-0.023
2/20/2020	2/27/2020	100.0	254.1	100.0	2.54	0.053	0.033		0.122	-0.069	0.073	-0.020
3/23/2020	3/27/2020	200.0	370.8	200.1	1.85	0.113	0.082		0.149	-0.036	0.096	0.017
4/7/2020	4/13/2020	150.0	283.3	150.0	1.89	0.117	0.082		0.216	-0.099	0.139	-0.022
4/20/2020	4/24/2020	150.0	265.1	150.0	1.77	0.091	0.040		0.215	-0.124	0.135	-0.044
5/8/2020	5/14/2020	300.0	313.7	300.0	1.05	0.012	-0.120		0.215	-0.203	0.133	-0.121
6/9/2020	6/15/2020	300.0	375.7	300.3	1.25	-0.114	-0.140	-0.140	0.208	-0.322	0.125	-0.239
7/3/2020	7/9/2020	300.0	616.2	300.0	2.05	-0.073	-0.092	-0.140	0.189	-0.262	0.110	-0.183
8/5/2020	8/7/2020	300.0	788.6	300.0	2.63	-0.034	-0.044	-0.130	0.187	-0.221	0.107	-0.141
9/9/2020	9/11/2020	300.0	542.7	300.1	1.81	-0.037	-0.053	-0.120	0.194	-0.231	0.112	-0.149

Panel B: Remaining maturities: (3,5] years

									Eligible bonds' secondary-market yields			
									AAA-BBB bonds		AAA-AA bonds	
		(a)	(b)	(c)		(d)	(e)	(f)	(g)		(h)	
Auction date	Settle. date	Planned purchase amount (bill. ¥)	Offer amount (bill. ¥)	Purchase amount (bill. ¥)	Offer-to-cover ratio: (b) / (c)	Average winning yield (%)	Lowest winning yield (%)	Lower limit on yields (%)	Average yield (%)	Diff: (d) - (g)	Average yield (%)	Diff: (d) - (h)
5/20/2020	5/26/2020	200.0	592.0	200.1	2.96	0.212	0.170		0.289	-0.077	0.183	0.029
6/23/2020	6/29/2020	200.0	617.7	200.0	3.09	0.192	0.170		0.287	-0.095	0.176	0.016
7/22/2020	7/28/2020	200.0	487.4	200.0	2.44	0.173	0.154		0.261	-0.088	0.153	0.020
8/21/2020	8/25/2020	200.0	471.6	200.0	2.36	0.181	0.165		0.283	-0.102	0.169	0.012
9/23/2020	9/25/2020	200.0	345.1	200.0	1.73	0.154	0.133		0.251	-0.097	0.140	0.014

This table summarizes the outcomes of the BOJ's reverse auctions of corporate bonds for the period from January to September 2020. The monetary unit is billion yen for Columns (a)–(c). The auction information is disclosed by the BOJ at `https://www.boj.or.jp/en/statistics/boj/fm/ope/index.htm/`. To compare the winning offer yields with the yields of comparable corporate bonds in the secondary market, I use the "rating matrix data" released by the JSDA. The data provide the average yield for each credit rating-remaining maturity pair on a given date. I rely on credit ratings provided by R&I. I retain only the credit rating-remaining maturity pairs whose remaining maturities correspond with the BOJ's target remaining maturities (i.e., [1,3] and (3,5] years). Column (g) ((h)) uses bonds rated at AAA–BBB (AAA–AA). I calculate the weighted average yields of these bonds on the auction dates, with the weight being the number of bonds of each credit rating-remaining maturity pair. Note that the JSDA's rating matrix data present the average yield for each of the remaining maturity bins of $[t,\ t+1)$ years, where t is 2, …, 19. On the other hand, bonds with remaining maturities of less than two years are put into one group (i.e., (0,2) years). Thus, the average yields for bonds with remaining maturities of [1,2) years were linearly interpolated, under the assumption that the numbers of bonds whose remaining maturities are (0,1) and [1,2) years are the same.

Appendix 2.C: Government debt supply during the COVID-19 pandemic in Japan

This paper argues that the BOJ's CBPP expansion induced disproportionate changes in firms' maturity choices around the eligibility threshold (five years). Nevertheless, one concern is possible contemporaneous shocks to government debt supply. This is because Greenwood et al. (2010), Badoer and James (2016), and Lugo and Piccillo (2019) suggest that some preferred-habitat investors regard government and corporate debt as substitutes. Section 2.5.1 of the main text notes that the following two types of shocks to government debt supply can generate the main finding of this paper: (a) an increase in issuances of JGBs with maturities of close to seven years and (b) a decrease in the BOJ's purchases of JGBs with remaining maturities of about seven years. In this section, I show that data on shocks to JGB supply—the BOJ's quantitative easing operations and JGB issuances—suggest that they did not drive the main finding of this paper.

The BOJ's quantitative easing operations

The BOJ's purchase amounts of JGBs with remaining maturities of close to seven years were relatively quite limited both before and after the COVID-19 crisis. Therefore, the BOJ's quantitative easing operations could not have largely changed the supply of JGBs in this maturity segment. The time series of the BOJ's JGB purchases is presented in Figure A2.1. The data source is the BOJ's historical JGB holding data.[47] During the sample period, the BOJ disclosed its most recent JGB holding data three times a month, with the last holding data date corresponding to the end of the month. I calculated the BOJ's monthly purchase amounts from the changes in its historical JGB holding amounts. Figure A2.1 aggregates the BOJ's JGB purchases at the level of the following remaining maturity bins: (0,1], (1,3], (3,5], (5,7], (7,10], and >10 years.[48] On the other hand, Figure A2.2 gives a more detailed breakdown of purchase amounts by remaining maturity. Note that in Figure A2.2 the BOJ's monthly purchase data are divided into three periods that roughly correspond to the pre-expansion, expansion1, and expansion2 periods, which are defined in the

[47] https://www.boj.or.jp/en/statistics/boj/other/mei/release/index.htm/

[48] Bonds in the bin of (0,1] years tend to have remaining maturities very close to one year.

main text. It is apparent from Figure A2.2 that the BOJ's quantitative easing operations did not mainly target JGBs with remaining maturities of between 5 and 8 years throughout the sample period.

Issuances of JGBs

It is also improbable that JGB issuance shocks drove the main finding of this paper. Note that there are two main channels through which the MOF issues JGBs: auctions for new bonds and reopenings and liquidity enhancement auctions. To begin with, no JGBs have original maturities of between 5 and 10 years; they are 2, 5, 10, 20, 30, and 40 years.[49] Therefore, there was no chance that auctions for new bonds and reopenings suddenly increased the supply of JGBs maturing in around seven years. The top graph of Figure A2.4 presents JGB issuance amounts via auctions for new bonds and reopenings.

JGBs are also issued via liquidity enhancement auctions (LEAs), in which the MOF reissues bonds that were issued typically at least years ago and whose liquidity has declined due to their limited supply in the secondary market Hattori (2019).[50] JGBs issued via LEAs tend to have remaining maturities markedly different from their original maturities, meaning that LEAs can suddenly increase the supply of JGBs maturing in the region of seven years. The data suggest, however, that LEAs did not cause this type of supply shock with a magnitude sufficient to generate my main finding. The bottom graph of Figure A2.4 documents JGB issuance amounts via LEAs. There are three key observations. First, the issuance amounts of JGBs via LEAs were much smaller than those via auctions of new bonds and reopenings. Second, although the JGB issuance amount via auctions of new bonds and reopenings shifted up in July 2020, that via LEAs did not—it rather slightly scaled down. Third, the composition of the remaining maturities of JGBs issued via LEAs was not stable.[51] Figure A2.5 presents the average monthly JGB issuance amounts via LEAs for

[49]The MOF also issues short-term government debt called treasury discount bills or "T-Bills." In my sample period, T-Bills issued had original maturities of around 3 months, 6 months, or 1 year. Figure A2.3 shows that T-Bill issuances dramatically increased to finance fiscal measures in response to the COVID-19 crisis.

[50]During my sample period, LEAs were held twice a month.

[51]This observation aligns with the stated objective of LEAs, i.e., improving liquidity of JGBs with diminished liquidity. To this end, the LEAs' JGB selection should be based primarily on liquidity and not other characteristics such as remaining maturities.

the three sub-periods same as those of Figure A2.2. On the one hand, Figure A2.5 shows that the average monthly reissuance amount of JGBs with remaining maturities of (6,7] years largely increased: from ¥63.6 billion in the pre-expansion months to ¥165.6 billion in the expansion2 months. On the other hand, the amounts of its neighboring remaining maturity bins *decreased* at the same time. As a result, the average monthly reissuance amount of JGBs with remaining maturities of (5,8] years increased only marginally: from ¥282.6 billion to ¥310.6 billion. Therefore, LEAs also did not cause a supply shock that seems to be large enough to drive the main finding of this paper.

Figure A2.1: The BOJ's new Japanese government bond (JGB) purchase amounts

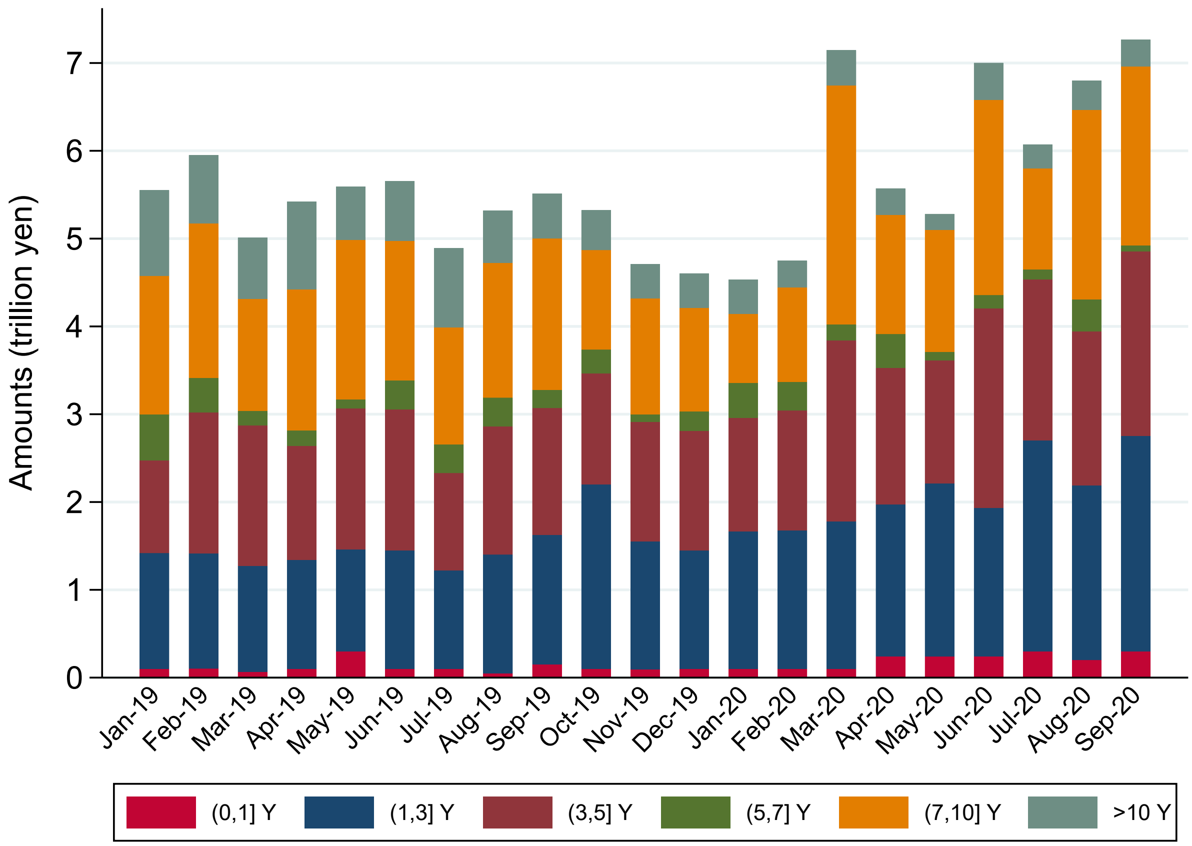

This figure shows the amounts of the BOJ's new JGB purchases (in trillions of yen) on a monthly basis from January 2019 to September 2020. The purchase amounts are categorized by the remaining maturities of the JGBs purchased. (JGBs have original maturities of 2, 5, 10, 20, 30, and 40 years.) The purchase amounts are calculated from the changes in the BOJ's historical JGB holding data available at `https://www.boj.or.jp/en/statistics/boj/other/mei/release/index.htm/`. Floating-rate JGBs and inflation-indexed JGBs are excluded. Also, note that the BOJ's purchases of treasury discount bills (T-Bills), which have original maturities of one year or less, are not included.

Figure A2.2: The BOJ's average monthly JGB purchase amounts by remaining maturity

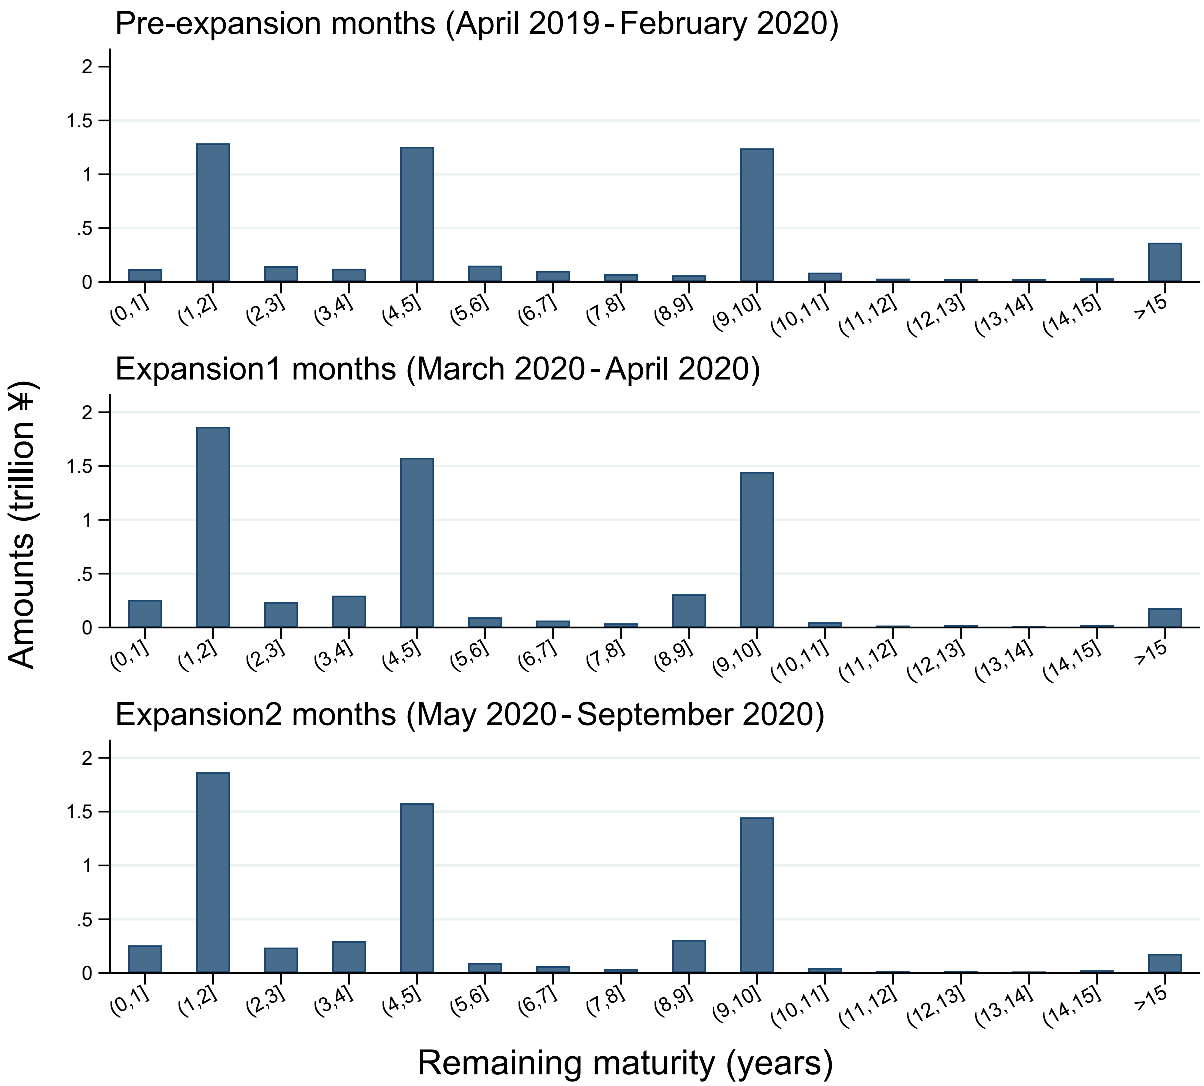

This figure shows the BOJ's average monthly JGB purchase amounts (in trillions of yen) by remaining maturity. The sample period is divided into three sets of months that roughly correspond to the pre-expansion, expansion1, and expansion2 periods. The purchase amounts are calculated from the changes in the BOJ's historical JGB holding data. The sample requirements are the same as those of Figure A2.1.

Figure A2.3: Issuance amounts of treasury discount bills (T-Bill)

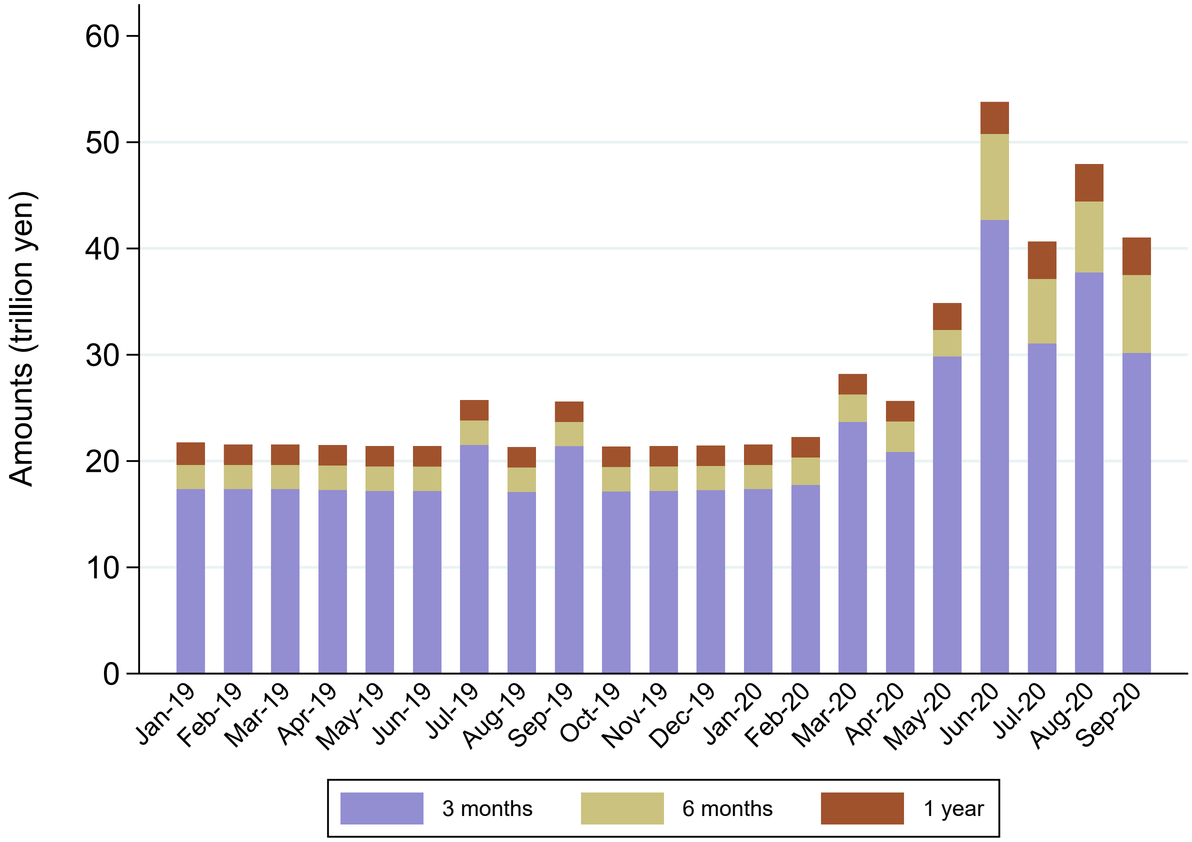

This figure shows the amounts of treasury discount bill (T-Bill) issuances (in trillions of yen) on a monthly basis from January 2019 to September 2020. The data were obtained from the MOF's website: `https://www.mof.go.jp/jgbs/reference/appendix/index.htm`.

Figure A2.4: Issuance amounts of JGBs

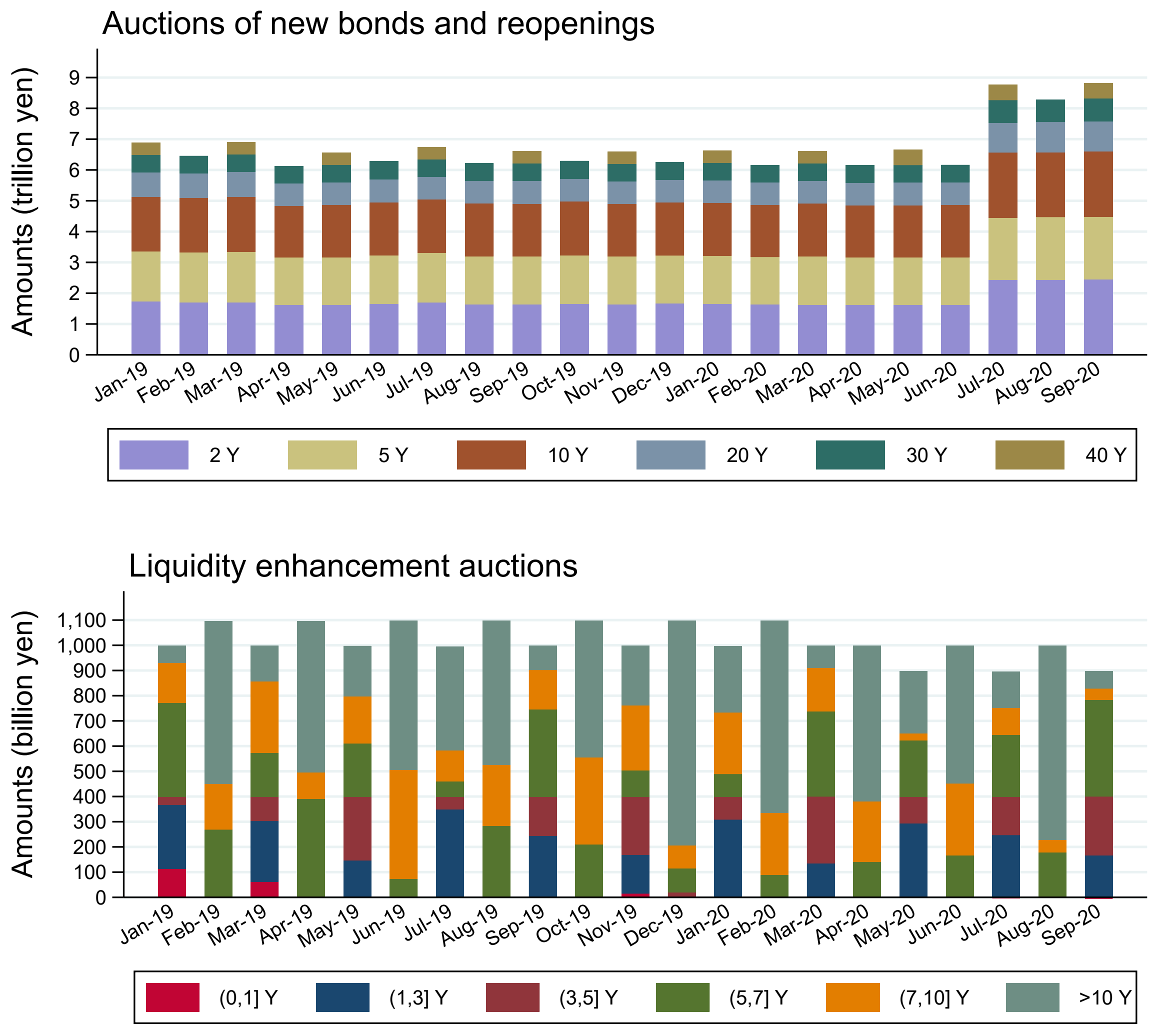

This figure shows the amounts of JGB issuances on a monthly basis from January 2019 to September 2020. The data were obtained from the MOF's website: `https://www.mof.go.jp/jgbs/reference/appendix/index.htm`. Floating-rate JGBs and inflation-indexed JGBs are excluded. The top graph shows the amounts issued via auctions of new JGBs and reopenings (in trillions of yen). JGBs have original maturities of 2, 5, 10, 20, 30, and 40 years. The amounts of JGBs issued via liquidity enhancement auctions (in billions of yen) are shown in the bottom graph. Here, the amounts are based on the face value and they are classified by remaining maturity bins.

Figure A2.5: Average monthly JGB issuance amounts via liquidity enhancement auctions by remaining maturity

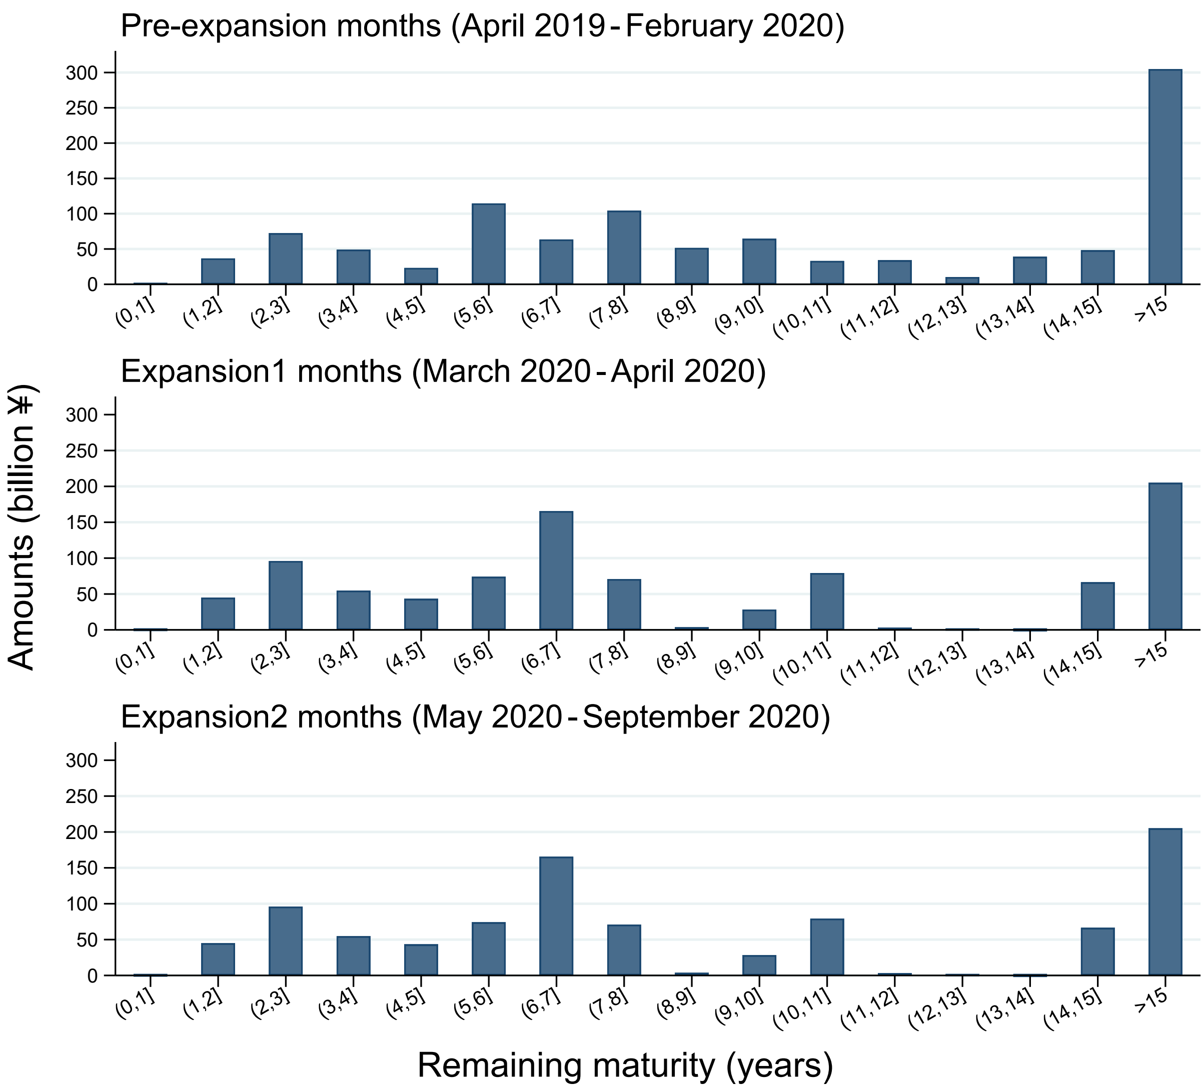

This figure shows the average monthly JGB issuance amounts via liquidity enhancement auctions (in billions of yen) by remaining maturity. The sample period is divided into three sets of months that roughly correspond to the pre-expansion, expansion1, and expansion2 periods. The sample is the same as that of the bottom graph of Figure A2.4. The issuance amounts are based on the face value.

Appendix 2.D: Supplementary results

Supplementary results are presented in this section.

Figure A2.6: Maturity distribution in the U.S.

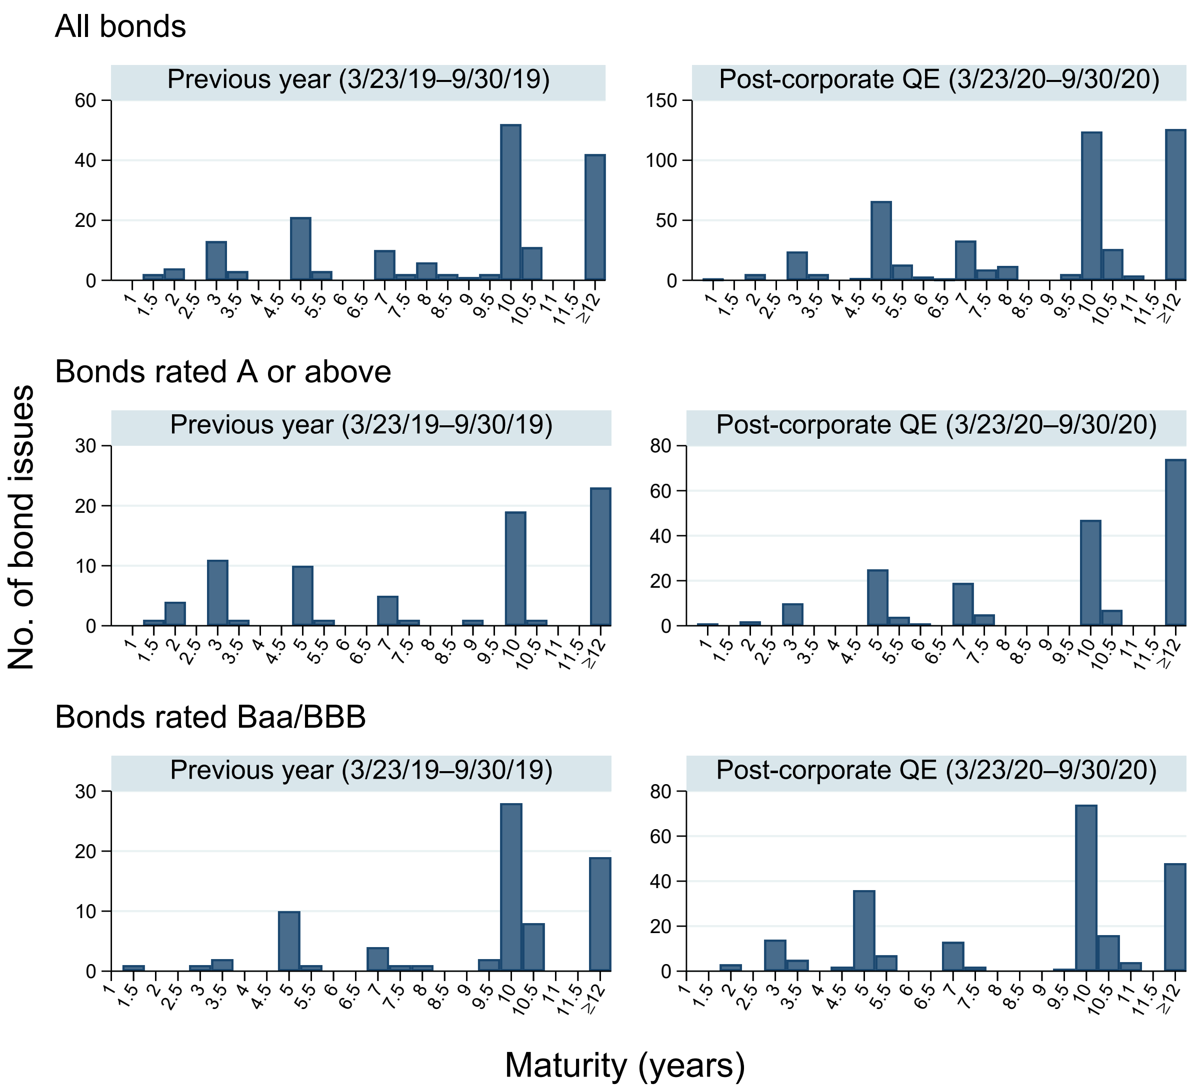

This figure shows the maturity distributions of corporate straight bond issues in the U.S. The issuers are public firms in non-financial and non-utility industries. Moody's bond ratings are used. The data were obtained from the SDC Platinum and Compustat.

Table A2.2: Descriptive statistics: Alternative samples

Panel A: Longer sample period (10/1/16–9/30/20)	Mean	S.d.	Min.	Median	Max.	N
Ln(maturity)	2.014	0.540	1.099	1.946	2.996	1107
Ln(proceeds)	23.266	0.580	22.333	23.026	24.818	1107
Offering yield (bps)	36.602	26.065	1.000	29.500	110.000	1107
Offering spread (bps)	37.904	15.338	15.400	35.100	88.500	1107
Credit rating: BBB	0.083	0.276	0.000	0.000	1.000	1107
Ln(total assets)	27.826	1.104	25.536	27.739	30.095	1089
Net book leverage	0.217	0.215	-0.196	0.217	0.723	1089
Profitability	0.087	0.036	0.026	0.082	0.177	1089
Asset tangibility	0.383	0.218	0.027	0.353	0.830	1089
Expansion1	0.021	0.143	0.000	0.000	1.000	1107
Expansion2	0.145	0.353	0.000	0.000	1.000	1107

Panel B: Including financial firms and utilities	Mean	S.d.	Min.	Median	Max.	N
Ln(maturity)	2.023	0.629	1.099	1.946	3.401	735
Ln(proceeds)	23.376	0.599	22.333	23.026	24.818	735
Offering yield (bps)	35.701	24.901	2.000	29.000	100.000	735
Offering spread (bps)	42.921	18.179	17.300	40.300	115.000	735
Credit rating: BBB	0.046	0.210	0.000	0.000	1.000	735
Ln(total assets)	28.225	1.181	25.638	28.308	30.665	726
Profitability	0.078	0.037	0.008	0.076	0.169	726
Asset tangibility	0.397	0.258	0.005	0.350	0.832	726
Expansion1	0.046	0.210	0.000	0.000	1.000	735
Expansion2	0.310	0.463	0.000	0.000	1.000	735
Industry: Financials	0.140	0.347	0.000	0.000	1.000	735
Industry: Utilities	0.181	0.385	0.000	0.000	1.000	735

Variable definitions are provided in Table 2.3. All variables are winsorized at the 2.5% and 97.5% levels.

Table A2.3: Multinomial logit regressions of corporate bond maturity choices: Longer sample period with month-by-year fixed effects

Panel A: Coefficients

	Model 1: Ref. cat. = (7,10] Y				Model 2: Ref. cat. = (3,5] Y	
	[1,3] Y	(3,5] Y	(5,7] Y	≥10 Y	(1,3] Y	(5,7] Y
Expansion1	-0.448	-0.399	-1.354	-0.097	-0.048	-0.954
	(0.965)	(0.727)	(1.012)	(0.752)	(1.029)	(1.082)
Expansion2	1.015***	0.382	-0.901**	-1.445***	0.633*	-1.283***
	(0.324)	(0.254)	(0.382)	(0.421)	(0.323)	(0.383)
Credit rating: BBB	0.396	0.586*	0.742**	-1.858***	-0.190	0.156
	(0.449)	(0.314)	(0.363)	(0.637)	(0.433)	(0.342)
Ln(total assets)	0.210*	-0.140	0.040	0.348***	0.350***	0.180*
	(0.118)	(0.085)	(0.099)	(0.115)	(0.118)	(0.100)
Net book leverage	2.399***	0.222	-0.925	0.627	2.177***	-1.147*
	(0.743)	(0.556)	(0.689)	(0.832)	(0.711)	(0.678)
Profitability	5.894	1.217	-1.111	-4.907	4.677	-2.327
	(4.449)	(3.060)	(3.547)	(4.671)	(4.390)	(3.542)
Asset tangibility	-1.029	-1.611***	0.287	3.151***	0.582	1.898***
	(0.661)	(0.542)	(0.657)	(0.674)	(0.664)	(0.669)
Industry FE	✓	✓	✓	✓	✓	✓
Month-by-year FE	✓	✓	✓	✓	✓	✓
N	1089				1089	
Pseudo R^2	0.108				0.108	

Panel B: Average marginal effects

	[1,3] Y	(3,5] Y	(5,7] Y	(7,10] Y	≥10 Y
Expansion1	-0.009	-0.014	-0.117*	0.110	0.030
	(0.063)	(0.117)	(0.061)	(0.129)	(0.089)
Expansion2	0.133***	0.099**	-0.110***	0.000	-0.122***
	(0.041)	(0.047)	(0.027)	(0.044)	(0.020)

This table reports the results of multinomial logit regressions using the bias-correction methods of Kosmidis and Firth (2011). The sample consists of corporate straight bonds issued for the period from October 1, 2016 to September 30, 2020. The dependent variable is a bond maturity category and the reference category is the maturity bin of (7, 10] years. Variable definitions are provided in Table 2.3. All the continuous variables are winsorized at the 2.5% and 97.5% levels. Standard errors are reported in parentheses below the estimated coefficients. ***, **, and * indicate that the difference is significant at 1%, 5%, and 10% levels, respectively.

Table A2.4: Multinomial logit: Including financial firms and utilities

Panel A: Coefficients						
	Model 1: Ref. cat. = (7,10] Y				Model 2: Ref. cat. = (3,5] Y	
	(1,3] Y	(3,5] Y	(5,7] Y	≥10 Y	(1,3] Y	(5,7] Y
Expansion1	0.032	-0.973	-1.628*	-0.463	1.005	-0.656
	(0.552)	(0.596)	(0.898)	(0.483)	(0.672)	(0.987)
Expansion2	0.705***	0.326	-1.069***	-0.615**	0.379	-1.395***
	(0.264)	(0.224)	(0.354)	(0.278)	(0.258)	(0.355)
Credit rating: BBB	0.684	0.866*	0.388	-0.877	-0.182	-0.479
	(0.606)	(0.498)	(0.664)	(0.966)	(0.525)	(0.605)
Ln(total assets)	-0.003	-0.080	0.004	0.148	0.077	0.083
	(0.122)	(0.100)	(0.127)	(0.131)	(0.117)	(0.126)
Profitability	-5.505	-0.608	-0.689	-2.619	-4.897	-0.080
	(4.616)	(3.625)	(4.583)	(4.700)	(4.572)	(4.617)
Asset tangibility	0.314	-1.024	0.472	2.090***	1.338*	1.496*
	(0.759)	(0.651)	(0.805)	(0.731)	(0.782)	(0.835)
Industry: Technology	1.697***	0.926*	1.052	-0.408	0.771	0.126
	(0.607)	(0.558)	(0.689)	(0.984)	(0.536)	(0.630)
Industry: Consumer Staples	0.999	0.208	1.284*	0.549	0.791	1.076
	(0.853)	(0.722)	(0.763)	(0.857)	(0.849)	(0.771)
Industry: Health Care	0.495	0.143	0.210	-0.833	0.352	0.067
	(1.168)	(0.903)	(1.176)	(1.768)	(1.168)	(1.184)
Industry: Financials	1.623***	0.957**	0.749	-0.631	0.666	-0.208
	(0.515)	(0.448)	(0.586)	(0.734)	(0.477)	(0.559)
Industry: Real Estate	-0.242	0.456	0.130	0.409	-0.697	-0.326
	(0.711)	(0.554)	(0.683)	(0.500)	(0.741)	(0.716)
Industry: Consumer Discretionary	-0.000	0.311	0.601	-0.401	-0.311	0.290
	(0.488)	(0.359)	(0.447)	(0.489)	(0.489)	(0.452)
Industry: Telecommunications	0.227	-0.051	0.676	0.186	0.278	0.728
	(0.735)	(0.606)	(0.695)	(0.734)	(0.750)	(0.726)
Industry: Basic Materials	-0.499	0.162	0.033	-0.545	-0.662	-0.129
	(0.507)	(0.332)	(0.447)	(0.466)	(0.513)	(0.456)
Industry: Energy	-1.353	0.307	0.426	-1.358	-1.660	0.119
	(1.806)	(0.921)	(1.200)	(1.785)	(1.805)	(1.211)
Industry: Utilities	-0.164	-0.555	-0.928	0.141	0.391	-0.372
	(0.472)	(0.455)	(0.581)	(0.391)	(0.540)	(0.642)
N	726				726	
Pseudo R^2	0.110				0.110	

Panel B: Average marginal effects					
	(1,3] Y	(3,5] Y	(5,7] Y	(7,10] Y	≥10 Y
Expansion1	0.070	-0.099	-0.106**	0.148*	-0.014
	(0.067)	(0.067)	(0.042)	(0.088)	(0.060)
Expansion2	0.102***	0.076**	-0.102***	0.006	-0.082***
	(0.031)	(0.036)	(0.022)	(0.036)	(0.027)

This table reports the results of multinomial logit regressions using the bias-correction methods of Kosmidis and Firth (2011). An R package `brglm2` (Kosmidis, 2020) is used for the estimation. The sample includes financial firms and utilities. Industry dummies based on the Industrial Classification Benchmark (ICB) are included with the reference category being Industrials. All variables are winsorized at the 2.5% and 97.5% levels.

Table A2.5 performs an analysis similar to Table 2.7. The difference is that Table A2.5 divides the sample firms by their reliance on bank debt. Specifically, I use *bank debt ratio*, the ratio of bank debt to total debt (measured at the beginning of the sample period), to split the sample.

The data source of *bank debt ratio* is the Capital IQ Capital Structure Summary database (accessed via WRDS). This database contains two items that can be used to construct *bank debt ratio*. Specifically, *TotBankDbt* measures the amount of total bank debt and *TotBankDbtPct* measures the percentage of total bank debt over total debt. Note that Capital IQ sometimes reports a non-missing value for only one of them. Therefore, I first looked at *TotBankDbtPct* and if it was missing, checked whether *bank debt ratio* could be computed by dividing *TotBankDbt* by total debt, which is the sum of *PrincipalAmtDbtOutstanding*, *TotAdjustments*, and *UnamortizedPremiumTot*.

Capital IQ provides historical snapshots of detailed capital structure information and normally, there exist multiple snapshots available for each fiscal year. I constructed *bank debt ratio* in the following manner. First, I downloaded all snapshots for the one-year period leading up to the sample start date (March 16, 2019). Second, I dropped snapshots if their *TotBankDbtPct* and *TotBankDbtPct* were both missing. Third, because it is not rare that the ratio of bank debt computed from this database slightly exceeds 100%, the following modifications were made: If the ratio was greater than 100% but not more than 110%, I retained the data point and treated it as 100%. In contrast, I dropped snapshots with a ratio of bank debt greater than 110%. Lastly, *bank debt ratio* was measured as the ratio of the bank debt in the snapshot whose *PeriodEndDate* (fiscal period end date) was the closest to the beginning of the sample period (March 16, 2019) among the remaining snapshots.

Bank debt ratio is missing for two sample firms that issued eight bonds in total during the sample period. This is because Capital IQ records neither *TotBankDbt* nor *TotBankDbtPct* for these firms during the one-year period preceding the sample start date. Panel A of Table A2.5 reports descriptive statistics of *bank debt ratio*. The mean (median) is 63.9% (68.8%). The median value is used to divide the sample.

Panel B of Table A2.5 fails to find any statistically significant differences in the AMEs of

Expansion2 between the two sub-samples. The AME of *Expansion2* is negative and significant for (5,7] year bonds for both of the high-*bank debt ratio* firms and the low-*bank debt ratio* firms. The AME for the maturity bin of (3,5] years is statistically significant only for the latter group, although the difference in this AME between the two sub-samples is not statistically significant. In summary, the result of Panel B of Table A2.5 is in line with that of Table 2.7 in that no clear cross-sectional differences in the effect of *Expansion2* are observed.

Table A2.5: Ratio of bank debt and the average marginal effects of Expansion2

Panel A: Summary statistics of the ratio of bank debt to total debt

	Mean	S.d.	Min.	Median	Max.	N
Bank debt ratio (%)	63.946	21.464	2.405	68.783	100.000	482

Panel B: Comparing the average marginal effects of Expansion2

	(1) *Bank debt ratio* $\geq$ median	(2) *Bank debt ratio* $<$ median	Diff.: (1) - (2)
[1,3]	0.114**	0.160***	-0.046
	(0.048)	(0.054)	(0.072)
(3,5] years	0.029	0.131*	-0.102
	(0.063)	(0.067)	(0.092)
(5,7] years	-0.093**	-0.135***	0.042
	(0.042)	(0.041)	(0.059)
(7,10] years	0.058	-0.022	0.080
	(0.064)	(0.065)	(0.091)
>10 years	-0.109**	-0.134***	0.026
	(0.042)	(0.035)	(0.054)

Panel A provides a descriptive statistic of *bank debt ratio*, which is the fraction of bank debt over total debt (expressed in percentage terms) at the beginning of the sample period. The data source is the Capital IQ Capital Structure Summary database (accessed via WRDS). In Panel B, the average marginal effects (AMEs) of *Expansion2* are reported together with standard errors in parentheses. The estimation method is the same as that of Table 2.7. The AMEs are calculated based on the MNLM where the dependent variable is the categorical variable for the maturity bins and the independent variables are *Expansion1*, *Expansion2*, an indicator variable that takes a value of one if the firm's *bank debt ratio* is greater than or equal to the sample median, and the interaction between *Expansion2* and the indicator variable. The AMEs are calculated separately based on the indicator variable values, and the differences in AMEs are also tested. ***, **, and * indicate that the difference is significant at 1%, 5%, and 10% levels, respectively.

Table A2.6: Changes in the determinants of offering spreads

	Pre-expansion period	Expansion2 period	Expansion2 vs. Pre-expansion	
			Diff. in coeff.	χ^2 stat. [p-value]
Maturity: (3,5] Y	8.14***	9.40***	1.26	0.14 [0.71]
	(2.12)	(2.93)		
Maturity: (5,7] Y	19.08***	20.07***	0.99	0.05 [0.82]
	(2.33)	(4.01)		
Maturity: (7,10] Y	14.15***	17.02***	2.87	0.75 [0.39]
	(2.13)	(2.78)		
Maturity: >10 Y	18.53***	13.43***	-5.10	1.72 [0.19]
	(2.60)	(3.16)		
Credit rating: A	8.94***	8.36***	-0.58	0.04 [0.84]
	(1.33)	(2.84)		
Credit rating: BBB	23.66***	20.67**	-2.99	0.10 [0.75]
	(2.70)	(9.63)		
Ln(proceeds)	2.89***	-0.30	-3.19	2.17 [0.14]
	(1.08)	(2.03)		
Ln(total assets)	-1.27*	-0.73	0.54	0.09 [0.76]
	(0.76)	(1.73)		
Net book leverage	10.96***	19.29**	8.32	0.99 [0.32]
	(3.90)	(7.99)		
Profitability	-30.07*	-50.34	-20.28	0.17 [0.68]
	(17.78)	(49.36)		
Asset tangibility	-13.34***	-14.72***	-1.38	0.05 [0.82]
	(3.73)	(5.29)		
Industry FE	✓	✓		
N	311	156		
Adjusted R^2	0.469	0.416		
Mean of dep. var.	44.55	34.76		
S.D. of dep. var.	13.41	14.01		

The dependent variable is the offering spread. Variable definitions are provided in Table 2.3. All the continuous variables are winsorized at the 2.5% and 97.5% levels. In the first two columns, heteroskedasticity-robust standard errors are reported in parentheses. The differences in coefficients are tested using Wald tests obtained by the "stacking" method (Weesie et al., 2000). ***, **, and * indicate that the difference is significant at 1%, 5%, and 10% levels, respectively.

Chapter 3

Coarse Pricing and Competition in QE Auctions

3.1 Introduction

In March 2009, the Fed launched a large-scale purchase program of long-term Treasury securities, now marked as the beginning of its quantitative easing (QE) purchases of government bonds. After more than a decade of the Fed's QE, there is currently abundant and growing literature on its macroeconomic effects. Yet, we still know little about the market where these massive purchases take place. This is rather puzzling given the Fed's enormous purchase sizes and the role of the U.S. Treasury market as the "single most important financial market in the world" (Group of Thirty, 2021). To fill this gap, this paper closely examines offer-level data from the Fed's reverse auctions (hereafter called QE auctions) for Treasury security purchases,[1] and document that primary dealers (PDs) of the Federal Reserve Bank of New York (FRBNY)—the only direct participants in QE auctions—exhibit coarse pricing behavior.

Specifically, in QE auctions the Fed sets the uniform tick size of 1/256th (i.e., 0.390625 cents)

[1]As discussed in Section 3.2, some of the reverse auctions in my sample period are better framed as traditional market operations rather than QE. However, for convenience, I use the term QE auctions to refer to all the Fed's reverse auctions of Treasury coupons.

per \$100 par value for all target Treasury notes and bonds, although quotes of those securities are seldom posted at this level of fineness in secondary trading venues. Figure 3.1 demonstrates the existence of coarse pricing among (successful) offers in QE auctions, and perhaps more surprisingly, a striking difference between two PDs: Morgan Stanley (with the second largest market share in my sample of QE auctions) and Credit Suisse (with the seventh largest market share). Whereas Morgan Stanley's price endings are almost uniformly distributed, Credit Suisse's distribution displays a strong clustering of price endings on 0, 4/256, 8/256,. . . , 252/256 (i.e., 1/64ths).[2] Coarse pricing has been documented in various financial (and non-financial) markets,[3] but it is particularly surprising in this setting. First, Treasury securities are one of the world's most liquid and heavily researched asset classes. Second, PDs (and institutions indirectly participating in QE auctions) are highly sophisticated investors with expertise in fixed-income valuation. Third, in QE auctions there is no incentive to sacrifice pricing precision for execution priority; offers are treated equally as long as they are submitted by the QE auction closure.

This paper studies the causes and consequences of this previously undocumented practice in QE auctions. Market microstructure literature suggests two possible reasons. First, dealers' coarse pricing might work as a mechanism to coordinate among themselves, thereby extracting higher profits. A well-discussed example is NASDAQ market makers' avoidance of odd-eighth quotes, first documented by Christie and Schultz (1994). This practice was controversial especially because virtually no market makers deviated from the practice for *certain* stocks (Christie and Schultz, 1994). While disagreement exists about what caused this apparently anti-competitive practice,[4] it is a fact that this practice mechanically limited the minimum possible bid-ask spread to two eighths for the affected stocks. Nevertheless, this collusion view is inconsistent with the coarse pricing patterns observed in QE auctions. As suggested in Figure 3.1, top dealers, who by definition

[2]Note that Credit Suisse's smaller sample size can lead to greater variability in realized proportions of price endings, but not clustering on a specific subset of price endings.

[3]They include stock exchanges (Harris, 1991; Ikenberry and Weston, 2008; Bhattacharya et al., 2012), gold markets (Ball et al., 1985), future markets (Ap Gwilym et al., 1998b,a; Kuo et al., 2015), and particularly relevant to this research are dealer markets such as NASDAQ (Christie and Schultz, 1994; Christie et al., 1994; Christie and Schultz, 1999) and municipal bond markets (Li, 2007; Griffin et al., 2023). Nikiforov and Pilotte (2017a,b, 2019) also document price-end clustering of Treasury coupon securities in the secondary market.

[4]Christie and Schultz (1994) interpret it as indicating implicit collusion, yet others point out institutional features limiting dealer competition (Demsetz, 1997; Kandel and Marx, 1997).

account for the bulk of transactions, engage in coarse pricing much less frequently, regardless of the Treasury security type.

An alternative view for dealers' coarse pricing is the "competitive theory of clustering" of Grossman et al. (1997), which was also proposed to explain the NASDAQ market maker behavior. According to this view, dealers, even when competing in the financial market, engage in coarse pricing due to information processing costs associated with increased pricing precision. Grossman et al. (1997, p. 25) state, "Finer units of trade allow for more accurate pricing. But this is a mixed blessing. It takes time and effort to obtain more precise valuations of assets," and as a result, "[t]he precise degree of coarseness chosen will depend on the balance between the benefits and costs of a finer grid." In the QE auction context, pricing precision's information costs can lead PDs to price on grids coarser than the Fed's 1/256ths grid, especially because secondary-market transactions are predominantly based on coarser grids (mainly 1/64ths or 1/128ths).

My results are consistent with the theory of Grossman et al. (1997). First, there exists a positive and strong association between the pricing fineness of (accepted) offers and the PD's market share in QE auctions. This result is expected if it is not costless for dealer banks to develop and employ sophisticated pricing technology tailored for this special QE market. (Note that the costs include not only capital investments but also, and perhaps more importantly, the attraction, retention, and deployment of human capital; highly-skilled traders are a critical asset in high finance). Because the return from pricing technology increases with the investment size (Arrow, 1987; Peress, 2004), the cost-benefit tradeoff indicates a positive association between the market share and sophisticated pricing.

The estimated between-PD difference is economically sizable. This paper proposes a method to quantify PD-level pricing fineness in a manner that is directly interpretable and comparable. It infers the proportion at which a particular PD used each of the following four possible pricing grids—1/32nds, 1/64ths, 1/128ths, and 1/256ths. Because offered Treasury security types vary by PD and time, this method controls for this heterogeneity, estimating the proportions in case the PD prices a 'typical' security in QE auctions: off-the-run Treasury note maturing in 5–10 years. Take

Figure 3.1 again. The estimated proportion of using the finest 1/256ths pricing grid is 94.9% for Morgan Stanley and 40.5% for Credit Suisse. More generally, according to my baseline specification, a one percentage point increase in QE auction market share translates into a 2.4 percentage point increase in the use of the finest 1/256ths grid.

Second, if the cost of increasing pricing precision drives dealers' coarse pricing, we should observe greater coarse pricing when it is more difficult—and therefore more costly—to precisely price the security (Ball et al., 1985; Grossman et al., 1997). I find a number of results in line with this prediction. First, coarse pricing is more pronounced for Treasury securities whose valuation is supposedly more difficult: those with longer remaining maturities (i.e., greater interest rate risk) and with high volatility. The second and perhaps more direct evidence is that offers for securities that are less finely priced in the secondary market are also priced less finely in QE auctions.[5] When a specific Treasury security's secondary-market prices are fine, it should be easier for PDs to price the security precisely in QE auctions—either because the fine secondary-market prices directly help them confidently price the security on a less coarse grid, or because fine prices in the secondary market indicate the intrinsic ease of precisely pricing the security. Moreover, I document that this effect is particularly large for non-top dealers.

Third, consistent with the view that coarse pricing per se is not an optimal auction strategy, I uncover one channel through which coarse pricing leads to a competitive disadvantage: getting undercut. It has been documented that retail traders tend to disproportionately submit prices ending with 0 or 5, and sophisticated traders strategically place orders one-tick below those clustered prices (Bhattacharya et al., 2012). I document similar patterns for price endings surrounding the coarsest 1/32nds pricing grid. Prices that are one-tick below this grid (i.e., 7/256, 15/256,..., 255/256) are observed more often than prices one-tick above it (i.e., 1/256, 9/256,..., 249/256) among successful offers. I further show that even after controlling for offer characteristics, *some* PDs (most notably Barclays) disproportionately submit undercutting offers. Overall, these results underscore the nature and competitive implications of coarse pricing in QE auctions.

Over time, PDs price more finely in QE auctions. This competition between PDs in upgrading

[5] Specifically, I collect daily closing ask prices of Treasury securities on the trading day preceding each QE auction.

pricing precision can be seen as direct micro-evidence of investor evolution in the adaptive markets view of Lo (2019). Notably, the Treasury security trading landscape has undergone remarkable changes during the analysis period. Joint Staff Report (2015) and Brainard (2018) indicate that in the 2010s dealers increasingly adopted fintech, such as systems for automated Treasury security trading, yet at different speeds.

The COVID-19 crisis period is, however, an exception to the trend toward greater pricing fineness. In massive Treasury purchase auctions of March 2020, pricing fineness temporarily plunged. This is a period in which the Treasury security market exhibited rare malfunctioning; in contrast to typical crisis episodes, the prices of long-term Treasury securities—arguably the world's safe haven—decreased in the wake of the COVID-19 crisis.[6] Yet, multivariate analysis indicates that the massive purchase sizes themselves, and the resulting low offer-to-cover ratios, contributed to the rise of coarse prices during this period. The inverse relationship between the offer-to-cover ratio and the prevalence of coarse pricing indicates the role of competition in curbing this phenomenon.

Two data limitations should be acknowledged. The first is that the offer-level data disclosed in accordance with the Dodd-Frank Act includes all *accepted* offers but not losing offers. This data curtailment precludes recovering some key structural parameters, such as dealers' marginal valuations (Boneva et al., 2020). In the context of this paper, this limitation's most direct consequence is that I can observe the pricing fineness of only winning offers, although the main interest lies in understanding how PDs price in general. This data curtailment can introduce a bias, because price-end fineness itself might affect the auction-winning probability. Indeed, the undercutting result does suggest that coarsely priced offers have (moderately) lower QE auction winning probabilities. In this sense, coarse pricing observed among *winning* offers should be regarded as a *lower* bound for the true extent of coarse pricing, unconditional on the QE auction outcome.

Second, although QE auctions permit PDs to submit not only their own offers but also their clients' offers, the data do not flag client offers, and therefore those offers cannot be discerned.[7] Consequently, the between-dealer variation might reflect different proportions of client offers. For

[6]The driver was investors dashing for cash, selling even safe Treasury securities to dealers at a massive scale (Duffie, 2020; Schrimpf et al., 2020; He et al., 2022; Vissing-Jorgensen, 2021).

[7]This is also the case of the proprietary U.K. QE auction data analyzed by Boneva et al. (2020).

example, the positive association between PD market share and pricing fineness can arise if client offers constitute a larger proportion for larger-market-share PDs *and* if client offers tend to be more finely priced. It appears highly unlikely that different client proportions solely explain the huge between-PD gap in pricing fineness.[8] Nevertheless, this issue should be kept in mind when interpreting the results of this paper.

The last part of this paper investigates coarse pricing from the viewpoint of the Fed's QE operation costs. Admittedly, the lack of losing offer data limits the thoroughness of this analysis. My empirical approach is to compare the levels of offer prices among accepted offers for the same Treasury security in the same QE auction. The fixed-effects regressions show that *conditional on winning*, coarsely priced offers are more highly priced. Note that this does not necessarily mean that coarsely priced offers are generally more highly priced. To win a QE auction an offer needs to have a sufficiently *low* price. Consequently, higher prices, conditional on winning, can be observed as a result of a reduced probability of winning the auction. Indeed, the undercutting result does strongly imply that coarsely priced offers have lower winning probabilities. In any case, this exercise indicates that coarse pricing has cost implications for the Fed.

This paper is policy-relevant. First of all, this paper sheds new light on the competition and microstructure of this important yet understudied market. In implementing QE-driven purchases, the Fed monitors "the performance of operations" and "the extent and concentration of dealer participation in operations" (Potter, 2013, p. 4). More broadly, there is an ongoing discussion about the optimal size—and "diversity"—of the Fed counterparties.[9] Based on the current primary dealer system, the Fed's market operations are solely intermediated by large Treasury security dealers designated as primary dealers. On one hand, Potter (2015) says that the system requires "established, regulated market participants" and "must be of an appropriate size to provide adequate execution capacity and competitive pricing." On the other hand, he also notes that "[s]taff time and

[8] At the same time, I do not exclude the possibility that PDs' precise pricing capabilities are influenced by client offers. For example, finely priced offers from clients, or just receiving more client offers, might be helpful for the PD to increase the precision of its own offers.

[9] The "diversity" debate has emerged due to the increasing presence of non-dealers (most notably the so-called "principal trading firms") in the secondary market.

resources are required to monitor and manage" relationships with counterparties.[10] This paper's results highlight the special role of topmost dealers in facilitating and enhancing price competition and market efficiency.

Second, the Fed's massive purchase scale means that coarse pricing can lead to substantial cost implications. In my sample period, the total amount of QE purchases of Treasury coupons is $3.92 trillion (with the average per auction being $4.12 billion). Then, one tick size change multiplied by the total purchase size amounts to $15.3 billion. The tick size is also economically significant in comparison to the Fed's transaction costs in QE auctions. Song and Zhu (2018) estimate that the weighted average cost (relative to the best ask price at auction closing) is merely 0.71 cents per $100 par value. This is roughly equivalent to two minimum ticks (0.78 cents per $100 par value).

This paper contributes to three strands of literature: First, this paper contributes to the literature on dealers' coarse pricing. In addition to the NASDAQ market maker behavior mentioned above, Li (2007) and Griffin et al. (2023) document that municipal bond dealers often price on a coarse grid (in eighths). Griffin et al. (2023) further show that dealers enjoy higher markups from such transactions. Note that the suspected mechanism here is the conflict of interest between well-informed sellers (dealers) and less-informed buyers (including retail customers) under decentralized bilateral transactions.[11] These market characteristics are quite different from QE auctions, where many PDs compete to sell more transparent Treasury securities. I show that in QE auctions dealers' coarse pricing is consistent with the theory of Grossman et al. (1997), in which they face a tradeoff between the pricing precision's cost and competitive advantage.

Second, this paper contributes to the emerging literature on the implementation mechanisms of QE. Song and Zhu (2018) study the Fed's preference in QE auctions and PDs' bidding behavior using the same data as this paper.[12] They first confirm that as the Fed's public disclosures suggest, the Fed prefers "cheap" Treasury securities, i.e., those whose secondary market prices are low

[10] The Fed's primary dealer pilot program (2013–2014), which temporarily allowed four relatively small dealers to participate in the Fed's operations, was one of the attempts to search for a better balance.

[11] The municipal bond market is unique among over-the-counter markets in that the participation of retail investors is high (see Bessembinder et al. (2020)).

[12] PDs' bidding behavior in Treasury issuance auctions is studied by Hortaçsu and Kastl (2012), Hortaçsu et al. (2018), and Allen et al. (2020).

relative to those implied by the yield-curve model. Song and Zhu (2018) then show that PDs extract higher profits when offering such securities. Boneva et al. (2020) analyze U.K. QE auction data containing both winning and losing offers. They structurally estimate U.K. primary dealers' marginal valuations of offered gilts, and show that these valuations are related to the interest rate risk and the regulatory capital requirements the dealers face.[13] This paper complements these papers in understanding price competition in QE auctions, taking cues from price-end patterns.

Third, this paper adds to the literature on the limited sophistication of sophisticated investors. Previous papers study institutional investors' herding (Wermers, 1999; Nofsinger and Sias, 1999; Griffin et al., 2003) and heuristic decisions (Akepanidtaworn et al., 2021; Wang, 2020) in stock trading.[14] This paper documents a manifestation of information processing constraints in the context in which it is arguably least expected—pricing of Treasury securities, the world's most heavily researched asset class, by the New York Fed's primary dealers, one of the most sophisticated and influential investor groups in the global economy (He et al., 2017; Goldberg, 2020).[15] Notably, pricing fineness of the market's most sophisticated investors is particularly interesting because it can even speak to the limit of market efficiency (Grossman and Stiglitz, 1980; Mondria et al., 2022).

3.2 Institutional background

3.2.1 Primary dealers

PDs are major dealers in U.S. Treasury securities designated by the FRBYNY as (sole) counterparties of the Fed's market operations. Most notably, PDs are allowed to submit competitive bids

[13] More broadly, the implementation cost of QE is also studied by D'Amico and King (2013) in the U.S., Breedon (2018) in the U.K., and Schlepper et al. (2020) in the Eurosystem. In addition, An and Song (forthcoming) study the Fed's purchase prices in its agency MBS operations.

[14] Wang (2020) shows that institutional investors, like retail investors but to a lesser extent, exhibit the "round-number bias"—i.e., the tendency to submit orders with rounded price endings such as ".0"—in stock trading, and that this bias is lower for larger institutions. Note that his research question is fundamentally different from mine. Price-end clustering measured this way likely reflects traders' cognitive reference points (also see Bhattacharya et al. (2012), whose method Wang (2020) follows).

[15] Fleming et al. (2005) and Goldreich (2015) also study bounded rationality of PDs. They document evidence of sub-optimal bids in Treasury bill issuance auctions (until an auction rule change in 2004). One notable difference between their datasets and mine is that theirs are aggregated auction-level data, precluding PD-level analysis.

in Treasury issuance auctions (and profit from selling the purchased securities in the secondary market). A primary dealer status can also help the dealer attract large customers such as foreign central banks (Rennison and McLannahan, 2016). Their main requirements are as follows. First, PDs are expected to bid a certain amount in every Treasury issuance auction. Second, for other Fed operations (including QE operations), each PD is expected to participate "at levels commensurate with its size and presence in the market."[16] The last requirement is assisting the Fed to formulate monetary policy. During my sample period, the number of PDs changed from 18 in 2010 to 24 in 2020.[17] Table 3.1 lists PDs in my sample period.

3.2.2 Phases of the Fed's QE

Table 3.2 lists major events among the Fed's QE purchases of Treasury securities. As shown in the last column, this paper divides the sample period into five sub-periods: *QE2* (August 17, 2010–September 20, 2020), *MEP* (September 21, 2011–December 11, 2012), *QE3* (December 12, 2012–October 29, 2014), *QE pause* (October 30, 2014–March 11, 2020), and *QE4* (March 12–June 29, 2020). Appendix 3.A provides a summary of the history of the Fed's QE. Figure 3.2 shows the time series of the Fed's purchases of Treasury coupon securities.

The following clarifications will prove useful. First, the first sub-period (*QE2*) centers around, but is not limited to, the so-called "QE2" round, which was implemented from November 3, 2010 through June 22, 2011. Before and after the round, the Fed reinvested the proceeds of maturing securities into Treasury securities to maintain its balance sheet size. Those purchases are included in the *QE2* sub-period. Second, the *MEP* and *QE3* sub-periods exactly match the periods of the Maturity Extension Program (MEP) and the so-called "QE3," respectively. Third, the MEP was different from the other QE rounds in that the Fed financed the purchases of long-term Treasury securities by selling short-term securities (thereby not changing the aggregate bank reserves). Fourth, although the *QE pause* sub-period starts in October 2014, most of the purchases took place in August 2019 or thereafter. During this period, the Fed resumed reinvesting in Treasury securities

[16] The Fed's "Administration of Relationships with Primary Dealers" statement published on March 24, 2016 (`https://www.newyorkfed.org/markets/pridealers_policies.html`)

[17] The current and historical lists of PDs are available at `https://www.newyorkfed.org/markets/primarydealers`.

to maintain a sufficient size of bank reserves. Lastly, the *QE4* sub-period covers QE purchases in response to the COVID-19 crisis (up to the sample conclusion of June 2020). Note that the primary purpose of the post-pandemic purchases was to tame the disruptions in the market (see Appendix 3.A for more discussions).

3.2.3 Structure of QE auctions

The Desk of the FRBNY administers QE auctions, as in the case of other Fed market operations. Auction-level summary statistics are provided in Table 3.3. Details of the QE auction protocol vary from phase to phase, yet a typical timeline is as follows. First, the QE auction schedule for the next monthly cycle is announced around the previous monthly cycle's end. The schedule lists the dates of auctions, together with their target maturity ranges (e.g., Treasury coupons with 4.5 to 7 years remaining to maturity) and expected purchase sizes.[18] Securities excluded from the auction, such as those scarce in the secondary market, are announced at the auction start. PDs are therefore well informed of which securities are included in the auction.

Second, on the auction date, PDs can submit up to nine offers per security. Each offer consists of price and quantity. Unlike Treasury issuance auctions, QE auctions are price-discriminatory; the offer price is the price at which the PD sells the security to the Fed if the offer is accepted. While PDs are the only direct participants, they can also place their clients' offers on their behalf. Both the minimum offer size and the price increment are set at $1 million. The tick size is 1/256th per $100 par value for all Treasury notes and bonds.

The Fed's QE auctions are multi-object auctions, as each auction accepts multiple Treasury securities within the target maturity range.[19] Therefore, the FRBNY adopts the following approach to rank offers for different securities. It first calculates the benchmark price for each auction-target security by applying its proprietary yield-curve model to secondary market prices (Sack, 2011).

[18] In early QE phases, the FRBNY disclosed the expected purchase *ranges*, while only the *maximum* purchase amounts are announced in later phases.

[19] In the average QE auction, 30.7 unique Treasury securities meet the basic eligibility criteria of QE auctions, with 3.8 of them being explicitly excluded for other reasons. Among the remaining 26.9 CUSIPs, 17.0 are purchased (Panel B of Table 3.3).

Then, offer prices normalized by their benchmark prices are directly comparable regardless of the differences in the offered securities. The FRBNY purchases from the offers with the lowest prices relative to the model-implied price until its desired total purchase amount is reached. Note that this protocol means that the Fed prefers securities whose market prices are below the prices implied by the Fed's model, i.e., "cheap" securities (Song and Zhu, 2018).

In my sample period, the median auction time is 45 minutes, and the auction close time ranges from 9:50 a.m. to 3:05 p.m., with 90.9% of them ending in the morning. It is exceedingly rare for an auction to end after 2 p.m. (1.6%).[20] The settlement is typically the next day, meaning that PDs can cover their short positions anytime in the afternoon of the auction date.

Lastly, the FRBNY releases the auction result in three steps. First, immediately after each auction, the FRBNY releases quantity-related information, such as the total offered and accepted amounts and the purchased amount per security. Participating PDs are notified of their auction outcomes at this time. Second, the price-related information is disclosed at the end of each monthly auction cycle. The information is aggregated to the (purchased) security level. The disclosed items are the total offered and purchased amounts, the weighted average and highest accepted prices, and the proportion of accepted offers among the highest price offers. Lastly, about two years after the auction, the FRBNY releases disaggregated data on accepted offers, which is the primary dataset used in this study.

3.3 Data

The FRBNY publicly discloses all *accepted* offers of QE auctions after the enactment of the Dodd-Frank Wall Street Reform and Consumer Protection Act on July 21, 2010.[21] The offer-level data includes the Treasury CUSIP, offer price, offered amount, and the identity of the PD who submitted the offer. I retrieved all accepted offers for nominal Treasury notes and bonds.[22] My sample period

[20] More specifically, until 2019, the vast majority (95.9%) of QE auctions ended at 11:00 a.m. In the post-COVID-19 period, the Fed tended to hold multiple QE auctions (for different target maturities) at different times on the same day, resulting in more diverse auction close time.

[21] `https://www.newyorkfed.org/markets/omo_transaction_data.html`

[22] Therefore, my sample does not include Treasury bills or TIPS.

spans from August 17, 2010 (the date of the first QE auction post-Dodd-Frank Act) to June 29, 2020 (the last auction date in 2020Q2). Treasury security information was obtained from TreasuryDirect.[23]

Table 3.4 reports the descriptive statistics of my sample offers. Security types dramatically vary over time. The vast majority of the purchased securities are off the run, with no on-the-run securities being purchased after the QE3 period. Also, the composition of remaining maturities varies strikingly over time. While 20–30 years account for 50.8% in the MEP period, 0–5 years is the majority (53.0%) in the QE4 period. These target security composition changes highlight the importance of controlling for security types in the empirical analysis. The bottom rows of Table 3.4 concern the market shares of PDs who submitted offers. The market share is based on trade amounts and calculated for each sub-period. Following Song and Zhu (2018), I group PDs into top five PDs and non-top five PDs.[24]

3.4 Result

3.4.1 Setup

The market convention for quoting Treasury coupon securities is per \$100 par value, with the decimal part being a multiple of 1/32 or a multiple of a fraction of 1/32 (that is, 1/64, 1/128, or 1/256).[25] In QE auctions, the tick size is set to one eight of 1/32, i.e., 1/256. Throughout the paper, a Treasury security prices' decimal part is denoted as $d/256$. That is, $d \in D = \{0, 1, 2, \ldots, 255\}$ and $d = 1$ means the price ending is 1/256.

To analyze price-end clustering, which indicates dealers' use of coarser pricing grids, I divide D

[23] https://www.treasurydirect.gov/instit/annceresult/annceresult_query.htm

[24] The remaining results are robust to the use of alternative cut-off points, such as top four or top seven. Section 3.4.7 conducts detailed PD-level analysis.

[25] For example, given a \$1,000 face value, a price quote of \$101.09375 (= \$101 + 3/32) indicates that the bond's price is \$1,010.9375.

into four mutually exclusive subsets:

$$D = \{0, 1, 2, \ldots, 255\} = X_{32} \cup X_{64} \cup X_{128} \cup X_{256}, \text{where} \begin{cases} X_{32} = \{0, 8, 16, \ldots, 248\} \\ X_{64} = \{4, 12, 20, \ldots, 252\} \\ X_{128} = \{2, 6, 10, \ldots, 254\} \\ X_{256} = \{1, 3, 5, \ldots, 255\} \end{cases} \tag{3.1}$$

In words, each subset represents the possible *coarsest* pricing grid that the PD could have used to arrive at the particular d. For example, if $d \in X_{64}$, the possible coarsest pricing grid used by the dealer is 1/64ths. This is because the dealer can arrive at $d \in X_{64}$ based on the pricing grid of 1/64ths, 1/128ths, or 1/256ths, but not on the 1/32nds grid. Note that only in the case of $d \in X_{256}$, the pricing grid used by the PD can be identified with certainty (to be 1/256ths).

3.4.2 Testing price-end clustering on coarser grids in QE auctions

I follow Kuo et al. (2015) and Bhattacharya et al. (2018) to statistically test the existence of price-end clustering on coarser grids. Specifically, for each price ending $d \in D = \{0, 1, 2, \ldots, 255\}$, I calculate the percentage of offers with this price ending among all offers. To facilitate interpretation, this percentage is then subtracted by the expected percentage under the uniform priced-end distribution, which is $1/256 \times 100 = 0.390625\%$. This outcome variable is regressed on the following variables:

$$Percent_d - 0.390625 = \alpha + \beta_1 D_{32} + \beta_2 D_{64} + \beta_3 D_{128} + \epsilon_{i,j}, \tag{3.2}$$

where D_{32}, D_{64}, and D_{128} take the value of one if the price ending d belongs to X_{32}, X_{64}, and X_{128}, respectively, and zero otherwise.

The regression results are reported in Panel A of Table 3.5. Column 1, which pools all sample (winning) offers, shows strong price-end clustering on coarser grids. According to the constant, the percentage of a d in $X_{256} = \{1, 3, 5, \ldots, 255\}$ being selected is 0.177% (as it is 0.213 percentage points less than what is expected under the uniform distribution). This is due to clustering on grids coarser than 1/256ths; the three price-end type dummy variables are positive and statistically significant, meaning that price endings in D_{32}, D_{64}, and D_{128} are more likely to be selected. Based on Column 1, the probability of d in X_{32} being chosen is 0.177% + 0.839% = 1.016%. Likewise, the probabilities of a price ending in X_{64} and X_{128} are 0.719% and 0.340%. Importantly, the differences between D_{32} and D_{64} and between D_{64} and D_{128} are statistically significant. This result implies that PDs tend to use *all* of the coarse pricing grids, namely 1/32nds, 1/64ths, and 1/128ths.[26]

3.4.3 Estimating the proportions in which different pricing grids are used

My fundamental interest is not the observed price-end clustering itself but inferring the pricing grids that dealers used from it. I thus infer the likelihood that PDs employed each of the four possible pricing grids from the coefficients of Specification 3.2. For now, assume that the price-end type itself does not affect the auction-winning probability. If PDs always use the finest 1/256ths pricing grid (*grid-256ths*), then $\Pr[d \in X_{32}] = P[d \in X_{64}] = 12.5\%$, $\Pr[d \in X_{128}] = 25\%$, and $\Pr[d \in X_{256}] = 50\%$. I denote these conditional probabilities as $\phi_{32}^{grid\text{-}256ths}$, $\phi_{64}^{grid\text{-}256ths}$, $\phi_{128}^{grid\text{-}256ths}$, and $\phi_{256}^{grid\text{-}256ths}$, respectively. Table 3.6 summarizes these conditional probabilities for all the pricing grids.

Now, suppose that dealers use the 1/32nds pricing grid (*grid-32nds*) with probability λ_{32}, and the pricing grids of 1/64ths, 1/128ths, and 1/256ths with probabilities λ_{64}, λ_{128}, and λ_{256}, respectively. Then, the unconditional probabilities of price-end types are:

$$\begin{aligned}
\phi_{32} &= \lambda_{32} \times \phi_{32}^{grid\text{-}32nds} + \lambda_{64} \times \phi_{32}^{grid\text{-}64ths} + \lambda_{128} \times \phi_{32}^{grid\text{-}128ths} + \lambda_{256} \times \phi_{32}^{grid\text{-}256ths}, \\
\phi_{64} &= \lambda_{32} \times \phi_{64}^{grid\text{-}32nds} + \lambda_{64} \times \phi_{64}^{grid\text{-}64ths} + \lambda_{128} \times \phi_{64}^{grid\text{-}128ths} + \lambda_{256} \times \phi_{64}^{grid\text{-}256ths}, \\
\phi_{128} &= \lambda_{32} \times \phi_{128}^{grid\text{-}32nds} + \lambda_{64} \times \phi_{128}^{grid\text{-}64ths} + \lambda_{128} \times \phi_{128}^{grid\text{-}128ths} + \lambda_{256} \times \phi_{128}^{grid\text{-}256ths}, \\
\phi_{256} &= \lambda_{32} \times \phi_{256}^{grid\text{-}32nds} + \lambda_{64} \times \phi_{256}^{grid\text{-}64ths} + \lambda_{128} \times \phi_{256}^{grid\text{-}128ths} + \lambda_{256} \times \phi_{256}^{grid\text{-}256ths}.
\end{aligned}$$

[26] For example, if PDs used only the pricing grids of 1/128ths and 1/256ths, the coefficients of D_{32} and D_{64} should not be greater than that of D_{128}.

The log-likelihood function can be defined as follows:

$$\ln \mathrm{L}(\lambda) = Fraction_{32} \times \ln\, \phi_{32} + Fraction_{64} \times \ln\, \phi_{64} + Fraction_{128} \times \ln\, \phi_{128} + Fraction_{256} \times \ln\, \phi_{256}, \tag{3.3}$$

where $Fraction_{32}$ is the observed fraction of $d \in X_{32}$, and $Fraction_{64}$, $Fraction_{128}$, and $Fraction_{256}$ are similarly defined.[27] I maximize this function to obtain the estimates of $\lambda = \{\lambda_{32}, \lambda_{64}, \lambda_{128}, \lambda_{256}\}$, with a constraint of each element of λ being between 0 and 1.

Panel B of Table 3.5 reports the predicted proportions in which the four pricing grids are used. The finest 1/256ths grid is predicted to be employed 45.4% of the time, and the coarsest 1/32nds grid 9.5% of the time. The rest is accounted for by the two pricing grids between these two extremes. I note, however, one caveat regarding this prediction. As discussed later, coarsely priced offers likely have a lower chance of winning a QE auction. This means that these proportions predicted from *accepted* offers are biased, in the direction of upward (downward) biasing the proportion of fine (coarse) grid use. As such, the predicted proportion of 9.5% for the coarsest grid use should be regarded as the lower bound if we consider dealers' grid use *un*conditional on the QE auction outcome.

3.4.4 Price-end clustering by time and by PD

Table 3.5 also shows that price-end clustering in QE auctions greatly varies by time and the PD submitting the offer. In Columns 2–4, I repeat Specification 3.2 separately for each sub-period. There are two notable findings. First, PDs tend to price more finely over time. While the predicted proportion of the 1/256ths grid use is only 14.5% in the QE2 period, it accounts for more than 50% in the subsequent periods. Second, the predicted proportion of the coarsest 1/32nds grid use jumped to a much higher value than the earlier periods (21.2%) in the QE4 period. More specifically, Figure 3.3 illustrates that this surge in the coarsest pricing took place during the massive purchases of March 2020. The extent to which changing auction characteristics can explain these time-series patterns is examined in the next section.

[27] These ratios can be computed from the coefficients of Specification 3.2.

The remaining columns of Table 3.5 show a stark difference in pricing fineness between top and non-top dealers. All of the coefficients of D_{32}, D_{64}, and D_{128} are greater for non-top five dealers than top dealers, and the differences in all of the coefficients are statistically significant at the 1% level (unreported). Non-top dealers thus engage in more coarse pricing practices. Panel B further illustrates the economic significance of the disparity, with the predicted proportion of the 1/256ths grid use being 66.2% for top dealers and 26.3% for non-top dealers.

3.4.5 Determinants of coarse pricing in QE auctions

To analyze heterogeneity in coarse pricing, I regress the fineness of offer prices on security-, offer-, and auction-characteristics. The key idea is that the price-end type, X_{32}, X_{64}, X_{128}, or X_{256}, indicates the likelihood with which coarser pricing grids are used, with X_{32} being the highest and X_{256} the least. (To be precise, the probability of coarse pricing is zero if d is in X_{256}.) I therefore perform the ordered logit analysis with the dependent variable being *Price-end fineness*, which takes the value of one if the offer's price ending is in X_{32}, two if it is in X_{64}, three if it is in X_{128}, and four if it is in X_{256}. This analysis aims to uncover which factors increase (or decrease) the likelihood of precise pricing. Table 3.7 lists the definitions of variables used in this offer-level analysis. The descriptive statistics are provided in Table 3.8.

One key explanatory variable is *Cheapness*, which is included to analyze the effect of the Fed's algorithm in ranking offers. As explained in Section 3.2.3, the Fed's protocol prefers securities that are deemed undervalued based on the Fed's yield curve model. Although the model is not publicly disclosed, it is a standard cubic spline model (Sack, 2011), and Song and Zhu (2018) find that their yield curve estimation result—and the ensuing cheapness measure—can explain the Fed's purchase behavior. I thus estimate a cubic spline model following Song and Zhu (2018), and define *Cheapness* as the percentage difference between the yield-curve-implied price and the actual secondary-market price. The details of the yield curve estimation are summarized in Appendix 3.B.

Table 3.9 shows the estimation results. Panel A reports the coefficients and the marginal effects on the probability of *Price-end fineness* = 4, that is, $D_{256} = 1$.[28] The coefficient of *Top five* is positive and significant in all models. According to Model 1, the probability of the price ending being the finest type (i.e., $D_{256} = 1$) is 14.9 percentage points higher for top dealers' offers. This is sizable given that only 22.7% of sample offers have this finest price-end type (Table 3.8). Models 2 and 3 show that the marginal effect remains stable even when additional control variables are added.

In Model 1, *Cheapness* is negative and significant, meaning that prices are coarser when the Treasury security is deemed undervalued and therefore preferred by the Fed. Because strategic dealers can extract higher profits when delivering cheaper securities in QE auctions (Song and Zhu, 2018), those dealers might be using coarse pricing as a device to set a higher price for cheaper securities. However, the evidence is not consistent with this view.

First, Panel B further investigates this association by splitting the sample by *Top five* and reporting the marginal effect of *Cheapness* on all possible values of *Price-end fineness*. It shows that the marginal effects of *Cheapness* are generally more pronounced for non-top dealers. Strikingly, the marginal effect on the probability of *Price-end fineness* = 1 is 0.357 for those dealers, meaning that a one standard deviation change of *Cheapness* translates into an 8.50 percentage point increase in this probability. The stronger association for *non*-top dealers does not align well with the narrative that the effect of *Cheapness* is a result of dealers' deliberate coarse pricing for cheap securities. Instead, a more plausible explanation is that when the security is deemed undervalued by the Fed, it is more likely that even coarsely priced—or less sophisticatedly priced—offers can win the QE auction. (In Section 3.4.9, I show that coarsely priced offers are more highly priced among accepted offers.) This is because offers for cheap securities have an intrinsically higher chance of winning a QE auction based on the Fed's algorithm to rank offers.

Second, the effect of *Cheapness* is sensitive to the inclusion of controls, with the coefficient even changing its sign. In particular, *Cheapness* is no longer negatively associated with *Price-end fineness* after controlling for three basic security types—on-the-run status, maturity, and Treasury bond dummy (Model 2 in Panel A of Table 3.9). On one hand, this result relates to the fact that

[28]The ordered logit model can derive marginal effects for each possible value of the outcome variable.

Cheapness is highly positively correlated with remaining maturity.[29] On the other hand, this result is inconsistent with the view that dealers use coarse pricing deliberately and specifically for cheap securities. Rather, this result indicates that certain security types (most notably maturity) are associated with *both* pricing fineness and *Cheapness*.

Model 3 provides three additional insights. First, *Voaltility*, the standard deviation of the offered security in the five-day period leading up to the auction date, is negatively associated with pricing fineness. This finding is consistent with the hypothesis of Ball et al. (1985) and Grossman et al. (1997) that valuation uncertainty, and the resulting greater cost in precise pricing, drives coarse pricing.

Second, the result of *Offer-to-cover* suggests the role of competition in restricting coarse pricing. According to Model 3, one standard deviation increase in *Offer-to-cover* translates into a 1.99 percentage point increase in the probability of $D_{256} = 1$. Note that this is the effect after controlling for period dummies, and moreover, this result is not driven by massive post-COVID-19 purchases; the coefficient remains significant at the 1% level even if the QE4 period is dropped.

Third, pricing fineness is greater in later phases than in the first QE2 phase (the model's baseline period), and this is more evident after controlling for security- and auction-characteristics. In Model 1, the coefficient of *QE4* is not significant with a low marginal effect on the probability of $D_{256} = 1$ (a 3.8 percentage point increase). Yet, the variable is significant at the 1% level in Model 3, which controls for security- and auction-type variables. Most notably, *Volatility* is high and *Offer-to-cover* is low in the QE4 period, and these factors contributed to deteriorated pricing fineness in the post-pandemic QE operations.

3.4.6 Pricing fineness in the secondary market for Treasury securities

I now turn my attention to the secondary market for Treasury securities with the purpose of better understanding the QE auction market's coarse pricing.

[29] The correlation between *Cheapness* and remaining maturity in years is 0.754.

Trading venues and tick sizes

Since the launches of eSpeed and BrokerTec in 1999 (the latter of which started service in 2000), trades of on-the-run Treasuries have almost entirely moved to fully electronic systems (Mizrach and Neely, 2006). These markets now accommodate non-dealers, with the most significant non-dealer participants being the so-called principal trading firms (PDFs) (Harkrader and Puglia, 2020). However, once a security goes off-the-run, the trading volume plunges, and most of the trading migrates to more traditional voice-assisted systems (Barclay et al., 2006). Off-the-run trading is also characterized by high market fragmentation, as opposed to the on-the-run trading arena, where two platforms, BrokerTec and Dealerweb, dominate others (McPartland, 2018).

Tick sizes vary by transaction venue and maturity. The main electronic venues of on-the-run securities, such as BrokerTec and eSpeed (then acquired by Dealerweb), set the tick sizes of 1/128th for 2-, 3-, and 5-year notes and 1/64th for 7- and 10-year notes and 30-year bonds in the period investigated in this section.[30] Conversely, in the case of off-the-run trading on voice-assisted markets, it is customary, albeit not a rule, that brokers display tick sizes coarser than 1/256th (such as 1/64th and 1/128th), except for securities nearing maturity.

Pricing fineness in the secondary market for Treasury securities

To understand pricing fineness in the secondary market, I repeat the price-end clustering regression with secondary-market price data. For each Treasury security purchased in a QE auction, I obtained the security's daily closing ask price on the preceding trading day from Bloomberg.[31] The result of repeating Specification 3.2 with the secondary-market price data is reported in Table 3.10. According to Column 1 of Panel B, the predicted proportion of the 1/256ths grid use is 13.4%. The sub-period analysis reveals that this value remained virtually zero for the first three sub-periods (*QE2*, *MEP*, and *QE3*). The 1/64ths and 1/128ths grids account for a large fraction, and, there is

[30] These venues lowered the tick size of 2-year notes to 1/256th on November 19, 2018 (Fleming et al., forthcoming). Note, however, that the Fed purchased on-the-run securities only until August 2014.

[31] I do not use the CRSP Treasury data for studying price endings due to its apparent anomaly. For unknown reasons, the proportion of price endings in D_{128} doubled on March 7, 2012; while the average proportion was 23% for the preceding ten trading days, it jumped to 48% in the next ten trading days, and this was not a temporary phenomenon but a permanent shift. Bloomberg data does not exhibit such dramatic and puzzling price-end patterns.

little evidence of the coarsest 1/32nds grid use.[32]

Table 3.11 then estimates the ordered logit model of the determinants of pricing fineness of the secondary-market price data. Importantly, maturity exhibits a strong nagative association with pricing fineness, as was the case for QE auction offers. (*On-the-run* and *Bond* have the same signs, but they are not statistically significant in this secondary-market data analysis.) For example, relative to the baseline maturity of five to ten years, maturity of less than five years is associated with a 9.7 (18.6) percentage point higher probability of $D_{256}^{Secondary} = 1$ ($D_{128}^{Secondary} = 1$). To sum up, Table 3.11 suggests a link in pricing fineness between the secondary market and QE auctions.

Relationship in pricing fineness between the secondary market and QE auctions

To directly test the association, Table 3.12 regresses QE auction offers' *Price-end fineness* on the price-end type dummies of the security's secondary market price; $D_{256}^{Secondary}$ is a dummy variable taking the value of one if the price ending of the Treasury security's closing ask on the trading day preceding the QE auction is in X_{256}, and I similarly define $D_{32}^{Secondary}$, $D_{64}^{Secondary}$, and $D_{128}^{Secondary}$. Specifically, I estimate the ordered logit model with the baseline category being $D_{32}^{Secondary} = 1$. The coefficients show whether finer prices in the secondary market indicate finer offer prices in QE auctions (i.e., a higher *Price-end fineness*).

The positive coefficients of $D_{128}^{Secondary}$ and $D_{256}^{Secondary}$ in Column 1 confirm the association. They remain statistically significant even when security-type controls (*On-the-run*, $Maturity_{\text{0-5}}$, $Maturity_{\text{10-20}}$, $Maturity_{\text{20-30}}$, and *Bond*) are added (Column 4). To demonstrate the economic significance, Figure 3.4 plots the predicted probabilities according to the models of Columns 2 and 3 of Table 3.12. In particular, the predicted probabilities for the two extreme price-end types, X_{256} and X_{32}, greatly vary by the secondary-market price-end type for non-top dealers (Panel B of Figure 3.4). When $D_{32}^{Secondary} = 1$, the predicted probability of X_{256} is merely 11.5%. However, it jumps to 29.2% (i.e., a 153.9% increase) when $D_{256}^{Secondary} = 1$. On the other hand, the predicted probability of X_{32} drops from 42.8% to 19.1%.

[32] Nikiforov and Pilotte (2017a) document a similar price-end distribution using tick-level data on Treasury notes.

These results strongly substantiate the costly nature of pricing precision as a driver of coarse pricing in QE auctions. There are two possible channels through which finer prices in the secondary market are associated with a lower cost of precise pricing for PDs. First, having fine secondary-market prices as inputs to the model makes it easier for dealers to confidently price securities on a fine grid. Second, securities with fine prices in the secondary market (e.g., short-maturity securities) are expected to be intrinsically easier to precisely price. That non-top dealers' pricing fineness is highly sensitive to the fineness of secondary market prices suggests that pricing easiness relaxes the limit of precise pricing that lagging PDs face.

3.4.7 Dealer-level analysis

PD-level pricing fineness

This section looks more closely at PD-level differences in coarse pricing. Pricing fineness varies by not only PD but also offered security type (Table 3.9). I thus quantify pricing fineness of each PD, after controlling for basic security characteristics. More specifically, I employ a variant of Specification 3.2, in which the right-hand side of the model includes dummy variables for three basic security characteristics—on-the-run status, remaining maturity, and Treasury note vs. bond—and PD $\times$ sub-period fixed effects. Appendix 3.C explains more details about this method.[33] This approach allows for predicting the probabilities of price endings in X_{32}, X_{64}, X_{128}, and X_{256} for each PD in each sub-period in the case of offering 'typical' security in QE auctions, namely, off-the-run Treasury note maturing in 5–10 years. These predicted probabilities are then fed into the method of Section 3.4.3 to estimate the proportions in which different price grids were used. The summary statistics of the estimated PD-level proportions can be found in Table A3.2.

[33]To include those dummy variables and fixed effects, the outcome variable, $Percent - 0.390625$, is calculated for each security type s, PD j, and sub-period t. Because this means that the ratios are calculated from different numbers of offers, the regression weights observations by the number of underlying offers.

Pricing fineness and market share

PD-level analysis also points to a strong association between pricing fineness and market share. Figure 3.5 plots dealer market share and the predicted proportion of using the finest 1/256ths grid. Panel A of Table 3.13 shows that market share is significantly and positively associated with the use of the finest pricing grid for all the sub-periods other than *QE2*. The association becomes somewhat weaker but still remains, even if I use the predicted proportion of using the 1/128ths or 1/256ths grid (Columns 5–8).

In Column 1 of Panel B, which pools the four sub-periods, *Market share* is statistically and economically significant. The coefficient indicates that a one percentage point increase in *Market share* translates into a 2.4 percentage point increase in the finest 1/256ths grid use. Notably, this result is not driven by a few of the largest PDs, as it is robust to the exclusions of Goldman Sachs (Column 2) and Morgan Stanley (Column 3). In Column 4, I also remove small PDs (with less than 2% market share in the period). The result remains highly similar. Panel C shows that the coefficients of *Market share* remain similar even if the regressions are performed on the first differences.

Having established a strong positive association between pricing fineness and QE auction market share, I also ask whether pricing fineness is related to more primitive dealer characteristics such as balance sheet size, location, and experience. Appendix 3.D summarizes the data collection and the analysis result. Panel A of Table A3.6 shows that the use of the finest 1/256ths grid is positively associated with the balance sheet size and its use is lower for foreign PDs (i.e., PDs that are a subsidiary of a foreign-based financial group). However, these results are statistically significant only at the 10% level and become insignificant if the dependent variable is replaced with the probability of using the 128ths or 256ths grid (Panel B). Therefore, these variables' explanatory power is weak at best. In contrast, the market share remains highly significant even if PD-level controls are added. One interpretation is that the key driver of the between-PD variation in the tendency of precise pricing in QE auctions is their presence in this specific market, rather than their overall size or experience. At the same time, I note substantial measurement issues of my dealer-level variables; dealer size is measured infrequently and noisily,[34] and experience is only roughly measured.

3.4.8 Undercutting

One mechanism through which coarse pricing can affect the auction outcome is a lower acceptance chance, being undercut by finely-priced offers. Figure 3.6 visually indicates this phenomenon. The number of price endings that are one-tick below X_{32} is greater than that of price endings that are one-tick above X_{32}, especially in 2012, 2013, and 2020.

Let $X_{UNDER32}$ refer to a set of d's that are one-tick below X_{32}, i.e., $X_{UNDER32} = \{7, 15, 23, \ldots, 255\}$. Similarly, define $X_{OVER32} = \{1, 9, 17, \ldots, 249\}$, that is, X_{OVER32} refers to price endings that are one-tick below X_{32}. The null hypothesis that the number of d's in $X_{UNDER32}$ is the same as the number of d's in X_{OVER32} is tested in Table 3.14. It is rejected at the 10% level in the QE2 period and at the 1% level in the subsequent MEP and QE3 periods. Interestingly, while undercutting-type price endings are no more prevalent in the QE-pause period, the proportion surpassed 50% once again in the QE4 period, in which price-end clustering (and in particular that on X_{32}) resurged.

Note that this imbalance among *winning* offers can occur even if *no* dealers are *strategically* submitting offers to undercut coarsely-priced offers. This is because different price endings can have different auction-winning probabilities. To gauge this effect, Appendix 3.E performs a simulation exercise under a setting akin to the average QE auction in my sample. It predicts that the proportion of d's in $X_{UNDER32}$ over d's in $X_{UNDER32} \cup X_{OVER32}$ based on *winning* offers would be 0.5201 or 0.5241, depending on a simulation assumption, even when *no* PDs are strategically submitting offers with price endings in $X_{UNDER32}$. Therefore, it is possible that $X_{UNDER32}$ is observed more frequently than X_{OVER32} even in the absence of strategic undercutting.

Yet, it is possible to indirectly test the existence of strategic undercutting by looking at between-PD variation. The idea is that, in the absence of a strategic undercutting motive, the proportion of d's in $X_{UNDER32}$ over d's in $X_{UNDER32} \cup X_{OVER32}$ is expected not to vary significantly by

[34] For most PDs, a separate balance sheet item for Treasury securities does not exist, and I simply analyze the balance sheet size.

PD, at least after controlling for offer characteristics. In contrast, if *some* PDs engage in strategic undercutting, PD fixed effects should have jointly significant explanatory power.

To operationalize the idea, the following regression is estimated on the sample of offers whose price endings are in $X_{UNDER32}$ or X_{OVER32}:

$$\begin{aligned} Undercut_{32} = \alpha + \beta_1 \textit{On-the-run} + \beta_2 Maturity_{\text{0-5}} + \beta_3 Maturity_{\text{10-20}} + \beta_4 Maturity_{\text{20-30}} \\ + \beta_5 Bond + \boldsymbol{\gamma} \cdot \text{Period dummies} + \boldsymbol{\delta} \cdot \text{PD dummies} + \epsilon_i, \end{aligned} \tag{3.4}$$

where the dependent variable, $Undercut_{32}$, takes the value of one if the price ending is in $X_{UNDER32}$, and zero otherwise (i.e., in X_{OVER32}), and *On-the-run*, $Maturity_{\text{0-5}}$, $Maturity_{\text{10-20}}$, $Maturity_{\text{20-30}}$, and *Bond* are defined in the same way as before. The null hypothesis of interest is that the elements of the vector $\boldsymbol{\delta}$ (the coefficients of PD dummies) are jointly zero. The model is estimated with the OLS with the standard errors triple clustered by CUSIP, QE auction date, and PD.

Table 3.15 demonstrates strong explanatory power of PD fixed effects. Because the realized proportion of d's in $X_{UNDER32}$ might easily take extreme values for PDs with limited sample sizes, I repeat the analysis using only PDs with at least 200 sample offers (i.e., offers with d's in $X_{UNDER32} \cup X_{OVER32}$). Models 2 and 3 show the results of running Specification 3.4. Regardless of the sample restriction, the null that each of the PD dummy coefficients equals zero is rejected at the 1% level.

To more strictly control for offered security heterogeneity, Models 2 and 4 include CUSIP $\times$ QE auction fixed effects, which subsume all the explanatory variables of Specification 3.4, except for PD dummies. Here, the identification of PD dummies comes from instances in which more than one PD submits offers with price endings in $X_{UNDER32}$ or $X_{UNDER32}$ for the same security on the same auction date. The joint tests are significant at the 1% and 5% levels in Models 2 and 4, respectively. Notably, the most striking result comes from Barclays (see Models 3–4). It indeed has 731 sample offers, among which 493 (67.4%) have price endings in $X_{UNDER32}$. This analysis thus points to deliberate undercutting, uncovering one channel through which dealers using coarse pricing get outmaneuvered.

3.4.9 Coarse pricing and the level of offer prices

Does coarse pricing relate to the *level* of offer prices, which ultimately determines the Fed's QE purchase costs? I examine this question by looking at differences in the level of prices *within* accepted offers for a given Treasury security in the QE auction. Specifically, I define *Price diff* as the percentage difference (in basis points) between the offer price and the lowest accepted offer price for the same security in the same QE auction, and run the following regression:

$$Price\ diff = \beta_1 D_{32} + \beta_2 D_{64} + \beta_3 D_{128} + \gamma + \epsilon_{i,j}, \tag{3.5}$$

where D_{32}, D_{64}, and D_{128} take the value of one if the (accepted) offer's price ending d belongs to X_{32}, X_{64}, and X_{128}, respectively, and zero otherwise, and γ is CUSIP $\times$ QE-auction fixed effects. I measure *Price diff* only when multiple accepted offers exist for the security in the QE auction.[35] Absorbing between-security variation, γ ensures that D_{32}, D_{64}, and D_{128} compare *Price diff*—the offer price normalized by the lowest accepted offer for the security—for accepted offers with different price-end types *within* accepted offers for the same security in the same QE auction.

Panel A of Table 3.16 shows that the mean *Price diff*, which is winsorized at the 2.5% and 97.5%, is 3.6 basis points. This is fairly small, being less than the average bid-ask spread for my sample QE auction securities, 4.3 basis points (Table 3.8). This result is consistent with the claim of Song and Zhu (2018) that the QE auctions' operation costs are fairly moderate.

The regression results are documented in Panel B of Table 3.16. Column 1 is the result of estimating Specification 3.5, and it shows that (accepted) offers with coarse price endings have significantly higher prices compared to those with price endings in X_{256}. The differences are also economically significant. The coefficient of D_{32} (0.282) indicates that, relative to the mean value (3.61), (accepted) offers with the coarsest price endings have, on average, a 7.81% higher *Price diff*, which is the percentage difference between the offer's price and the minimum accepted offer price for the security in the QE auction.

[35] This restriction only slightly decreases the number of sample (accepted) offers: from 90,022 to 86,611. Conditional on having multiple accepted offers, the average (median) number of accepted offers for a Treasury security in a QE auction is 12.4 (10).

Therefore, coarsely priced offers on average have higher prices, *conditional on winning*. The lack of losing offer data precludes exactly pinpointing the origin. However, note that this result itself does not necessarily mean that coarsely priced offers, unconditional on the QE auction outcome, are on average more highly priced. For an offer to win a QE auction, its price needs to be sufficiently *low* relative to competing offers. Therefore, coarsely priced offers can have a higher average price conditional on winning simply because they have a lower chance of winning a QE auction in the neighborhood of the lowest acceptable price. Indeed, this is what the undercutting analysis (Section 3.4.8) and the simulation exercise (Appendix 3.E) suggest.

The remaining columns in Panel B of Table 3.16 show that *who* submits offers also matters. Note that in the presence of CUSIP $\times$ QE-auction fixed effects, I can add only variables that vary at the *offer* level to the right-hand side of Specification 3.5. While the offer amount (*Ln(offer amount)*) has only an insignificant effect, *Top five*, the dummy for offers submitted by a top five dealer, is negative and significant at the 5% level in Columns 3 and 4. Importantly, the inclusion of *Top five* materially lowers the effects of coarse pricing; the coefficients of D_{32} and D_{64} in Column 3 are slightly less than half of those in Column 1.

To summarize, Table 3.16 demonstrates that coarse pricing does indeed matter from the viewpoint of the Fed's QE operation costs. At the same time, prices are lower for accepted offers submitted by top dealers, and the association between pricing fineness and the offer price level is nearly halved when the top five dealer dummy is added to the right-hand side of the model. Therefore, the association partially reflects the fact that coarsely priced offers are more likely to come from non-top five dealers, whose accepted prices tend to be more highly priced.

3.5 Discussion of client offers

The positive association between market share and pricing fineness is consistent with the notion that larger-market-share PDs are more willing to seek precise pricing; the information processing costs of pricing precision can lead to increasing returns from pricing technology (Arrow, 1987; Peress, 2004). This section discusses how the existence of client offers can affect this interpretation.

In QE auctions non-PDs (such as hedge funds and money managers) are allowed to submit offers indirectly through PDs (Sack, 2011). Unfortunately, these client offers cannot be discerned in my data. Consequently, one concern is that the positive association between pricing fineness and market share might reflect different proportions of client offers *and* different extent of coarse pricing between PDs and clients. Consider a (seemingly more plausible) situation in which top PDs route more client offers due to their more extensive customer network. In such a case, in order for client offers alone to produce the positive association between PD market share and pricing fineness, clients must have a tendency to price more finely than PDs. There is one institutional reason to suspect this. In these primary dealer-intermediated markets, clients are aware that PDs can revise their own offers after observing client offers, possibly deliberately undercutting them (Hortaçsu and Kastl, 2012). As such, coarse pricing can be particularly costly for clients.

The observed between-PD variation, however, still rejects client offers as a sole explanation. Consider the MEP period as an example (Panel B of Figure 3.5). In this period, the estimated proportion of using the 1/256ths pricing grid reached 100% for Morgan Stanley, the market share leader. Of course, this should be the case only if *both* the PD (Morgan Stanley) and its clients always used the finest pricing grid. This result, however, is inconsistent with the assumption that clients tend to price more finely than PDs. Therefore, there must be at least *some* between-PD differences in the tendency of coarse pricing.[36]

3.6 Conclusions

This paper sheds new light on price competition in a massive, important, yet still understudied market: QE reverse auctions of Treasury securities. To my knowledge, this paper is the first to document dealers' practice of submitting coarsely priced offers in this market. As such, it complements the work of Song and Zhu (2018), who conduct auction theory-based analysis of dealer behavior in the Fed's QE auctions. On one hand, coarse pricing is on a downward trend. On the other hand, coarse pricing surged during the Fed's massive pandemic-driven purchases in

[36]Notably, this logic does not exclude the possibility that PDs *learn* from client offers; observing a larger number of client offers might help top PDs to price more finely than others.

March 2020. I also show that *conditional on winning a QE auction*, coarsely priced offers are more highly priced. Therefore, this practice has relevance for policymakers designing and monitoring this market.

The cross-sectional analysis of pricing fineness reveals that it varies by the ease of precisely pricing the security. Offer prices are more finely priced when valuation uncertainty is limited (e.g., low volatility) and when the security is finely priced in the secondary market prior to the auction. This paper also documents strong association between PD market share and pricing fineness—offers submitted by larger-market-share PDs are more finely priced. The relationship is economically significant; my benchmark specification indicates a one percentage point increase in market share translates into a 2.39 percentage point increase in the estimated proportion of using the finest 1/256ths grid. Collectively, my results are consistent with Grossman et al.'s (1997) theory that information costs of increasing pricing precision lead to coarse pricing of dealers. In contrast, I do not find evidence that coarse pricing works as a coordinating mechanism for PDs to maintain high spreads. Yet, the results do imply that competition plays a role in constraining coarse pricing.

This paper thus illustrates the special importance of the topmost dealers in the Fed's counterparty framework—they can uniquely facilitate price competition and market efficiency in Fed operations. Also, from a theoretical standpoint, this paper presents rare micro-level evidence that information processing costs can lead to a trade-off in pricing precision even among highly sophisticated investors.

Figures

Figure 3.1: Price-end clustering of QE auction offers: Morgan Stanley vs. Credit Suisse

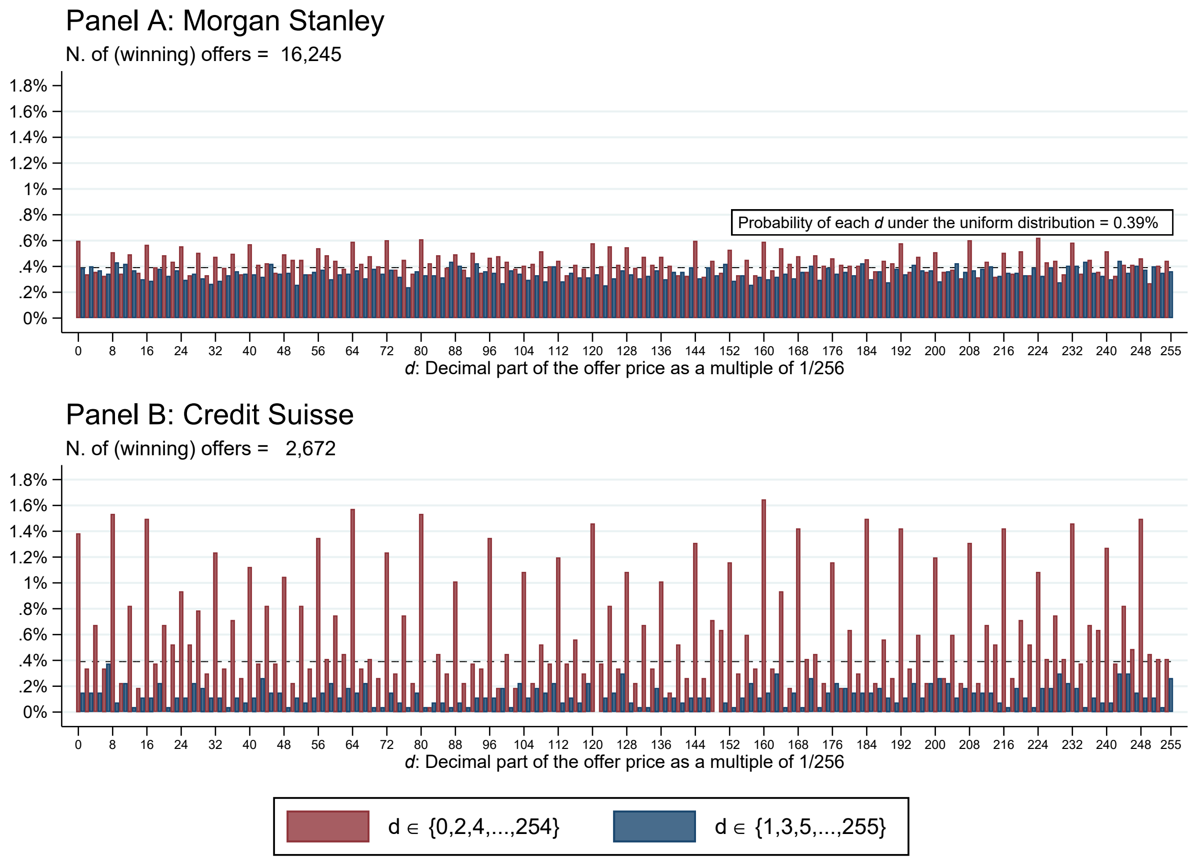

This figure shows the distributions of the decimal part of the offer price as a multiple s of 1/256 (d). The sample is (winning) offers in QE auctions from August 17, 2010 to June 29, 2020. The data is obtained from the FRBNY.

Figure 3.2: QE operations of Treasury notes and bonds

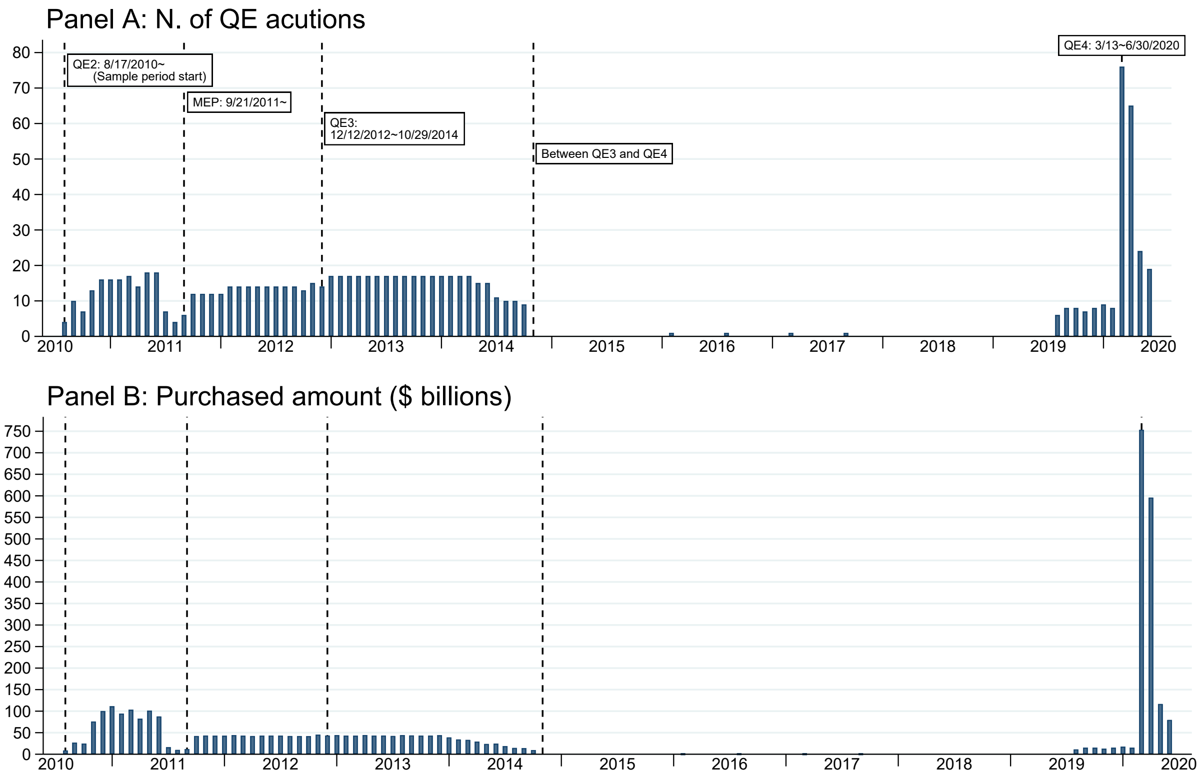

This figure shows the time series of the Fed's QE purchases of Treasury notes and bonds from August 17, 2010 to June 29, 2020. The data source is the Treasury securities operation results disclosed by the FRBNY at `https://www.newyorkfed.org/markets/desk-operations/treasury-securities`.

Figure 3.3: Price-end types of (successful) offers in QE auctions

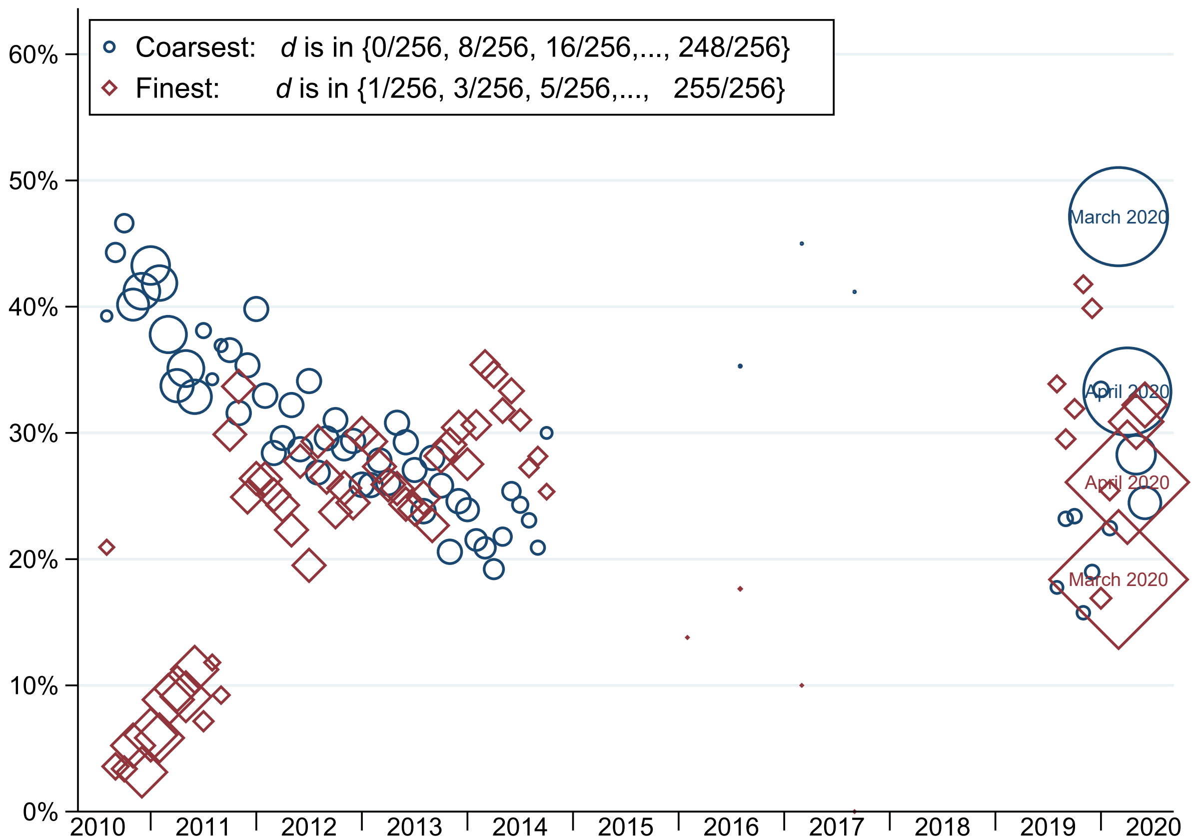

In this figure, successful offers in QE auctions are aggregated monthly, and each displayed symbol's size corresponds to the amount purchased in the month. For each month, the price-end types of (successful) offers are computed. The sample period is from August 17, 2010 to June 29, 2020.

Figure 3.4: Predicted probabilities of price-end types of QE auction offers based on Models 2 and 3 of Table 3.12

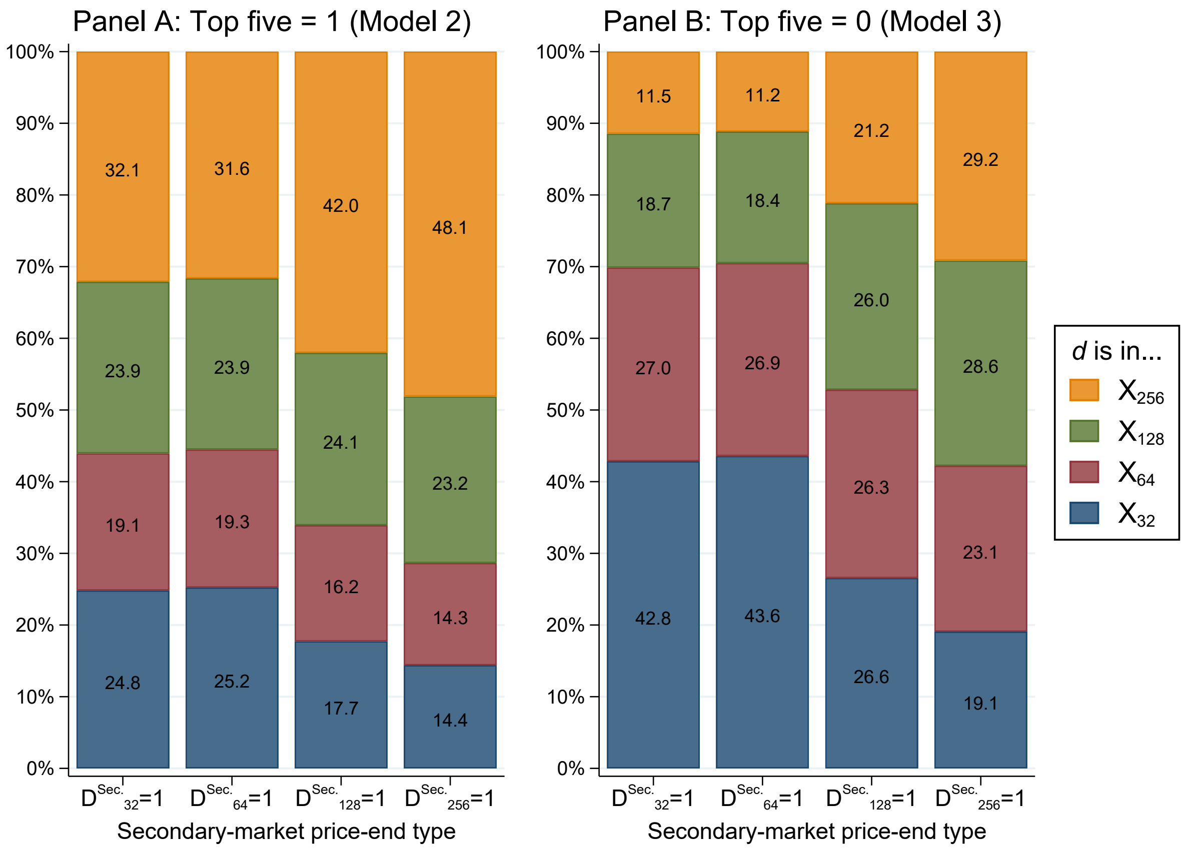

This figure shows the predicted probabilities of price-end types of QE auction offers based on Models 2 and 3 of Table 3.12. Model 2 uses the sample of $Top\ five = 1$, and Model 3 that of $Top\ five = 0$. The predicted probabilities are calculated for each of the four possible types of the secondary-market price ending: $D_{32}^{Secondary} = 1$, $D_{64}^{Secondary} = 1$, $D_{128}^{Secondary} = 1$, and $D_{256}^{Secondary} = 1$. $D_{32}^{Secondary}$ takes the value of one if the decimal part of the Treasury security's closing ask quote on the trading day preceding the QE auction is in X_{32}, and zero otherwise. $D_{64}^{Secondary} = 1$, $D_{128}^{Secondary} = 1$, and $D_{256}^{Secondary} = 1$ are similarly defined.

Figure 3.5: Market share and the predicted proportion of using the finest 1/256ths pricing grid

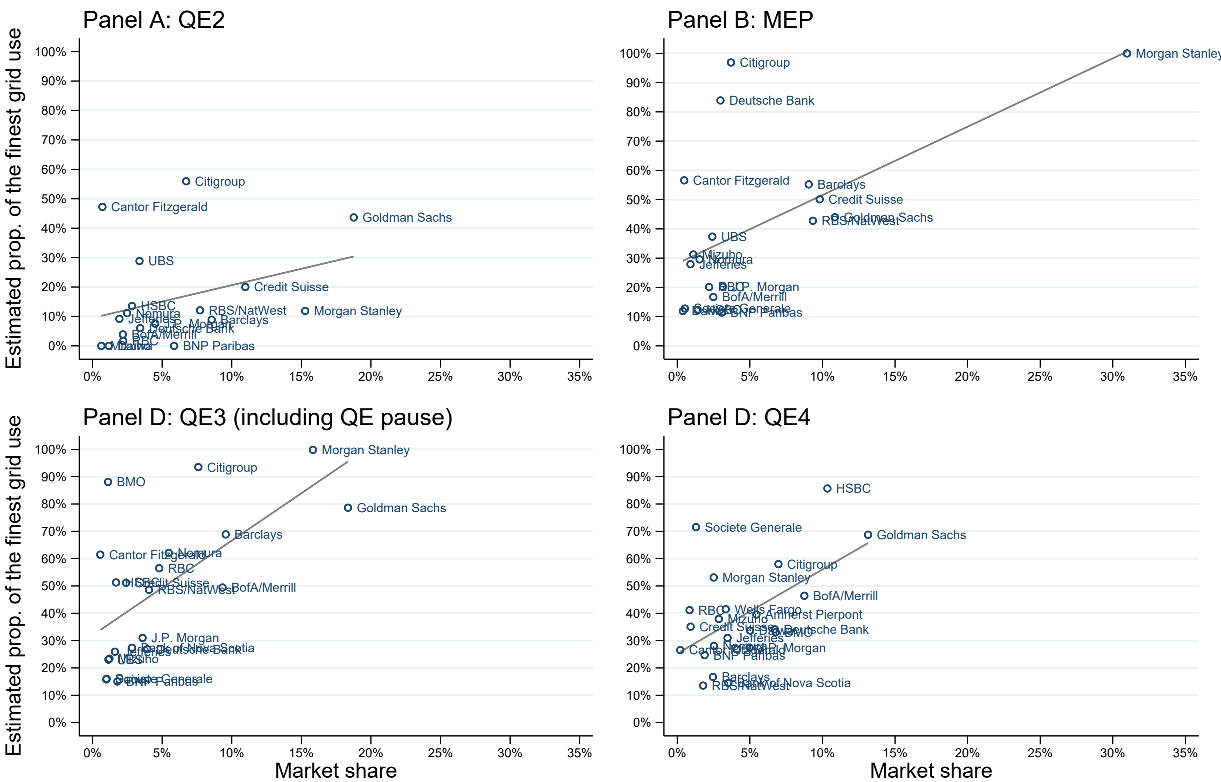

The x-axis measures the (traded amount-based) market share in QE auctions in the sub-period. The y-axis represents the estimated proportion of using the finest 1/256ths pricing grid. A PD is included in a panel only if it has at least 100 winning offers in the sub-period and if it was designated as a PD at the beginning of the sub-sample period. Each panel includes the fitted line.

Figure 3.6: Undercutting the most coarsely priced offers

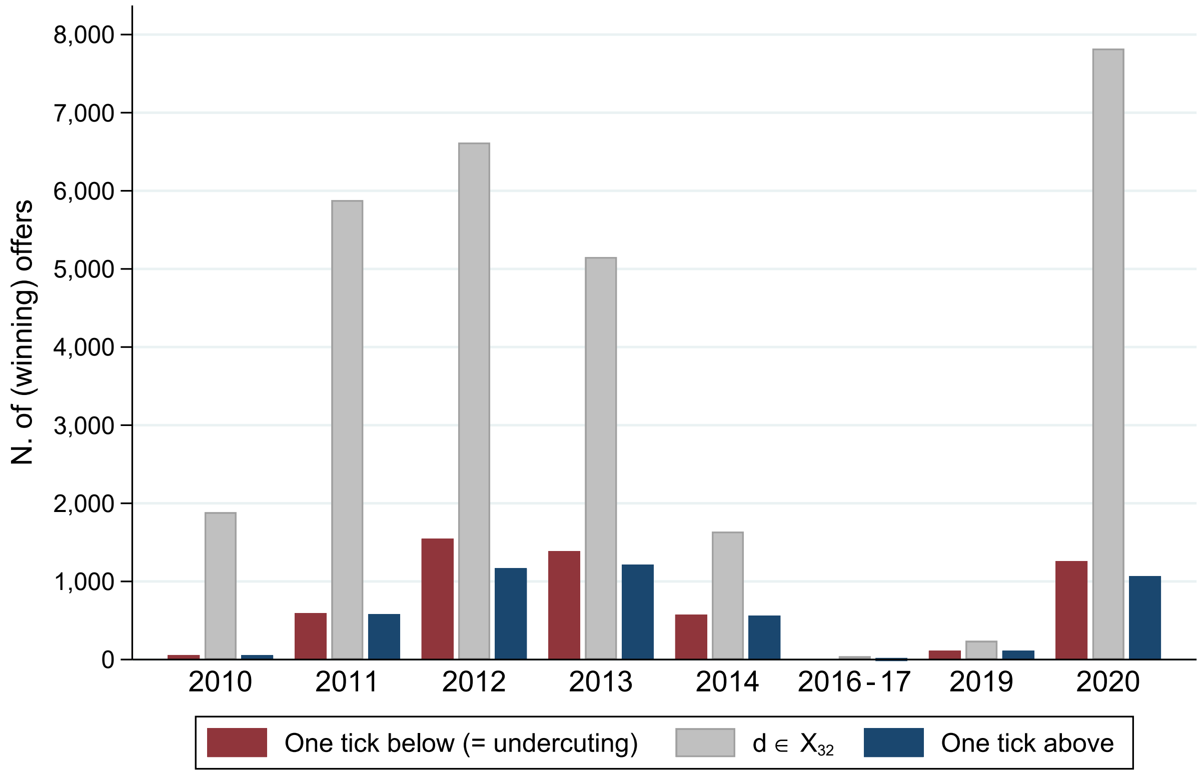

This figure examines whether offers whose price endings are on the coarsest 1/32nds grid (i.e., the most coarsely priced offers) are undercut by finely priced offers, by comparing the numbers of (winning) offers whose price endings are one tick below X_{32} (i.e., $\{7, 15, 23, \dots, 255\}$) or one tick above X_{32} (i.e., $\{1, 9, 17, \dots, 249\}$). For reference, the numbers of (winning) offers with price endings in X_{32} are also plotted. The sample period is from August 17, 2010 to June 29, 2020.

Tables

Table 3.1: Primary dealer list

Name	Parent location	PD designation date (post-8/17/2010)	Amount sold ($ billions)	N. of offers accepted
Goldman Sachs	Domestic		595.787	8,612
Morgan Stanley	Domestic		503.859	16,245
Citigroup	Domestic		256.365	4,103
Barclays	Foreign		253.379	6,915
BofA/Merrill	Domestic		253.098	5,725
HSBC	Foreign		207.759	2,930
Credit Suisse	Foreign		190.467	2,672
RBS/NatWest	Foreign		187.855	2,820
Deutsche Bank	Foreign		186.736	4,112
J.P. Morgan	Domestic		167.096	2,836
BMO	Foreign	10/4/2011–	126.949	2,896
Nomura	Foreign		119.014	2,782
BNP Paribas	Foreign		114.229	11,117
Daiwa	Foreign		98.555	1,577
Jefferies	Domestic		90.034	1,903
Bank of Nova Scotia	Foreign	10/4/2011–	89.314	1,711
RBC	Foreign		88.638	1,610
Amherst Pierpont	Domestic	5/6/2019–	85.405	1,572
Mizuho	Foreign		67.256	1,637
TD	Foreign	2/11/2014–	66.321	424
Wells Fargo	Domestic	4/18/2016–	56.648	883
UBS	Foreign		56.207	1,607
Societe Generale	Foreign	2/2/2011–	36.135	1,531
Cantor Fitzgerald	Domestic		17.174	1,338
MF Global	Domestic	2/2/2011–10/31/2011	2.847	110
Pilot program primary dealers (July 2013–July 2014)				
Cabrera			0.477	162
G.X. Clarke			0.315	154
Loop			0.044	23
Mischler			0.023	15

Primary dealers during my sample period (2010Q3–2020Q2) are listed. Note that four dealers participated in the primary dealer pilot program in July 2013–July 2014. See 3.6 for data sources of the parent location and period as a primary dealer. The amount sold and the number of winning offers are based on my sample QE auctions of Treasury notes and bonds.

Table 3.2: Timeline of major events in the Fed's QE purchases of Treasury securities

Date	**Event**	**Period**
March 18, 2009	Announcement of "QE1" of Treasury securities: The Fed would purchase up to \$300 billion of Treasury coupon securities in the next six months.	
March 25, 2009	Purchases of Treasury coupon securities began.	
October 29, 2009	QE1 ended.	
August 10, 2010	The FOMC announced the plan to reinvest proceeds of maturing Treasury securities, agency debt, and agency MBS in Treasury coupon securities.	
August 17, 2010	The Fed started disclosing detailed operation result data (roughly two years after the operation date).	*QE2*
November 3, 2010	Announcement of "QE2": The Fed would purchase \$600 billion of Treasury coupon securities by June 2011.	
June 22, 2011	The FOMC announced the end of QE2 and the plan to reinvest proceeds of maturing debt in Treasury coupon securities.	
September 21, 2011	Announcement of the Maturity Extension Program: The Fed would purchase \$400 billion of long-term Treasury securities (maturing in 6 to 30 years) based on proceeds from selling short-term ones (maturing in less than 6 years) by June 2012.	*MEP*
June 20, 2012	The termination date of the Maturity Extension Program was extended to December 2012. The Fed would continue the purchases of long-term Treasuries (and sales of short-term Treasuries) at the same pace.	
December 12, 2012	Announcement of "QE3" of Treasury securities: The Fed would continue the purchases of long-term Treasury securities, but unlike in the previous Maturity Extension Program, it would not match the purchase amounts with the proceeds from selling short-term securities.	*QE3*
October 29, 2014	QE3 ended. However, the Fed would continue reinvesting in Treasury coupon securities to maintain its balance sheet size at \$4.5 trillion.	*QE pause*
June 14, 2017	The FOMC announced the intention of initiating the balance sheet normalization program (reducing reinvestment in Treasury securities) this year, if the economic condition allows.	
September 20, 2017	The FOMC announced the start of the balance sheet normalization program in October 2017.	
July 31, 2019	The balance sheet normalization program was concluded. The Fed would reinvest up to \$20 billion per month in Treasury securities.	
March 12, 2020	Beginning of QE4" round: The Fed would purchase Treasury securities of various maturities to address highly unusual disruptions in Treasury financing markets associated with the coronavirus outbreak." The massive purchases started on March 13, 2020.	*QE4*
March 15, 2020	Announcement of QE4: The Fed would purchase at least \$500 billion of Treasury securities "over coming months."	

Sources: Announcements and events listed on the website of the Federal Reserve.

Table 3.3: Descriptive statistics of QE auctions

Panel A: N. of QE auctions per auction date

Period	Mean number of QE auctions	Distribution of the number of QE auctions						
		1	2	3	4	5	6	Total
QE2	1.019	159	3	0	0	0	0	162
MEP	1.049	176	9	0	0	0	0	185
QE3	1.012	343	4	0	0	0	0	347
QE pause	1.017	58	1	0	0	0	0	59
QE4	2.493	39	9	4	4	5	12	73
Total	1.153	775	26	4	4	5	12	826

Panel B: N. of eligible, included, and purchased securities per QE auction

Period	N	Eligible securitis			Included securities			Purchased securities		
		Mean	S.d.	Median	Mean	S.d.	Median	Mean	S.d.	Median
QE2	165	27.4	6.08	28.0	25.2	5.28	26.0	13.2	5.20	13.0
MEP	194	19.1	4.11	19.0	16.5	4.54	17.0	13.6	4.02	14.0
QE3	351	22.2	3.07	22.0	18.3	2.76	19.0	12.9	4.64	13.0
QE pause	60	38.0	7.84	40.0	33.8	8.56	36.5	11.4	6.41	10.0
QE4	182	60.0	30.41	52.0	53.9	29.12	49.0	33.8	19.68	29.0
Total	952	30.7	20.42	24.0	26.9	19.23	20.0	17.0	12.61	14.0

Panel C: Submitted and purchased amounts per QE auction

Period	N	Submitted amount ($ milllions)			Accepted amount ($ millions)			Offer-to-cover ratio		
		Mean	S.d.	Median	Mean	S.d.	Median	Mean	S.d.	Median
QE2	165	20,641	8,481	20,949	5,182	2,414	6,260	4.996	4.212	3.899
MEP	194	9,223	4,914	6,486	3,193	1,406	2,512	2.838	0.599	2.774
QE3	351	8,004	4,250	5,870	2,264	1,268	1,575	3.775	1.373	3.521
QE pause	60	9,549	4,172	9,530	1,768	520	1,801	5.549	2.279	5.189
QE4	182	18,797	12,725	14,270	8,477	5,967	6,000	2.344	0.853	2.104
Total	952	12,603	9,227	10,426	4,116	3,797	3,000	3.634	2.301	3.159

Panel D: Winning offers and dealers per QE auction

Period	N	N. of winning offers			N. of winning dealers			Mean N. of winning offers per winning dealer		
		Mean	S.d.	Median	Mean	S.d.	Median	Mean	S.d.	Median
QE2	165	94.2	55.3	89	15.2	3.33	16	5.96	3.12	5.61
MEP	194	128.4	58.5	117	16.2	3.64	17	7.71	2.75	7.38
QE3	351	80.1	42.9	72	15.6	3.88	16	5.03	2.30	4.64
QE pause	60	30.9	23.8	24	11.0	4.33	11	2.65	1.35	2.16
QE4	169	115.8	66.0	110	19.3	4.09	20	5.66	2.59	5.13
Total	939	95.9	58.3	86	16.0	4.26	17	5.71	2.86	5.18

This table presents desctiptive statistics of QE auctions of Treasury notes and bonds held from August 17, 2010 to June 29, 2020. The sample period is divided into five sub-periods based on QE phases: *QE2* (8/17/2010–9/19/2011), *MEP* (9/23/2011–12/10/2012), *QE3* (12/13/2012–10/27/2014), *QE pause* (2/23/2016–3/3/2020), and *QE4* (3/13/2020–6/29/2020). Panel A reports the number of separate QE auctions (based on the Fed's operation ID) that the Fed conducted per QE auction date. Panels B–D report QE auction-level characteristics. The sample size for Panel D is slightly smaller because there were 13 instances in which two separate QE auctions targeting the same set of Treasury securities were held on the same date. Panel D treats those pairs of QE auctions as a single observations, because the publicly disclosed offer-level data does not indicate which of the two auctions each offer belongs to in such cases.

Table 3.4: Descriptive statistics of (accepted) offers in QE auctions

	QE2	MEP	QE3	QE pause	QE4	Total
N	15,549	24,912	28,126	1,857	19,578	90,022
Offer size ($ millions)						
Mean	55	24.9	28.3	57.1	78.8	43.5
S.d.	121.1	78.2	65.5	129.8	128.1	99.7
Min	1	1	1	1	1	1
Median	25	5	10	25	49	15
Max	5,008	1,875	1,500	2,150	4,749	5,008
On-the-runs vs. off-the-runs (%)						
Off-the-run	89.0	96.0	95.8	100.0	100.0	95.7
On-the-run	11.0	4.0	4.2	0.0	0.0	4.3
Remaining maturities (years; %)						
0–5	34.4	0.0	10.7	46.3	53.0	21.8
5–10	50.8	37.3	37.9	30.7	22.7	36.5
10–20	6.6	7.0	4.7	5.4	2.1	5.1
20–30	8.2	55.8	46.7	17.5	22.2	36.6
Security types (%)						
2Y notes	1.0	0.0	0.0	7.5	8.6	2.2
3Y notes	10.2	0.0	0.0	8.5	11.8	4.5
5Y notes	18.3	0.0	6.3	18.1	18.8	9.6
7Y notes	22.6	15.3	15.7	22.5	17.5	17.3
10Y notes	23.8	18.3	21.6	19.3	18.0	20.3
30Y bonds	24.0	66.4	56.5	24.0	25.2	46.2
Market share of the primary dealer (%)						
Top five dealers	33.6	57.0	56.1	31.1	37.4	47.9
Non-top five dealers	66.4	43.0	43.9	68.9	62.6	52.1

This table reports the types of the sample offers for QE auctions of Treasury notes and bonds from August 17, 2010 to June 29, 2020. Primary dealer market shares are measured based on the trade amount in QE auctions in the sub-period.

Table 3.5: Price-end clustering on coarser grids

Panel A: Regression results								
	All	Sub-period					PD market share	
		QE2	MEP	QE3	QE pause	QE4	Top five	Non-top five
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
D_{32}	0.839***	1.155***	0.794***	0.584***	0.571***	1.037***	0.493***	1.157***
	(0.017)	(0.024)	(0.020)	(0.017)	(0.052)	(0.048)	(0.012)	(0.023)
D_{64}	0.542***	0.872***	0.536***	0.492***	0.378***	0.374***	0.332***	0.735***
	(0.008)	(0.020)	(0.016)	(0.014)	(0.042)	(0.013)	(0.007)	(0.011)
D_{128}	0.163***	0.322***	0.093***	0.156***	0.196***	0.134***	0.117***	0.206***
	(0.004)	(0.011)	(0.007)	(0.008)	(0.025)	(0.007)	(0.005)	(0.005)
Constant	-0.213***	-0.334***	-0.189***	-0.174***	-0.168***	-0.210***	-0.132***	-0.288***
	(0.002)	(0.002)	(0.003)	(0.003)	(0.012)	(0.004)	(0.003)	(0.002)
N	256	256	256	256	256	256	256	256
Adjusted R^2	0.980	0.967	0.956	0.937	0.542	0.910	0.954	0.981
$D_{32} - D_{64}$	0.297***	0.282***	0.259***	0.093***	0.194***	0.663***	0.161***	0.423***
	[265.535]	[79.976]	[107.412]	[18.534]	[9.090]	[181.815]	[154.367]	[264.990]
$D_{64} - D_{128}$	0.378***	0.550***	0.443***	0.335***	0.182***	0.240***	0.215***	0.528***
	[1977.001]	[578.012]	[684.574]	[467.250]	[15.948]	[296.327]	[866.277]	[1894.439]
Panel B: Estimated proportions in which different pricing grids are used (%)								
grid-32nds	9.52	9.04	8.28	2.96	6.19	21.21	5.15	13.53
grid-64ths	24.22	35.2	28.36	21.47	11.63	15.36	13.76	33.82
grid-128ths	20.89	41.23	11.84	19.99	25.09	17.15	14.92	26.37
grid-256ths	45.38	14.53	51.53	55.58	57.08	46.28	66.17	26.28

Panel A reports the results of estimating Specification 3.2. The dependent variable, $Percent_d - 0.390625$, is the percentage of offers with price endings being d among all offers, minus 0.390625. The right-hand side variables are D_{32}, D_{64}, and D_{128}, which take the value of one if the price ending d belongs to X_{32}, X_{64}, and X_{128}, respectively, and zero otherwise. Column 1 uses all sample (winning) offers. In Columns 2–6, the ratios are calculated separately for each sub-period: *QE2*, *MEP*, *QE3*, *QE pause*, and *QE4*. Columns 7 and 8 repeat the analysis for offers of top five PDs and non-top five PDs, respectively. PDs are ranked based on trade amounts in QE auctions in each sub-period.Heteroskedasticity-robust standard errors are reported in parentheses. At the bottom, differences between coefficients are tested and their F statistics are reported in bracket parentheses. ***, **, and * indicate significance at 1%, 5%, and 10% levels. Panel B reports the predicted proportions in which PDs used the 1/32nds, 1/64ths, 1/128ths, and 1/256ths pricing grids, based on the method described in Section 3.4.3.

Table 3.6: Price-end type distributions conditional on a pricing grid used

	Conditional on using $grid\text{-}g = \ldots$			
	$grid$-$32nds$	$grid$-$64ths$	$grid$-$128ths$	$grid$-$256ths$
$\phi_{32}^{grid\text{-}g} = \Pr[d \in X_{32} \mid grid\text{-}g]$	1	0.5	0.25	0.125
$\phi_{64}^{grid\text{-}g} = \Pr[d \in X_{32} \mid grid\text{-}g]$	0	0.5	0.25	0.125
$\phi_{128}^{grid\text{-}g} = \Pr[d \in X_{128} \mid grid\text{-}g]$	0	0	0.5	0.25
$\phi_{256}^{grid\text{-}g} = \Pr[d \in X_{256} \mid grid\text{-}g]$	0	0	0	0.5

This table summarizes the conditional probabilities of price ending d in X_{32}, X_{64}, X_{128}, and X_{256}, for each of the four pricing grids, *grid-32nds*, *grid-64ths*, *grid-128ths*, *grid-256ths*.

Table 3.7: Definitions of the variables for the analysis of the determinants of pricing fineness

Variable	Description	Data source
D_{32}	Dummy variable: 1 if the decimal part of the offer price is in $X_{32} = \{0/256, 8/256, 16/256, \dots, 248/256\}$, 0 otherwise.	FRBNY
D_{64}	Dummy variable: 1 if the decimal part of the offer price is in $X_{64} = \{4/256, 12/256, 20/256, \dots, 252/256\}$, 0 otherwise.	FRBNY
D_{128}	Dummy variable: 1 if the decimal part of the offer price is in $X_{128} = \{2/256, 6/256, 10/256, \dots, 254/256\}$, 0 otherwise.	FRBNY
D_{256}	Dummy variable: 1 if the decimal part of the offer price is in $X_{256} = \{1/256, 3/256, 5/256, \dots, 255/256\}$, 0 otherwise.	FRBNY
$Topfive$	Dummy variable: 1 if the PD is among the top five dealers based on the trade amount in QE auctions in the sub-period, 0 otherwise.	FRBNY
Cheapness	Yield-curve-implied price minus market mid-price, divided by market mid-price (based on the previous trading day's close price; in percent; winsorized at the 2.5% and 97.5%).	CRSP
On-the-run	Dummy variable: 1 if the security is an on-the-run security, 0 otherwise.	TreasuryDirect
$Maturity_{0\text{-}1}$	Dummy variable: 1 if the remaining maturity is in (0,1)Y, 0 otherwise.	TreasuryDirect
$Maturity_{1\text{-}2}$	Dummy variable: 1 if the remaining maturity is in [1,2)Y, 0 otherwise.	TreasuryDirect
$Maturity_{2\text{-}3}$	Dummy variable: 1 if the remaining maturity is in [2,3) years, 0 otherwise.	TreasuryDirect
$Maturity_{3\text{-}4}$	Dummy variable: 1 if the remaining maturity is in [3,4) years, 0 otherwise.	TreasuryDirect
$Maturity_{4\text{-}5}$	Dummy variable: 1 if the remaining maturity is in [4,5) years, 0 otherwise.	TreasuryDirect
$Maturity_{5\text{-}10}$	Dummy variable: 1 if the remaining maturity is in [5,10) years, 0 otherwise.	TreasuryDirect
$Maturity_{10\text{-}20}$	Dummy variable: 1 if the remaining maturity is in [10,20) years, 0 otherwise.	TreasuryDirect
$Maturity_{20\text{-}30}$	Dummy variable: 1 if the remaining maturity is in [20,30) years, 0 otherwise.	TreasuryDirect
Bond	Dummy variable: 1 if the security is a Treasury bond, 0 if it is a Treasury note.	TreasuryDirect
Bid-ask	Bid-ask spread, divided by the mid quote (based on the previous trading day's end-of-the-day bid and ask quotes; in percent; winsorized at the 2.5% and 97.5% levels).	CRSP
Volatility	Standard deviation of the security's returns in the previous five trading days (percent; winsorized at the 1% and 99% levels).	CRSP
Ln(offer amount)	Natural logarithm of the offered amount.	FRBNY
Ln(outstanding)	Natural logarithm of the publicly-held outstanding par value of the offered CUSIP.	CRSP
Ln(total purchases)	Natural logarithm of the total purchase amount of the QE auction.	FRBNY
Offer-to-cover	Offer-to-cover ratio of the QE auction (winsorized at the 2.5% and 97.5% levels).	FRBNY

This table lists the definitions and data sources of the variables used in the subsequent cross-sectional analysis of price-end fineness of QE auction offers.

Table 3.8: Descriptive statistics of the variables for the analysis of the determinants of pricing fineness

	Mean	S.d.	Min.	Median	Max.	N
D_{32}	0.325	0.468	0.000	0.000	1.000	90,022
D_{64}	0.230	0.421	0.000	0.000	1.000	90,022
D_{128}	0.218	0.413	0.000	0.000	1.000	90,022
D_{256}	0.227	0.419	0.000	0.000	1.000	90,022
$D_{128or256}$	0.445	0.497	0.000	0.000	1.000	90,022
$Topfive$	0.479	0.500	0.000	0.000	1.000	90,022
Cheapness	0.180	0.238	-0.251	0.101	0.723	89,964
On-the-run	0.043	0.204	0.000	0.000	1.000	90,022
$Maturity_{0\text{-}5}$	0.218	0.413	0.000	0.000	1.000	90,022
$Maturity_{5\text{-}10}$	0.365	0.481	0.000	0.000	1.000	90,022
$Maturity_{10\text{-}20}$	0.051	0.220	0.000	0.000	1.000	90,022
$Maturity_{20\text{-}30}$	0.366	0.482	0.000	0.000	1.000	90,022
Bond	0.462	0.499	0.000	0.000	1.000	90,022
Bid-ask	0.043	0.015	0.012	0.044	0.074	89,964
Volatility	0.552	0.455	0.027	0.434	2.079	89,281
Ln(offer amount)	23.153	1.782	20.723	23.431	29.242	90,022
Ln(outstanding)	24.013	0.620	21.409	24.122	25.041	85,974
Ln(total purchases)	22.023	0.797	19.052	22.032	23.942	90,022
Offer-to-cover	2.968	0.993	1.459	2.773	5.778	90,022

This table reports the descriptive statistics of the variables used in the subsequent cross-sectional analysis of price-end fineness of QE auction offers. See Table 3.7 for variable definitions. *Cheapness*, *Bid-ask*, *Volatility*, and *Offer-to-cover* are winsorized at the 2.5% and 97.5% levels.

Table 3.9: Ordered logit regression of the determinants of price-end fineness of QE auction offers

Panel A: Ordered logit regression of *Price-end fineness*						
	(1)		(2)		(3)	
	Coef.	Marg. eff.	Coef.	Marg. eff.	Coef.	Marg. eff.
$Topfive$	0.895***	0.149***	0.929***	0.151***	0.966***	0.156***
	(0.217)	(0.039)	(0.216)	(0.038)	(0.228)	(0.038)
Cheapness	-1.072***	-0.179***	0.429***	0.069***	0.252*	0.041*
	(0.226)	(0.029)	(0.142)	(0.021)	(0.138)	(0.022)
On-the-run			-0.385***	-0.062***	-0.532***	-0.086***
			(0.122)	(0.024)	(0.118)	(0.024)
$Maturity_{0\text{-}5}$			0.623***	0.117***	0.410***	0.072***
			(0.089)	(0.020)	(0.076)	(0.016)
$Maturity_{10\text{-}20}$			-0.419***	-0.063***	-0.405***	-0.059***
			(0.113)	(0.018)	(0.152)	(0.019)
$Maturity_{20\text{-}30}$			-0.456***	-0.068***	-0.139	-0.022
			(0.096)	(0.016)	(0.134)	(0.021)
Bond			-0.443**	-0.072**	-0.462**	-0.075***
			(0.219)	(0.030)	(0.201)	(0.027)
Bid-ask					0.484	0.078
					(0.657)	(0.105)
Volatility					-0.484***	-0.078***
					(0.104)	(0.017)
Ln(offer amount)					-0.066	-0.011
					(0.060)	(0.010)
Ln(outstanding)					-0.006	-0.001
					(0.039)	(0.006)
Ln(total purchases)					-0.012	-0.002
					(0.058)	(0.009)
Offer-to-cover					0.125***	0.020***
					(0.027)	(0.004)

(*Continued*)

Table 3.9: *Continued*

MEP	0.476	0.070	0.919***	0.136***	0.900***	0.129***
	(0.309)	(0.052)	(0.261)	(0.050)	(0.227)	(0.043)
QE3	0.812***	0.132***	1.011***	0.153***	0.930***	0.134***
	(0.214)	(0.038)	(0.197)	(0.035)	(0.186)	(0.033)
QE pause	0.898***	0.149***	0.871***	0.127***	0.749***	0.103***
	(0.207)	(0.037)	(0.203)	(0.033)	(0.206)	(0.032)
QE4	0.270	0.038	0.238	0.029	0.613***	0.081***
	(0.193)	(0.028)	(0.172)	(0.022)	(0.212)	(0.031)
N	89,964		89,964		85,866	
Pseudo R^2	0.036		0.053		0.060	

Panel B: The association of *Cheapness* and *Price-end fineness* for *Top five* = 1 and *Top five* = 0

		Marginal effect of *Cheapness* on the probability of *Price-end fineness* =...			
		1	2	3	4
	Coef.	($\iff D_{32} = 1$)	($\iff D_{64} = 1$)	($\iff D_{128} = 1$)	($\iff D_{256} = 1$)
Sample: *Top five* = 1					
Cheapness	-0.633***	0.112**	0.038***	-0.013	-0.136***
	(0.210)	(0.044)	(0.011)	(0.016)	(0.040)
Period dummies	✓				
N	43,077				
Pseudo R^2	0.017				
Sample: *Top five* = 0					
Cheapness	-1.540***	0.357***	-0.026	-0.157***	-0.173***
	(0.207)	(0.045)	(0.018)	(0.032)	(0.032)
Period dummies	✓				
N	46,887				
Pseudo R^2	0.017				
Test of the difference in the marginal effects					
χ^2 statistics		17.33***	6.11**	22.92***	0.94

This table estimates the ordered logit model in which the dependent variable is *Price-end fineness*, which takes the value of one if $D_{32} = 1$, two if $D_{64} = 1$, three if $D_{128} = 1$, and four if $D_{256} = 1$. Panel A uses all sample (accepted) offers in QE auctions and reports the coefficients and marginal effects on the probability of *Price-end fineness* = 4 (i.e., $D_{256} = 1$). For the definitions of the variables, see Table 3.7. In Panel B, the sample is split by *Top five*. In addition to the coefficients, the table reports the marginal effects on the probabilities of *Price-end fineness* = 1, 2, 3, and 4. Standard errors are three-way clustered by CUSIP, auction date, and PD. They are reported in parentheses. Standard errors for marginal effects are obtained by using the delta method. At the bottom, the differences in the marginal effects between $Top\ five = 1$ and $Top\ five = 0$ are tested. To preferm this test, I pool the two sub-samples and run the ordered logit model in which the independent variables are *Top five*, *Cheapness*, the dummies for the sub-periods, the interaction term of *Top five* and *Cheapness*, and those of *Top five* and the sub-period dummies. ***, **, and * indicate significance at 1%, 5%, and 10% levels.

Table 3.10: Price-end clustering on coarser grids of secondary market prices

Panel A: Regression results						
	All	Sub-period				
		QE2	MEP	QE3	QE pause	QE4
	(1)	(2)	(3)	(4)	(5)	(6)
$D_{32}^{Secondary}$	1.159***	1.426***	1.465***	1.481***	0.848***	0.708***
	(0.017)	(0.045)	(0.030)	(0.031)	(0.067)	(0.025)
$D_{64}^{Secondary}$	1.173***	1.304***	1.513***	1.500***	0.748***	0.768***
	(0.015)	(0.043)	(0.034)	(0.038)	(0.059)	(0.025)
$D_{128}^{Secondary}$	0.187***	0.186***	0.074***	0.072***	0.350***	0.309***
	(0.005)	(0.013)	(0.007)	(0.006)	(0.030)	(0.010)
Constant	-0.338***	-0.388***	-0.391	-0.391	-0.287***	-0.262***
	(0.002)	(0.001)	(.)	(.)	(0.012)	(0.004)
N	256	256	256	256	256	256
Adjusted R^2	0.990	0.947	0.978	0.976	0.674	0.927
$D_{32}^{Secondary} - D_{64}^{Secondary}$	-0.014	0.122*	-0.048	-0.019	0.100	-0.059*
	[0.400]	[3.807]	[1.126]	[0.146]	[1.315]	[2.950]
$D_{64}^{Secondary} - D_{128}^{Secondary}$	0.986***	1.119***	1.439***	1.427***	0.399***	0.459***
	[4052.203]	[621.373]	[1693.223]	[1397.315]	[38.994]	[310.069]
Panel B: Estimated proportions in which different pricing grids are used (%)						
grid-32*nds*	0.00	3.90	0.00	0.00	3.21	0.00
grid-64*ths*	62.66	71.60	90.56	90.75	25.51	27.46
grid-128*ths*	23.95	23.76	9.44	9.25	44.75	39.55
grid-256*ths*	13.38	0.74	0.00	0.00	26.53	32.99

Panel A tests the price-end clustering on coarser grids of my sample Treasury securities in the secondary market. For each Treasury security purchased in a QE auction, I obtain that security's closing ask price on the trading day preceding the QE auction from Bloomberg. The regression specifications are identical to Specification 3.2, except that this analysis uses not offer prices in QE auctions but secondary market prices. The dependent variable, $Percent_d^{Secondary} - 0.390625$, is the percentage of price endings being d, minus 0.390625. The right-hand side variables are $D_{32}^{Secondary}$, $D_{64}^{Secondary}$, and $D_{128}^{Secondary}$, which take the value of one if the price ending d belongs to X_{32}, X_{64}, and X_{128}, respectively, and zero otherwise. Heteroskedasticity-robust standard errors are reported in parentheses. At the bottom, differences between coefficients are tested and their F statistics are reported in bracket parentheses. ***, **, and * indicate significance at 1%, 5%, and 10% levels. Panel B reports the predicted proportions in which PDs used the 1/32nds, 1/64ths, 1/128ths, and 1/256ths pricing grids, based on the method described in Section 3.4.3.

Table 3.11: Ordered logit regression of the determinants of price-end fineness of secondary market prices

Dependent variable: *Price-end fineness*^{Secondary}		Marginal effect on the probability of *Price-end fineness*^{Secondary} =...			
	Coefficient	1 ($\Longleftrightarrow D_{32}^{Sec.}=1$)	2 ($\Longleftrightarrow D_{64}^{Sec.}=1$)	3 ($\Longleftrightarrow D_{128}^{Sec.}=1$)	4 ($\Longleftrightarrow D_{256}^{Sec.}=1$)
On-the-run	-0.179	0.034	-0.009	-0.015	-0.011
	(0.136)	(0.026)	(0.006)	(0.012)	(0.008)
$Maturity_{0\text{-}5}$	1.506***	-0.269***	-0.014	0.186***	0.097***
	(0.117)	(0.019)	(0.014)	(0.012)	(0.014)
$Maturity_{10\text{-}20}$	-0.345**	0.081**	-0.042**	-0.030**	-0.009**
	(0.145)	(0.034)	(0.019)	(0.012)	(0.004)
$Maturity_{20\text{-}30}$	-0.435***	0.102***	-0.055***	-0.036***	-0.011***
	(0.133)	(0.031)	(0.017)	(0.012)	(0.004)
Bond	-0.131	0.025	-0.006	-0.011	-0.008
	(0.135)	(0.026)	(0.006)	(0.012)	(0.008)
MEP	0.696***	-0.145***	0.059***	0.060***	0.026***
	(0.118)	(0.024)	(0.012)	(0.009)	(0.005)
QE3	0.556***	-0.116***	0.050***	0.047***	0.019***
	(0.118)	(0.024)	(0.013)	(0.009)	(0.004)
QE pause	1.177***	-0.238***	0.075***	0.109***	0.054***
	(0.177)	(0.035)	(0.013)	(0.017)	(0.011)
QE4	1.250***	-0.251***	0.075***	0.116***	0.059***
	(0.133)	(0.027)	(0.013)	(0.012)	(0.008)
N	15,872				
Pseudo R^2	0.114				

This table reports the result of estimating the ordered logit model of the price-end fineness of my sample Treasury securities in the secondary market. For each Treasury security purchased in a QE auction, I obtain that security's closing ask price on the trading day preceding the QE auction from Bloomberg. Therefore, the outcome variable, *Price-end fineness*^{Secondary}, is the price-end fineness of the secondary market price and takes the value of one if $D_{32}^{Secondary}=1$, two if $D_{64}^{Secondary}=1$, three if $D_{128}^{Secondary}=1$, and four if $D_{256}^{Secondary}=1$. The remaining columns show the marginal effects on the probabilities of *Price-end fineness*^{Secondary} = 1, 2, 3, and 4. For the definitions of the explanatory variables, see Table 3.7. Standard errors are two-way clustered by CUSIP and auction date, and they are reported in parentheses. Standard errors for marginal effects are obtained by using the delta method. ***, **, and * indicate significance at 1%, 5%, and 10% levels.

Table 3.12: Ordered logit regression of the price-end fineness of QE auction offers on the price-end fineness of secondary market prices

Sample:	All	$Top\ five = 1$	$Top\ five = 0$	All	$Top\ five = 1$	$Top\ five = 0$
	(1)	(2)	(3)	(4)	(5)	(6)
$D_{64}^{Secondary}$	-0.017	-0.023	-0.030	-0.014	-0.027	-0.018
	(0.014)	(0.015)	(0.024)	(0.012)	(0.020)	(0.022)
$D_{128}^{Secondary}$	0.594***	0.440***	0.734***	0.114***	0.065	0.132**
	(0.075)	(0.096)	(0.093)	(0.042)	(0.062)	(0.058)
$D_{256}^{Secondary}$	0.904***	0.695***	1.164***	0.336***	0.214	0.468***
	(0.137)	(0.176)	(0.152)	(0.104)	(0.149)	(0.117)
Security-type controls				✓	✓	✓
Period dummies	✓	✓	✓	✓	✓	✓
N	90,022	43,096	46,926	90,022	43,096	46,926
Pseudo R^2	0.014	0.018	0.014	0.031	0.026	0.045

This table reports the coefficients of the ordered logit regressin in which the dependent variable is *Price-end fineness*, which takes the value of one if the QE auction offer's price ending is in X_{32}, two if it is in X_{64}, three if it is in X_{128}, and four if it is in X_{256}. $D_{64}^{Secondary}$ takes the value of one if the price ending of the Treasury security's closing ask on the trading day preceding the QE auction is in X_{64}, and zero otherwise. $D_{128}^{Secondary}$ and $D_{256}^{Secondary}$ are similarly defined. The secondary market price-end data come from Bloomberg. In Columns 4–6, the following scurity-type controls are included: *On-the-run*, $Maturity_{\text{0-5}}$, $Maturity_{\text{10-20}}$, $Maturity_{\text{20-30}}$, and *Bond*. Standard errors are three-way clustered by CUSIP, auction date, and PD. ***, **, and * indicate significance at 1%, 5%, and 10% levels.

Table 3.13: Market share and the predicted proportion of using fine pricing grids

Panel A: Sub-period analysis

Dependent var.:	*Pricing* $grid_{256}$				*Pricing* $grid_{128or256}$			
Period	QE2	MEP	QE3	QE4	QE2	MEP	QE3	QE4
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Market share	1.006	2.395***	3.384***	2.970**	1.332	2.071***	2.158***	1.193
	(0.774)	(0.293)	(0.859)	(1.069)	(0.975)	(0.519)	(0.504)	(0.993)
Constant	15.166**	25.941***	32.692***	31.388***	45.762***	48.700***	67.630***	56.978***
	(6.143)	(6.623)	(6.919)	(6.336)	(6.024)	(6.825)	(5.084)	(5.479)
N	18	19	21	23	18	19	21	23
Adjusted R^2	0.031	0.326	0.363	0.244	0.110	0.266	0.305	0.044

Panel B: Pooled analysis in levels

Dependent var.:	*Pricing* $grid_{256}$				*Pricing* $grid_{128or256}$			
Sample	All	Excl. Goldman Sachs	Excl. Morgan Stanley	Excl. < 2% mkt share	All	Excl. Goldman Sachs	Excl. Morgan Stanley	Excl. < 2% mkt share
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Market share	2.393***	2.431***	2.438***	2.402***	1.836***	1.703***	2.264***	1.665***
	(0.313)	(0.386)	(0.521)	(0.295)	(0.405)	(0.435)	(0.490)	(0.439)
Period dummies	✓	✓	✓	✓	✓	✓	✓	✓
N	81	77	77	55	81	77	77	55
Adjusted R^2	0.383	0.344	0.299	0.486	0.360	0.317	0.335	0.444

Panel C: Pooled analysis in first differences

Dependent var.:	Δ*Pricing* $grid_{256}$				Δ*Pricing* $grid_{128or256}$			
Sample	All	Excl. Goldman Sachs	Excl. Morgan Stanley	Excl. < 2% mkt share	All	Excl. Goldman Sachs	Excl. Morgan Stanley	Excl. < 2% mkt share
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ*Market share*	2.431***	2.413***	2.297**	2.574***	2.137***	2.284***	2.260***	1.908***
	(0.457)	(0.523)	(0.893)	(0.330)	(0.400)	(0.471)	(0.736)	(0.410)
Period dummies	✓	✓	✓	✓	✓	✓	✓	✓
N	63	60	60	38	63	60	60	38
Adjusted R^2	0.278	0.266	0.156	0.473	0.383	0.375	0.330	0.439

This table reports the results of the OLS regressions in which the dependent variable is the predicted proportion of using the finest 1/256ths pricing grid (*Price* $grid_{256}$) or that of using the pricing grids of either 1/128ths or 1/256ths (*Price* $grid_{128or256}$). A PD is included in the sample if it has at least 100 winning offers in the sub-period and if it was designated as a PD at the beginning of the sub-sample period. *Market share* is the trade amount-based market share of the PD in the sub-period. In Panel A, heteroskedasticity-robust standard errors are reported in parentheses. In Panels B and C, standard errors clustered for PD are reported in parentheses. ***, **, and * indicate significance at 1%, 5%, and 10% levels.

Table 3.14: Binomial tests for the existence of offers undercutting the most coarsely priced offers

Period	QE2	MEP	QE3	QE pause	QE4	Total
(a) N. of offers: one tick below X_{32}	309	1,786	2,065	138	1,234	5,532
(b) N. of offers: one tick above X_{32}	264	1,448	1,865	142	1,034	4,753
(c) Fraction: (a) over (a) + (b)	0.539	0.552	0.525	0.493	0.544	0.538
(d) p value of the binomial test						
H_0: (c) = 0.5	[0.066]	[0.000]	[0.001]	[0.858]	[0.000]	[0.000]

This table presents the results of the binomial tests of the null hypothesis that there are equal numbers of (winning) offers whose price endings are one tick below X_{32} (i.e., $\{7, 15, 23, \ldots, 255\}$) and those whose price endings are one tick above X_{32} (i.e., $\{9, 17, 25, \ldots, 249\}$).

Table 3.15: Between-PD variation in the proportion of undercutting offers

	All				PDs w/ at least 200 sample offers			
	(1)		(2)		(3)		(4)	
	Coef.	S.e.	Coef.	S.e.	Coef.	S.e.	Coef.	S.e.
On-the-run	0.126***	(0.042)			0.113*	(0.054)		
$Maturity_{0\text{-}5}$	-0.007	(0.016)			0.001	(0.017)		
$Maturity_{10\text{-}20}$	0.121***	(0.030)			0.118***	(0.032)		
$Maturity_{20\text{-}30}$	0.055***	(0.019)			0.051*	(0.023)		
Bond	-0.063	(0.037)			-0.062	(0.042)		
Period dummies								
MEP	0.025	(0.031)			0.026	(0.036)		
QE3	-0.013	(0.034)			-0.018	(0.041)		
QE pause	-0.050	(0.047)			-0.051	(0.055)		
QE4	0.017	(0.034)			0.008	(0.042)		
PD dummies								
(Ref. cat. = Morgan Stanley)								
Amherst Pierpont	-0.025	(0.036)	0.050	(0.068)				
BMO	0.000	(0.014)	0.028	(0.037)	0.001	(0.015)	0.002	(0.038)
BNP Paribas	0.134***	(0.036)	0.059	(0.070)				
Bank of Nova Scotia	0.037	(0.028)	-0.023	(0.057)				
Barclays	0.162***	(0.012)	0.169***	(0.036)	0.163***	(0.012)	0.182***	(0.040)
BofA/Merrill	0.024	(0.023)	0.023	(0.047)	0.026	(0.024)	0.001	(0.055)
Cabrera	0.409***	(0.038)	0.347***	(0.079)				
Cantor Fitzgerald	0.045*	(0.025)	0.047	(0.044)				
Citigroup	-0.013	(0.010)	0.014	(0.025)	-0.012	(0.011)	0.015	(0.027)
Credit Suisse	0.037**	(0.016)	0.055	(0.036)	0.037**	(0.015)	0.059	(0.037)
Daiwa	0.049	(0.032)	0.035	(0.068)				
Deutsche Bank	0.005	(0.018)	0.030	(0.046)	0.004	(0.015)	0.027	(0.051)
G.X. Clarke	0.271***	(0.071)	0.057	(0.035)				
Goldman Sachs	-0.002	(0.021)	0.027	(0.038)	-0.002	(0.023)	0.021	(0.044)
HSBC	0.022	(0.021)	0.003	(0.059)	0.026	(0.020)	0.011	(0.072)
J.P. Morgan	0.044	(0.031)	0.074	(0.057)				
Jefferies	-0.092***	(0.030)	-0.016	(0.059)				
MF Global	-0.070	(0.073)	0.013	(0.027)				
Mizuho	0.015	(0.035)	-0.043	(0.064)				
Nomura	0.051**	(0.024)	0.072	(0.051)	0.051*	(0.025)	0.060	(0.060)
RBC	0.086***	(0.030)	0.193***	(0.053)				
RBS/NatWest	0.006	(0.025)	0.105**	(0.047)				
Societe Generale	-0.013	(0.032)	-0.059	(0.083)				
TD	0.016	(0.047)	0.139	(0.118)				
UBS	0.018	(0.025)	0.001	(0.042)				
Wells Fargo	0.120***	(0.035)	0.068	(0.079)				

(*Continued*)

Table 3.15: *Continued*

Constant	0.514***	(0.031)		0.517***	(0.037)	
CUSIP × QE auction FE			✓			✓
N (excl. singleton obs.)	10,285		7,026	8,641		5,784
Adjusted R^2	0.011		0.351	0.010		0.385
Joint test of H_0:						
Each PD dummy = 0						
F stat	181.67***		5.51***	338.58***		3.49**
p value	[0.000]		[0.000]	[0.000]		[0.040]

This table reports the OLS regression results for Specification 3.4. The sample consists of offers whose price endings are in $X_{UNDER32} = \{7, 15, 23, \ldots, 255\}$ or in $X_{OVER32} = \{9, 17, 25, \ldots, 249\}$. Models 1 and 2 use all PDs, whereas Models 3 and 4 use only PDs with at least 200 sample (winning) offers for this analysis. The dependent variable, $Undercut_{32}$, takes the value of one if the price ending is in $X_{UNDER32}$, and zero if it is in X_{OVER32}. In Models 1 and 3, the right-hand side variables are *On-the-run*, $Maturity_{0\text{-}5}$, $Maturity_{10\text{-}20}$, $Maturity_{20\text{-}30}$, and *Bond*, sub-period dummies, and PD dummies. Models 2 and 4 have PD dummies and CUSIP × QE auction fixed effects. Standard errors are three-way clustered by CUSIP, QE auction date, and PD. At the bottom of the table, the null that each of the PD dummy coefficients is equal to zero is tested.

Table 3.16: Coarse pricing and the level of prices among accepted offers in QE auctions

Panel A: Summary statistics of *Price diff*

	N	Mean	S.d.	Min	25%tile	Median	75%tile	Max
Price diff	86,611	3.61	3.93	0	0.81	2.62	4.82	18.50

Panel B: Regressions of *Price diff*

	(1)	(2)	(3)	(4)
D_{32}	0.282*	0.278**	0.138	0.127
	(0.139)	(0.127)	(0.117)	(0.107)
D_{64}	0.247**	0.243**	0.119	0.108
	(0.117)	(0.108)	(0.110)	(0.100)
D_{128}	0.181**	0.178**	0.116*	0.109*
	(0.073)	(0.068)	(0.062)	(0.056)
Ln(offer amount)		0.035		0.050
		(0.052)		(0.045)
Top five			-0.415**	-0.431**
			(0.177)	(0.165)
CUSIP × QE auction FEs	✓	✓	✓	✓
N	86,611	86,611	86,611	86,611
Adjusted R^2	0.566	0.566	0.569	0.569
$D_{32} - D_{64}$	0.035	0.035	0.020	0.019
	[0.484]	[0.481]	[0.195]	[0.175]
$D_{64} - D_{128}$	0.066	0.066	0.002	-0.001
	[0.978]	[1.062]	[0.001]	[0.000]

Panel A shows the summary statistics of *Price diff*, which is the percentage difference (in basis points) between the offer price and the minimum accepted price of offers for the same security in the same QE auction. *Price diff* is defined only when there exist multiple winning offers for the security in the QE auction. This variable is winsorized at the 2.5% and 97.5%. Panel B reports the OLS regression results with the dependent variable *Price diff*. All the models include CUSIP × QE-auction fixed effects. D_{32}, D_{64}, and D_{128} take the value of one if the price ending d belongs to X_{32}, X_{64}, and X_{128}, respectively, and zero otherwise. *Ln(Total assets)* is the natural logarithm of the offer amount. *Top five* takes the value of one if the PD is one of the top five dealers based on the trade amount in QE auctions in the sub-period, and zero otherwise. Standard errors are three-way clustered by CUSIP, auction date, and PD. They are reported in parentheses. At the bottom, differences between coefficients are tested and their F statistics are reported in bracket parentheses. ***, **, and * indicate significance at 1%, 5%, and 10% levels.

Appendices

Appendix 3.A: A brief history of the Fed's QE purchases of Treasury securities

Table 3.2 lists major events of the Fed's QE purchases of Treasury securities. The phases can be classified as follows. First, on March 18, 2009, the Fed announced the commencement of purchasing up to $300 billion of Treasury coupon securities. The purchases of this so-called "QE1" round began on March 25 and ended as planned six months later (on October 29).[37] The Fed then announced on August 10, 2010 that it would reinvest proceeds from maturing Treasury securities and other debt into Treasury coupon securities to maintain its balance sheet size. The Fed resumed Treasury coupon security purchases following the announcement.

Second, the "QE2" phase was announced on November 3, 2010. The announced purchase size of Treasury coupon securities was $600 billion. As planned, QE2 ended in June 2011, with the last purchase taking place on June 22. Proceeds from maturing debt continued to be reinvested into Treasury securities.

Third, the Maturity Extension Program (MEP) was announced on September 21, 2011.[38] While the MEP still purchased long-term Treasury securities, the MEP differed from the preceding programs in that it did not change the size of bank reserves—the purchases were funded by sales of shorter-term Treasury securities (maturing in less than 3 years). The original plan was to purchase $400 billion of long-term Treasury securities by June 2012. It was announced on June 20, 2012, however, that the Fed would continue the purchases (and sales of shorter-term Treasuries) at the same pace until December 2012.

Fourth, the Fed announced the replacement of the MEP with "QE3" on December 12, 2012. Like QE1 and QE2, QE3 did not entail sales of shorter-term Treasury securities. Unlike them, QE3 was open-ended as it specified the (initial) monthly purchase amount, $45 billion, but neither the total purchase size nor the termination date.[39] QE3 was concluded on October 29, 2014.

[37] The purchases of agency MBS and agency debt continued until March 2010.

[38] The MEP is also referred to as the "Operation Twist" after a similar program the Fed launched in 1961.

[39] The announcement says, "If the outlook for the labor market does not improve substantially, the Committee will continue its purchases of Treasury and agency mortgage-backed securities" (https://www.federalreserve.gov/

Fifth, in the years following the QE3 conclusion, purchases of Treasury coupon securities were sporadic and quite small. Moreover, purchases of long-term Treasuries were completely halted following the initiation of the balance sheet normalization program in October 2017. The balance sheet normalization program was then concluded in August 2019, and the Fed resumed reinvesting proceeds of maturing securities in Treasury coupon securities (with the first purchase taking place on August 15, 2019). The reinvestment continued at the pace of roughly $15 billion per month until February 2020. The FOMC emphasized, however, that the purpose of Treasury security purchases during this period was to maintain a sufficient level of bank reserves, instead of restarting QE (Bernanke, 2022, pp. 248–252).

Lastly, the COVID-19 crisis led the Fed to launch massive operations in the Treasury market. On March 12, 2020, the Fed announced that, among other interventions, it would purchase Treasury securities of various maturities from the next day. The massive new round of QE, which this paper dubs "QE4," was announced on March 15. According to it, the Fed would purchase at least $500 billion of Treasury securities "over coming months." While previous QE rounds had the main policy objective of stimulating the economy through lowering long-term rates, that of QE4 was different. The March 12 announcement clarified that the purpose was "to address highly unusual disruptions in Treasury financing markets associated with the coronavirus outbreak."[40]

Appendix 3.B: Fitting a Fed-style yield curve model

Song and Zhu (2018) estimate a yield curve model based on the Fed's public information. Unless stated otherwise, I replicate their method to the greatest extent possible. The model is a standard cubic spline model. There are two important choices to consider. First, some prefer setting a positive smoothness parameter, which effectively prioritizes curve smoothness at the expense of the model's fit. Following the baseline approach of Song and Zhu (2018), I set the smoothness parameter to zero. The second consideration is the number and location of knots. Again, I employ their choice: 2, 5, 10, 20, and 30 years.

newsevents/pressreleases/monetary20121212a.htm).

[40] https://www.newyorkfed.org/markets/opolicy/operating_policy_200312a

Like Song and Zhu (2018), I exclude Treasury bills and on-the-run securities from the yield curve estimation. One difference between their sample selection and mine is the remaining maturity. Song and Zhu (2018) discard securities whose remaining maturities are less than one year. This selection makes sense in their setting; Treasury securities nearing maturity occasionally exhibit large, idiosyncratic price fluctuations, and during the period they examine, the Fed did not purchase any securities maturing in less than 1.5 years. In contrast, in my sample period, the shortest maturity purchased by the Fed was 32 days. I therefore set a minimum remaining maturity restriction of one month.

While Song and Zhu (2018) use the Fed's internal secondary-market price data (New Price Quote System), my data source is the CRSP Treasury data. For each QE auction, I estimate the yield curve using the daily closing mid price of the previous trading day. I then measure Treasury securities' cheapness by comparing the model-implied prices and their actual closing mid prices on the day of the yield curve estimation. Note that cheapness is measured even for on-the-run securities, which are not included in the yield curve estimation.

Appendix 3.C: Dealer-level estimation of the proportions of pricing grids used

To quantify the extent of coarse pricing for each PD while controlling for offered security heterogeneity, I employ a more elaborate version of Specification 3.2. The dependent variable, $Percent_{s,j,t,d} - 0.390625$, is the percentage of offers with price endings being d among all offers for security type s submitted by PD j during period t, minus 0.390625. More specifically, three basic Treasury security characteristics are taken into account: on-the-run status, remaining maturity, and Treasury note vs. bond. Therefore, the ratio is calculated for each possible combination of security type s, PD j, and period t. The full regression specification is:

$$\begin{aligned} Percent_{s,j,t,d} - 0.390625 = \beta_1 D_{32} + \beta_2 D_{64} + \beta_3 D_{128} + \beta_4 On\text{-}the\text{-}run_s + \beta_5 Maturity_{\text{0-5}Y} \\ + \beta_6 Maturity_{\text{10-20}Y} + \beta_7 Maturity_{\text{20-30}Y} + \beta_8 Bond \\ + (\beta_9 D_{32} + \beta_{10} D_{64} \beta_{11} D_{128}) \times On\text{-}the\text{-}run \\ + (\beta_{12} D_{32} + \beta_{13} D_{64} + \beta_{14} D_{128}) \times Maturity_{\text{0-5}Y} \\ + (\beta_{15} D_{32} + \beta_{16} D_{64} + \beta_{17} D_{128}) \times Maturity_{\text{10-20}Y} \\ + (\beta_{18} D_{32} + \beta_{19} D_{64} + \beta_{20} D_{128}) \times Maturity_{\text{20-30}Y} \\ + (\beta_{21} D_{32} + \beta_{22} D_{64} + \beta_{23} D_{128}) \times Bond \\ + \gamma Z_{j,t} + (\delta D_{32} + \zeta D_{64} + \eta D_{128}) \times Z_{j,t} + \epsilon_{i,d}, \end{aligned} \tag{3.6}$$

where D_{32}, D_{64}, and D_{128} take the value of one if the cell's price ending d ($\in D = \{0, 1, 2, \ldots, 255\}$) belongs to X_{32}, X_{64}, and X_{128}, respectively, and zero otherwise; $On\text{-}the\text{-}run_s$ is the dummy variable for on-the-run securities; $Maturity_{\text{0-5}Y}$, $Maturity_{\text{10-20}Y}$, and $Maturity_{\text{20-30}Y}$ are dummy variables for remaining maturities being in (0, 5), [10, 20), and [20, 30) years, respectively; *Bond* is the dummy variable for Treasury bonds; $Z_{j,t}$ is the fixed effects for PD $\times$ period. (In this PD-level analysis I merge the QE-pause period into the QE3 period due to its small sample size.) Since $Percent_{s,j,t,d}$'s are calculated from different numbers of offers, the regression weights the observations by the number of offers in the cell. The regression result is reported in Table A3.1.

Based on Specification 3.6, the baseline security is an off-the-run Treasury note maturing in [5, 10) years, which is the most common security type in my sample of QE auction offers. Moreover, by combining the coefficients reported in Table A3.1 with the coefficients of $Z_{j,t}$ and the interaction terms of $Z_{j,t}$ with D_{32}, D_{64}, and D_{128}, I can predict the price-end type distributions at the PD $\times$ period level in the case of offering the most common Treasury security type. These predicted PD $\times$ period-level values are then fed into the maximum likelihood procedure detailed in Section 3.4.3 for estimating the proportions in which the four pricing grids were used.

Table A3.2 summarizes the estimated PD $\times$ period-level proportions in which the four pricing grids were used. There are two notable time-series patterns (which are in line with Table 3.5).

First, PDs tend to price more finely over time. The median PD increases the use of the 1/256ths pricing grid from 15.6% in *QE2* to 29.6% in *MEP* and 48.5% in *QE3*. Second, the pricing fineness deteriorated in the QE4 period. The median probability of using the coarsest 1/32nds grid increased from 0% in *QE3* to 18.2%.

Table A3.1: OLS estimation of Specification 3.6

	Coefficient	Standard error
D_{32}	1.132***	(0.037)
D_{64}	1.038***	(0.023)
D_{128}	0.492***	(0.013)
On-the-run	-0.062***	(0.020)
$Maturity_{\text{0-5}Y}$	0.054***	(0.012)
$Maturity_{\text{10-20}Y}$	-0.025**	(0.012)
$Maturity_{\text{20-30}Y}$	-0.027*	(0.015)
Bond	-0.029	(0.018)
On-the-run $\times D_{32}$	0.185*	(0.106)
On-the-run $\times D_{64}$	0.168	(0.106)
On-the-run $\times D_{128}$	0.071*	(0.042)
$Maturity_{\text{0-5}Y} \times D_{32}$	-0.440***	(0.073)
$Maturity_{\text{10-20}Y} \times D_{32}$	0.192**	(0.092)
$Maturity_{\text{20-30}Y} \times D_{32}$	0.109**	(0.053)
$Maturity_{\text{0-5}Y} \times D_{64}$	-0.165***	(0.039)
$Maturity_{\text{10-20}Y} \times D_{64}$	0.011	(0.077)
$Maturity_{\text{20-30}Y} \times D_{64}$	0.042	(0.058)
$Maturity_{\text{0-5}Y} \times D_{128}$	0.087***	(0.033)
$Maturity_{\text{10-20}Y} \times D_{128}$	-0.001	(0.030)
$Maturity_{\text{20-30}Y} \times D_{128}$	0.033	(0.034)
Bond $\times D_{32}$	0.415***	(0.093)
Bond $\times D_{64}$	0.074	(0.055)
Bond $\times D_{128}$	-0.130***	(0.024)
$Z_{j,t}$	✓	
$Z_{j,t} \times (\delta D_{32} + \zeta D_{64} + \eta D_{128})$	✓	
Number of observations before weighting	161,024	
Number of observations after weighting	23,045,632	
Adjusted R^2	0.324	

This table reports the OLS estimation of Specification 3.6. The dependent variable, $Percent_{s,j,t,d} - 0.390625$, is the percentage of offers with price endings being d among all offers for security type s submitted by PD j during period t, minus 0.390625. D_{32}, D_{64}, and D_{128} take the value of one if the price ending d belongs to X_{32}, X_{64}, and X_{128}, respectively, and zero otherwise. *On-the-run* takes the value of one if the offered Treasury security is on-the-run. $Maturity_{\text{0-5}Y}$, $Maturity_{\text{10-20}Y}$, and $Maturity_{\text{20-30}Y}$ are dummy variables for remaining maturities being in (0, 5), [10, 20), and [20, 30) years, respectively. *Bond* takes the value of one if the offered security is a Treasury bond and zero if it is a Treasury note. $Z_{j,t}$ is the fixed effects for PD $\times$ period. Given that the ratios are calculated from different numbers of offers, the OLS regression is weighted by the respective number of offers. Standard errors clustered by price ending (d) and PD (j) are reported in parentheses. ***, **, and * indicate significance at 1%, 5%, and 10% levels.

Table A3.2: Predicted proportions of using the 1/32nds, 1/64ths, 1/128ths, and 1/256ths pricing grids for each PD

	N	Mean	S.d.	Min	p25	Median	p75	Max
Panel A: Estimated proportion of using the 1/32nds grid								
QE2	19	8.70	8.40	0.00	3.83	5.54	13.11	35.78
MEP	21	10.29	14.95	0.00	0.00	6.50	15.11	62.80
QE3	23	3.44	6.47	0.00	0.00	0.00	4.52	20.64
QE4	23	18.87	10.82	0.00	9.54	18.15	25.25	47.89
Panel B: Estimated proportion of using the 1/64ths grid								
QE2	19	38.21	14.45	8.39	27.83	38.02	46.60	70.05
MEP	21	29.06	22.62	0.00	8.86	29.09	46.30	70.35
QE3	23	18.14	14.60	0.00	8.14	16.47	30.78	53.65
QE4	23	18.98	7.46	1.19	14.77	19.36	23.88	33.76
Panel C: Estimated proportion of using the 1/128ths grid								
QE2	19	32.75	12.82	13.51	25.38	30.53	39.11	63.52
MEP	21	21.01	9.21	0.00	15.95	22.00	25.52	40.21
QE3	23	29.74	15.21	0.93	19.49	31.94	35.50	72.98
QE4	23	17.88	10.98	0.00	11.35	16.54	23.35	45.42
Panel D: Estimated proportion of using the 1/256ths grid								
QE2	19	20.34	16.94	3.22	7.73	15.60	25.20	60.04
MEP	21	39.64	28.77	8.37	18.36	29.59	52.24	100.00
QE3	23	48.68	25.79	14.04	26.55	48.48	66.17	99.07
QE4	23	44.28	18.27	15.44	33.30	41.01	51.79	88.00

This table reports the estimated proportions of using the 1/32nds, 1/64ths, 1/128ths, and 1/256ths pricing grids, for each PD in each sub-period, in the case of offering an off-the-run Treasury note maturing in [5, 10) years. Therefore, N refers to the number of sample PDs in the sub-period. Note that in this analysis the QE3 period incorporates the QE-pause period due to a small sample size of the latter. A PD is included in the sample only if it has at least 100 winning offers in the sub-period. See text for the method of estimating these proportions.

Appendix 3.D: Dealer characteristics and coarse pricing

Data

Does the degree of pricing fineness relate to the PD's characteristics (other than QE auction market share)? To explore this question, I look at three dimensions of PD characteristics.

Location of the parent bank: First, I examine whether PDs have a foreign parent company or not, referring to the PD lists of He et al. (2017) and Giannone and Robotti (2022). *Foreign* takes the value of one if the PD has a foreign parent bank, and zero otherwise.

Number of years as a primary dealer: Second, To shed light on varying experiences of PDs in the Treasury market, I measure the number of years as a PD. One challenge is how to handle mergers & acquisitions, which have been pervasive from time to time in this industry. I thus checked the historical primary dealer lists of the FRBNY. Table A3.3 summarizes the designation dates of domestic PDs, and Table A3.4 those of foreign PDs. The identification of the designation date is particularly challenging for some foreign PDs. Therefore, for some foreign PDs, I assign both the baseline and earliest start dates. *Old* takes the value of one if the (baseline) PD designation date is earlier than the median date, and zero otherwise. The use of the earliest start dates, instead of the baseline dates, does not much change the result.

Balance sheet size: Third, I hand-collected each PD's balance sheet information from Form X-17A-5, as done by Gupta (2022). An advantage of using this form is that I can obtain balance sheet data of dealer subsidiaries, as opposed to the group-wide consolidated financial data reported in 10-Ks, etc. Note that this data is also available for foreign bank-affiliated PDs because they are still incorporated as a U.S. subsidiary, and large U.S. broker-dealers are mandated to file (audited) Form X-17A-5. I was able to locate all my sample PDs' Form X-17A-5, except for that of the Bank of Nova Scotia, although there exist some gap years and non-reported items for some PDs.[41] *Ln(Total assets)* is the natural logarithm of the total assets measured at the beginning of each sub-period.

[41] Ownership structures of PDs were examined using the National Information Center repository (`https://www.ffiec.gov/NPW`). See Avraham et al. (2012) for its details.

Table A3.3: Domestic primary dealers

Name	PD designation period	Note
Amherst Pierpont	5/6/2019–	
BofA/Merrill	5/19/1960–	Merrill Lynch Government Securities was a primary dealer from 5/19/1960 to 2/11/2009. In 2009, Merrill Lynch was acquired by Bank of America, which also owned a primary dealer, Banc of America Securities. Banc of America Securities had roots in CRT Government Securities, which was designated as a primary dealer on 12/22/1987.
Cantor Fitzgerald	8/1/2006–	
Citigroup	5/19/1960–	Citigroup was created as a result of a merger between Citicorp and Travelers Group in 1998. Citicorp owned Citicorp Securities, which was designated as a primary dealer under the name of Citibank on 6/15/1961. Travelers Group owned Salomon Smith Barney, which had roots in Salomon Brothers. Salomon Brothers was designated as a primary dealer on 5/19/1960.
Goldman Sachs	12/4/1974–	
Jefferies	6/18/2009–	
JP Morgan	5/19/1960–	
MF Global	2/2/2011– 10/31/2011	MF Global filed for bankruptcy on 10/31/2011.
Morgan Stanley	2/1/1978–	
Wells Fargo	4/18/2016–	

This table summarizes domestic primary dealers during my sample period (2010Q3–2020Q2). The four (domestic) dealers that participated in the primary dealer pilot program for July 2013–July 2014 are not included. For domestic PDs, the period as a PD is the same whether the start date is based on the baseline or earliest one. See text for data sources.

Table A3.4: Primary dealers owned by foreign parent companies

Name	PD designation period: Baseline	PD designation period: Earliest	Note
Bank of Nova Scotia	10/4/2011–	10/4/2011–	
Barclays	12/7/1989–	12/7/1989–	Barclays De Zoete Wedd was designated as a primary dealer on 12/7/1989.
BMO	10/4/2011–	10/4/2011–	For a brief period (2/15/2000–3/31/2002), BMO Nesbitt, a subsidiary of BMO, was a primary dealer. Given the brevity of the period and the time gap, the start date is determined to be that of BMO.
BNP Paribas	5/1/1997–	5/1/1997–	Paribas was designated as a primary dealer on 5/1/1997. The name was changed to BNP Paribas on 9/15/2000 as a result of a merger.
Credit Suisse	12/23/1988–	5/19/1960–	On 12/23/1988, First Boston was merged into CS First Boston, which was established as a London-based joint venture in 1978. CS Holding effectively controlled the new CS First Boston with a 44.5% equity stake. (Credit Suisse's ownership was further increased to 60% in 1990.) First Boston had been a primary dealer since 5/19/1960.
Daiwa	12/11/1986–	12/11/1986–	
Deutsche Bank	12/13/1990–	5/19/1960–	Deutsche Bank Government Securities was designated as a primary dealer on 12/13/1990. In 1999, Deutsche Bank acquired Bankers Trust, which had been a primary dealer since 5/19/1960. (Note that this acquisition occurred well after Deutsche Bank became a primary dealer.)
HSBC	12/2/1983–	9/29/1976–	Marine Midland Bank, whose majority equity stake was owned by the Hongkong and Shanghai Banking Corporation (current HSBC), acquired Carroll McEntee & McGinley on 12/2/1983. Carroll McEntee & McGinley had been a primary dealer since 9/29/1976.
Mizuho	12/28/1989–	2/13/1980–	Following Fuji Bank's acquisition, Kleinwort Benson Government Securities was renamed Fuji Securities on 12/28/1989. Kleinwort Benson was designated as a primary dealer on 2/13/1980. (Fuji Securities was renamed Mizuho Securities on 4/1/2002.)
Nomura	12/11/1986 –11/30/2007, 7/27/2009–	12/11/1986 –11/30/2007, 7/27/2009–	Nomura once quit as a primary dealer on 11/30/2007 and regained a primary dealer status on 7/27/2009. Given the brevity of the discontinuation period, the start date is considered to be the date on which Nomura was first designated as a primary dealer (12/11/1986).
RBC	7/8/2009–	7/8/2009–	
RBS/NatWest	3/6/2000–	7/31/1984–	On 3/6/2000, RBS acquired NatWest, which owned a primary dealer Greenwich Capital. Greenwich Capital had been a primary dealer since 7/31/1984, and it was acquired by NatWest in 1996.
Societe Generale	2/2/2011–	2/2/2011–	From 7/1/1999 to 10/31/2001, its subsidiary (SG Cowen) was a primary dealer. Given the brevity and the time gap, the start date is determined to be that of SG Americas.
TD	2/11/2014–	2/11/2014–	
UBS	12/7/1989–	6/24/1988–	UBS Securities was designated as a PD on 12/7/1989. In 1998, SBC and UBS merged, with the new entity name being UBS. Before the merger, UBS Securities, a subsidiary of UBS, had been a primary dealer since 12/7/1989. SBC also owned S. G. Warburg & Co., which had been a primary dealer since 6/24/1988, as a result of its acquisition in 1995.

This table summarizes primary dealers owned by foreign parent companies during my sample period (2010Q3–2020Q2). See text for data sources.

Result

Table A3.5 shows the correlation matrix of PD-level variables. The results of pooled regressions are reported in Table A3.6.

Table A3.5: Correlation matrix of PD-level variables

	Price $grid_{256}$	*Price* $grid_{128 or 256}$	*Market share*	*Ln(Total assets)*	*Foreign*	*Old*
Price $grid_{256}$	1					
Price $grid_{128 or 256}$	0.843*** [0.000] (81)	1				
Market share	0.457*** [0.000] (81)	0.428*** [0.000] (81)	1			
Ln(Total assets)	0.254** [0.030] (73)	0.271** [0.020] (73)	0.615*** [0.000] (73)	1		
Foreign	–0.300*** [0.006] (81)	–0.169 [0.132] (81)	–0.313*** [0.004] (81)	–0.158 [0.181] (73)	1	
Old	0.211* [0.059] (81)	0.076 [0.498] (81)	0.363*** [0.001] (81)	0.488*** [0.000] (73)	–0.343*** [0.002] (81)	1

This table presents the correlation matrix of PD-level variables, pooling the data from the four sub-periods. *Price* $grid_{256}$ is the predicted proportion of using the finest 1/256ths pricing grid, and *Price* $grid_{128or256}$ is that of using the pricing grid of either 1/128ths or 1/256ths. *Market share* is the trade amount-based market share of the PD in the sub-period. See Section 3.6 for the definitions and data sources for the other variables. A PD is included in the sample only if it has at least 100 winning offers in the sub-period. The number of observations used for calculating the pair-wise correlation is shown in round brackets and p-values in square brackets.

Table A3.6: PD characteristics and the predicted proportions of using fine pricing grids

Panel A: Dependent variable: *Pricing* $grid_{256}$

	(1)	(2)	(3)	(4)	(5)
Market share	2.393***				2.084***
	(0.313)				(0.414)
Ln(Total assets)		7.513*			
		(4.112)			
Foreign			-16.004*		-8.603
			(7.958)		(7.469)
Old				12.025	1.492
				(7.330)	(5.738)
MEP	18.436***	17.370***	17.700***	17.689***	18.514***
	(4.625)	(5.909)	(5.991)	(5.999)	(4.790)
QE3	29.844***	29.541***	28.799***	28.769***	30.178***
	(5.301)	(6.276)	(5.437)	(5.505)	(5.385)
QE4	26.378***	28.257***	23.528***	24.874***	26.156***
	(5.230)	(6.796)	(4.825)	(4.981)	(5.367)
Constant	7.512	-169.328	30.495***	14.702***	13.731
	(5.202)	(106.113)	(7.406)	(4.574)	(9.768)
N	81	73	81	81	81
Adjusted R^2	0.383	0.223	0.232	0.191	0.397

Panel B: Dependent variable: *Pricing* $grid_{128 or 256}$

	(6)	(7)	(8)	(9)	(10)
Market share	1.836***				1.878***
	(0.405)				(0.490)
Ln(Total assets)		5.826			
		(3.386)			
Foreign			-8.062		-3.277
			(6.335)		(6.752)
Old				4.008	-3.981
				(5.772)	(5.357)
MEP	6.910	5.591	6.259	6.199	6.890
	(5.653)	(6.360)	(6.530)	(6.537)	(5.751)
QE3	26.165***	25.414***	25.129***	24.967***	26.096***
	(4.519)	(5.537)	(4.713)	(4.746)	(4.517)
QE4	11.204**	11.952**	9.027*	9.482**	10.812**
	(4.108)	(5.134)	(4.546)	(4.322)	(4.140)
Constant	42.982***	-93.763	58.037***	51.106***	46.745***
	(4.340)	(86.682)	(6.698)	(3.851)	(8.699)
N	81	73	81	81	81
Adjusted R^2	0.360	0.226	0.185	0.158	0.354

This table reports the results of the regressions in which the dependent variable is the predicted proportion of using the finest pricing grid of 1/256ths (*Price* $grid_{256}$) or that of using the pricing grids of either 1/128ths or 1/256ths (*Price* $grid_{128 or 256}$). A PD is included in the sample only if it has at least 100 winning offers in the sub-period. *Market share* is the trade amount-based market share of the PD in the sub-period. *MEP*, *QE3*, and *QE4* are indicator variables for the sub-periods. (*QE2* is the reference category.) See Section 3.6 for the definitions and data sources for the other variables. Standard errors clustered for PD are reported in parentheses. ***, **, and * indicate significance at 1%, 5%, and 10% levels.

Appendix 3.E: Simulation analysis of price-end fineness and QE auction-winning probability

The Fed publicly discloses only *accepted* offers, and my analysis is based on this sub-sample. However, my primary interest lies in understanding the extent of PDs' coarse pricing, irrespective of the QE auction outcomes of the submitted offers. Unfortunately, this data truncation introduces a bias in this inference, as offers with different price-end fineness can have different auction-winning probabilities. This section employs a simple simulation exercise to gauge this possible bias.

The average QE auction accepts 95.9 offers, with the offer-to-cover ratio being 3.63 (Table 3.3). The simulation therefore assumes that in each auction 360 offers are submitted and 100 are accepted. I assume that 10% of offers come from the 1/32nds pricing grid, 25% from the 1/64ths grid, 20% from the 1/128ths grid, and the remaining 45% from the 1/256ths grid. Again, these numbers are chosen so that the simulated environment is close to the average QE auction.[42] These pricing-grid proportions indicate that the distribution of price-end types based on *all* offers should be 33.125% in X_{32}, 23.125% in X_{64}, 21.25% in X_{128}, and 22.5% in X_{256}.

Each PD's unrounded valuation of a Treasury security is assumed to be independently and normally distributed.[43] For simplicity, I assume that the mean is the same as the security's benchmark price derived from the Fed's (proprietary) yield-curve model.[44] Recall that the FRBNY ranks offers based on the price difference from the benchmark price, combining offers for all target securities in the particular QE auction. Therefore, under the assumption that each offer price comes from an independent and normal distribution around the benchmark price, the fact that a QE auction accepts multiple CUSIPs becomes irrelevant for analyzing auction outcomes—I can analyze as if

[42] Column 1 of Panel B in Table 3.5 reports the predicted grid-use proportions based on all sample QE offers. Those proportions are estimated to be 9.5% for the 1/32nds grid, 24.2% for the 1/64ths grid, 20.9% for the 1/128ths grid, and 45.4% for the 1/256ths grid.

[43] In practice, coarsely-pricing PDs do not necessarily observe these unrounded valuations; they might always think based on their coarser pricing grids. This formalization is adopted only to permit simpler modeling of the valuation distribution; it circumvents the need to model a discrete valuation distribution for each pricing grid type by considering an "underlying" continuous distribution.

[44] Because offers whose prices are lowest relative to their benchmark prices are accepted first, this assumption means that the average markup of accepted offers would be negative. Shifting the distribution mean can make the average price markup positive, but this change would not affect the main implications of this simulation exercise.

all offers are perfect substitutes.[45]

The distribution's standard deviation is determined in such a way that the standard deviation of the *accepted* offers' price markup (relative to the benchmark price) in this simulation matches what I can infer from observable data. Specifically, I calculate the standard deviation of the following measure: the accepted offer price minus the security's end-of-day ask price in the secondary market. That is, I use the standard deviation of this measure as a proxy for the standard deviation of winning offers' markups relative to the (unobservable) Fed benchmark prices. The end-of-day ask prices are obtained from CRSP. After trimming the bottom and top 0.5%, the standard deviation of this price difference measure is 0.0578.

The distribution of PDs' unrounded valuations of Treasury securities has been specified. Offer prices are based on realized valuations, but with different fineness. In the case of PDs using the 1/256ths pricing grid, the offer price is the realized valuation rounded to the nearest 1/256ths grid. In contrast, coarsely-pricing PDs' offer prices are further rounded. I assume they round price endings in either of the following two ways: (i) rounding to the nearest coarse price grid and (ii) rounding up to the closest coarse price grid. For illustration, consider how a PD using the 1/32nds pricing grid would round if the realized unrounded valuation is $\$105 + 3/256$. Based on the first method, this PD would truncate the decimal part, whereas the second method would round up the price to $\$105 + 1/32$.[46]

Table A3.7 summarizes the result of simulating this hypothetical QE auction market 50,000 times. In each auction, the 100 lowest-priced offers are accepted. Panels A shows that the average proportions of the price-end types of winning offers (in simulation) are different but still fairly similar to the proportions of all offers (which are simulation input parameters). The differences are particularly small when coarsely-pricing PDs employ near-rounding. For example, the difference in the proportion of X_{256} is only 0.06 percentage points. Indeed, in this case, the differences remain small for a relatively wide range of parameters, and even the direction of the bias (i.e.,

[45] The assumption is violated if offers for certain Treasury securities tend to have higher/lower price markups (relative to the security's benchmark price) than offers for other securities. Although Song and Zhu (2018) show that this is the case (higher markups for "cheaper" securities), this issue is omitted to simplify the analysis.

[46] The up-rounding case can be regarded as coarsely-pricing PDs being particularly loss averse.

whether finely-priced offers are over- or under-weighted in the sample of winning offers) depends on input parameters (most particularly the standard deviation of the distribution of the unrounded valuations).

Second, if coarsely-pricing PDs round up offer prices, the differences are slightly larger than in the previous case. This result is to be expected because if coarsely-pricing PDs employ up-rounding, the mean of their post-rounding valuation distribution should be greater than the mean of their pre-rounding valuation distribution. At the same time, the magnitudes of the differences still remain small; the difference in the proportion of X_{256} is still 1.01 percentage points.

To facilitate interpretation, Panel B predicts dealers' grid-use proportions from the values in Panel A based on the method described in Section 3.4.3. When estimated from *winning* offers (Columns 2 and 3), the predicted proportion of using the coarsest 1/32nds pricing grid is moderately downward biased, whereas that of using the finest 1/256ths grid is slightly upward biased. However, the sizes of the biases are still fairly small. Whereas the 'true' proportion of using the finest 1/256ths pricing grid is 45%, the proportion estimated from winning offers is 45.13% for the near-rounding case and 47.02% for the up-rounding case.

Panel C documents the bias with respect to the analysis of whether offers whose price endings are on the coarsest 1/32nds grid are undercut by finely priced offers. Recall that in this simulation *no* dealers strategically submit offers whose price endings are one-tick below X_{32} (i.e., $d \in X_{UNDER32} = \{7, 15, 23, \ldots, 255\}$). Nevertheless, if we look at *winning* offers, there are more price endings in $X_{UNDER32}$ compared to the price endings in $X_{OVER32} = \{9, 17, 25, \ldots, 249\}$ (Columns 2 and 3 of Panel C). This result reflects their different auction-winning probabilities.

Table A3.7: Simulation analysis of the discrepancy between all offers and winning offers

	All offers (input parameters)	Simulation result of winning offers	
		Near-rounding	Up-rounding
	(1)	(2)	(3)
Panel A: Price-end type distribution (%)			
X_{32}	33.125	32.53	31.72
X_{64}	23.125	23.54	22.94
X_{128}	21.25	21.37	21.82
X_{256}	22.5	22.56	23.51
Panel B: Predicted grid-use proportions (%)			
Pricing grid = 1/32nds	10	8.98	8.78
Pricing grid = 1/64ths	25	25.72	24.06
Pricing grid = 1/128ths	20	20.17	20.14
Pricing grid = 1/256ths	45	45.13	47.02
Panel C: Undercutting of price endings in X_{32}			
$\frac{\text{N. of offers with price endings in } X_{UNDER32}}{\text{N. of offers with price endings in } X_{UNDER32} \text{ or } X_{OVER32}}$	0.5	0.5201	0.5241

This table presents the results of simulating the hypothetical QE auction 50,000 times. Column 1 lists the values based on all offers, and they are not simulation results but simulation input parameters. Column 2 shows the simulation result when coarsely-pricing PDs round offer prices to the nearest value on their pricing grid. Column 3 is the simulation result when coarsely-pricing PDs always round up prices. Panel A reports the average distribution of price-end types from the 50,000 simulation rounds. Panel B presents the results of predicting proportions of PDs' grid use based on the values in Panel A using the method described in Section 3.4.3. Panel C reports the proportion of offers whose price endings are in $X_{UNDER32} = \{7, 15, 23, \ldots, 255\}$ relative to offers whose price endings are in either $X_{UNDER32}$ or $X_{OVER32} = \{9, 17, 25, \ldots, 249\}$, based on the 50,000 simulated auctions.

Chapter 4

Empowering Women by Index Membership: Evidence from a Unique Experiment from Japan

"*Creating an environment in which women find it comfortable to work [...] is no longer a matter of choice for Japan. It is instead a matter of the greatest urgency.*"

—Abe Shinzō, speaking to the United Nations in September 2013

4.1 Introduction

Investing with a social good has gained tremendously over the last decade.[1] Matos (2020) cites changing societal preference as a strong driving force behind such a shift. In this paper, we ask a simple question: can specially crafted equity indices bring about real changes in corporate social behavior? Specifically, we focus on the MSCI Empowering Women Index (WIN) in Japan that uses "*women's participation and advancement in the workforce*" as its primary membership crite-

[1]See for e.g. Bansal et al. (2021) for a recent review. Their findings indicate a pro-cyclical preference for socially responsible investments. See also Berk and van Binsbergen (2021) for the impact of socially responsible divestitures on a firm's cost of capital.

rion (MSCI, 2019a). The WIN was launched in 2017 with the support of the Government Pension Investment Fund of Japan (GPIF) to encourage Japanese companies to improve women's participation in the workforce.[2] The index features a quasi-tournament-like structure in that it is hived off the top half of the MSCI Japan IMI Top 700 Index, roughly corresponding to the S&P 500. Each firm in the IMI 700 is ranked on its MSCI Gender Diversity Score relative to its industry, and the top 50% are included in the WIN. To shed light on whether indexation can bring about *real* changes to corporate social behavior, we examine whether the creation of the WIN has resulted in significant improvements in women's participation in the workforce. In our empirical tests we compare gender diversity performance for the marginal firm that either gains inclusion in the index or just misses it vis-à-vis firms that rank sufficiently low that exclusion from the index is a *fait accompli*. Thus, this difference-in-differences methodology affords us a plausible identification strategy in establishing causality.

Why would belonging to the WIN lead to changes in firm behavior and practices, or even be desirable? We posit two main channels through which this can happen. First, changing shareholder preferences regarding ESG commitments may compel firms to improve the share of women in the workforce. To wit, if the marginal shareholder cares more about gender diversity than they did before, they will place a higher value on firms with a greater gender diversity. In response, firms might oblige, as in the catering models of Baker et al. (2009).[3] But this raises the question of why firms would wait for the creation of the WIN before they invest in improvements in women's participation in the workforce. Couldn't a firm have done this independent of the WIN? We posit that status quo habits are difficult to change, especially in a country like Japan with tight cultures,[4] and often an external nudge may prove more effective in expediting change than slower moving secular trends might achieve on their own. The crux of the nudge argument is familiar to readers

[2] In 2016, among developed nations, Japan had one of the highest workforce gender gaps of 20.3%. In comparison, this was almost double the workforce gender gap in the United States (12.6%) and triple the average gap in Scandinavian countries (7.2%).

[3] Dyck et al. (2019) show that institutional investors drive corporate sustainability performance of firms, and especially, foreign investors domiciled in countries with high social norms towards sustainability. This channel is likely not the main driving force of real sustainability change in Japan as institutional investor ownership in Japan is low with total and foreign institutional ownership of 13.5% and 8.4%, respectively (see Table 1 in Dyck et al. (2019)).

[4] See Glosserman (2019) who notes a unique resistance to changes in Japanese culture.

of the literature[5] —for instance, Bhargava and Loewenstein (2015) point to psychological biases such as "*motivated disbelief, the ostrich effect, confirmation bias, present-bias, adaptation, and intangibility*" (p.p. 399) as roadblocks to quick adaptations towards optimal choices. Our results are consistent with such a conclusion.

Second, belonging to the WIN may increase the firm's visibility to investors, especially to those who share the goals of the index—for example, large investors such as the GPIF in Japan, the world's largest pension fund, with ¥191 trillion (≈$1.75 trillion) in assets under management in 2021,[6] or large foreign institutional investors pledging to consider a firm's sustainability performance in their investment decisions. The GPIF is on record in 2017 to allocate ¥1 trillion towards investing in indices based on three ESG factors, one of which is the WIN. The increased investor attention for firms with greater social performance is expected to lead to greater institutional ownership and a lower cost of capital (Pastor et al., 2021).[7]

There is also the possibility that managers experience personal disutility from being excluded from the WIN. If the index is structured as a tournament, and exclusion from the index is associated with 'shame,' firms will compete to be included. For instance, Chattopadhyaya et al. (2020) show that firms competed to be included in the JPX-Nikkei400 index because the inclusion bestowed the status of Japan's *Best Run* companies.

Our analysis focuses on the real effects of the creation of the WIN—that is, whether the WIN brought about real changes in Japanese firms' social behaviour, specifically in their gender performance in the workplace. Because there may be secular trends regarding the advancement and empowerment of women across all firms in Japan, we employ a regression discontinuity specification to provide directional evidence. To that end, we compare firms around the inclusion threshold

[5]See Thaler and Sunstein (2003), Camerer et al. (2003), and especially the eponymous book by Thaler and Sunstein (2008), for an insight into the literature on nudge and directed choices. The essence of this literature is that seemingly innocuous but purposefully designed nudges may induce far reaching and consequential choice switching among a vast class of agents. Put simply, nudges lead agents towards welfare enhancing choices without restricting their options. Nudges are the ultimate positive NPV investments.

[6]The GPIF introduced a new code in 2020 that specifically includes a focus on diversity and inclusion and requires firms to disclose gender statistics in line with the *2015 Act on Promotion of Women's Participation and Advancement in the Workplace*.

[7]In addition, incremental buying pressure by investors might elevate prices for index members under the assumption of downward sloping demand curves. See, among others, Shleifer (1986) and Kaul et al. (2000) for evidence supporting downward sloping demand curves for stocks.

to firms farther away from the threshold with little chance to be included in the WIN. Given the tournament-like structure of the WIN, the identification assumption is that firms around the threshold (those that barely made it or missed it) have an incentive to improve their gender diversity performance to be included in the WIN at the next rebalancing date (whereas the firms farther away from the threshold have little chance of being included).

Because MSCI only discloses inclusion in the WIN as a binary variable, we obtain the workforce gender diversity data from MSCI and re-create the gender diversity scores used by MSCI for each of the IMI 700 firms and assign ordinal ranks based on these scores for each firm in a given industry (MSCI GICS sector) and year.[8] We confirm that this synthetic ranking approach accurately predicts the actual WIN members with a correlation of 94%. Using these synthetic ranks, we identify *treated* firms as those that rank in the vicinity of the inclusion threshold (ranked between the 40th to 60th percentile; the threshold is the median), and *control* firms as those with a much lower probability of gaining inclusion (ranked between the 40th to 10th percentile). The difference-in-differences analysis compares the differences of various workforce gender diversity measures in these two groups between the years before the WIN's inauguration in July 2017 and the years after 2017 (we exclude the inauguration year). The sample period is 2013 to 2020.

We measure workforce gender performance with data obtained from the Toyo Keizai CSR Workforce database. Toyo Keizai, founded in 1895, is among the top two prominent publishers in Japan along with Nikkei that has published economic and business news for more than a century. Each year, Toyo Keizai launches a survey inviting companies listed on Japanese stock exchanges to participate. Based on the survey data as well as publicly available disclosures, Toyo Keizai compiles a panel of Japanese firms' workforce characteristics. The database contains rich workforce data with more than 200 line items in aggregate and many line items broken down by gender—for example, the number of employees, turnover of employees, number of employees by position in the workforce, and maternity/paternity leaves, to name a few. These data allow us to construct various workforce gender diversity outcome measures. We augment the workforce data with financial and

[8] See MSCI (2019a) for a description of the detailed index methodology and MSCI (2019b) for a description of the workforce gender diversity database.

accounting data from Worldscope and Datastream and institutional ownership from Bloomberg.

We begin our analysis by examining whether the WIN brought about real changes to firms' workforce gender diversity. Using our difference-in-differences design, controlling for firm characteristics and firm and time fixed effects, we find that treated firms (compared to control firms) significantly improved the fraction of women in the workforce following the launch of the WIN. In terms of economic significance, treated firms improved their fraction of women in the workforce by about 5% per year compared to control firms. A visual parallel trends analysis and regressions in event time confirm that the change happened in the years after the WIN was created. Thus, the increase is not due to long-run secular trends either across all firms, or within the treated firms themselves, lending credit to a directional interpretation of our results—–the WIN incentivised firms to significantly improve their overall gender diversity in the workforce. These effects are novel to the literature, documenting the social power of index creations.

A greater fraction of women in the workforce can be obtained by hiring more women or by terminating more men employees (while keeping the number of women employees constant). If the latter was true, observing a greater fraction of women in the workforce would indicate some 'greenwashing' behaviour of firms and would not be an indication of an improvement of women's participation in the workforce. To explore this further, we repeat our baseline tests focusing on the changes in the number of women and men in the workforce. Our results show that treated firms significantly increased the number of women in the workforce in the post period (compared to control firms). In contrast, the change in men employees as well as the turnover in men employees are insignificantly different from zero. Thus, the increase in women in the workforce is particularly explained by firms hiring more women in the post period.

In additional tests, we find that the increase in women in the workforce is specifically concentrated in supervisory positions (entry positions obtained by employees with a fresh undergraduate degree) as well as at senior levels, including general managers, executive officers, and directors. Thus, firms do not just hire more women at the lowest levels, which is promising for firms' future improvements in the workforce as Matsa and Miller (2011) document a trickle-down effect where firms with women directors are more likely to recruit women executives. We also document a sign of a shift toward a more family-friendly workplace culture. For instance, we find that men employees in treated firms are more likely to take paternity leaves in the post period. This indicates a shift in culture in that men employees are not afraid of losing their jobs because they take parental leaves (it is now more socially accepted), and this shift can further help women to stay in the workforce.

Our results thus far show the social power of index creation—firms significantly improved their workplace gender diversity performance after the launch of the WIN. Next, we examine changes in investor interest in firms included in the WIN. In particular, we shed light on whether WIN firms' institutional ownership and performance changed. To that end, we compare differences between WIN and non-WIN firms before and after WIN inclusion. Our main result is that institutional ownership growth is stronger for the included group vis-à-vis the excluded group of firms. In terms of economic significance, WIN firms' institutional ownership increased by 3 percentage points compared to non-WIN firms in the years after the launch of the WIN. Despite the increase in institutional ownership, we are unable to detect an increase in valuation (measured with Tobin's q and annual stock returns) for the included firms relative to excluded firms. While we find no positive valuation effects, our results also do not document that being included in the WIN and investing in greater workforce gender diversity hurts shareholders, an often-debated fact when it comes to greater investments in firms' social performance.

Our paper makes several contributions. Our findings inform investors, regulators, and academics interested in mechanisms capable of bringing about real change in corporate social behavior. We identify one specific channel, the creation of a tournament-style equity index linked to a firm's gender diversity performance, through which firms significantly improved their workforce gender diversity contributing to empowering women.

We add to the literature on gender diversity in corporations. Many studies focus on the effects of women in leadership positions on corporate cultures with greater gender equality (Tate and Yang, 2015; Kunze and Miller, 2017), more conservative investment and financing policies (Huang and Kisgen, 2013; Faccio et al., 2016), fewer incidents of lawsuits (Adhikari et al., 2019), and having a less sexist corporate culture, which is associated with higher firm valuations (Lins et al., 2022). These studies focus mostly on women in leadership positions and on the *outcomes* of having more women in such positions, whereas we highlight a potential first-order *channel* through which firms *improve* gender diversity in the workforce, including management and executive positions, potentially leading to the above documented outcomes.

We also contribute to the literature on board gender diversity and board gender quotas as a regulatory force to improve gender equality in the board room (Adams and Ferreira, 2009; Adams and Funk, 2012; Ahern and Dittmar, 2012; Kim and Starks, 2016; Bertrand et al., 2019). Our results highlight a *capital market* channel through which firms are motivated (e.g., by lowering the cost of capital or by increasing institutional ownership) to improve gender diversity in the workforce above and beyond the board of directors.

Our findings also advance the literature on the drivers behind the improvement of firms' sustainability. The literature has documented the roles of institutional investors (Dyck et al., 2019; Krueger et al., 2020), private engagements by investors (Dimson et al., 2015), governance mechanisms to renew the thinking of the board (Dyck et al., 2023), directors with foreign sustainability expertise (Iliev and Roth, forthcoming), legal environment (Liang and Renneboog, 2017; Ioannou and Serafeim, 2012), litigation (Freund et al., 2021), propagation through the supply chain (Dai et al., 2021; Schiller, 2018), and the influence of management (Cronqvist and Yu, 2017). Most of these studies focus on aggregate sustainability performance, and it is unclear whether specific material sustainability improvements are obtained. We focus on a narrow, well-defined outcome—greater women participation in the workforce—–and show that through a purposefully crafted equity index that ties inclusion to the desired outcome, firms change their corporate social behaviour and improve their gender diversity.

Lastly, we also extend the literature on index inclusion, which finds that inclusion in large equity indices, such as the S&P 500 or Russell 1000, increases institutional ownership, improves monitoring, and lowers the cost of capital through capital inflows by large index funds who need to purchase newly added securities. We specifically add to the debate in Bebchuk and Hirst (2019), who argue that index funds are insufficient drivers of stewardship to improve governance. Our findings show that indices, specifically created to focus on one narrowly defined goal, can have a positive impact on firms' social behaviour and performance—in our context, improving women's advancement in the workforce.

4.2 Institutional Details, the WIN, and Empirical Strategy

4.2.1 Institutional Details in Japan

Since the mid-2010s, the Government Pension Investment Fund of Japan (GPIF), the world's largest pension fund with ¥191 trillion ($\approx$ \$1.75 trillion) in assets under management in 2021, has increased its focus on investor stewardship and ESG substantially. First, the GPIF signed the UN Principles for Responsible Investment (PRI) in September 2015.[9] Given the GPIF's large presence in the domestic market, Kawaguchi (2015) discusses a possible spillover effect that promotes ESG investments by other Japanese institutional investors. Second, the GPIF revised its investment principles in October 2017, with the new principle explicitly stating that its stewardship responsibilities include "the consideration of ESG (Environmental, Social, and Governance) factor" (GPIF, 2018).

To further enhance the GPIF's stewardship role, the pension fund was at the forefront of creating various ESG indices. In September 2016, the GPIF requested proposals for ESG indices of Japanese public equities with the promise of significant investments in chosen indices. Fourteen investment companies responded and proposed 27 indices, thereof the GPIF selected three indices to launch in July 2017. They are the FTSE Blossom Japan Index, MSCI Japan ESG Select Leaders Index, and MSCI Japan Empowering Women Index (WIN).[10] The WIN is the only thematic index, the other two are integrated indices based on broad ESG metrics.[11]

[9] https://www.gpif.go.jp/en/investment/pdf/signatory_UN_PRI_en.pdf

[10] The assets under management as of March 31, 2018 in the FTSE Blossom Japan Index, MSCI Japan ESG Select Leaders Index, and MSCI Japan Empowering Women Index (WIN) were ¥527 billion, ¥623 billion, and ¥388 billion, respectively (GPIF, 2018).

[11] In September 2018, the GPIF launched two additional indices focusing on carbon emissions, one for domestic equities and the other for foreign equities.

The GPIF's identity as a long-term, large, and diversified owner of domestic companies in Japan means that it needs to be attentive to curbing economy-wide negative externalities. The pension fund states that by creating ESG indices it aims to promote Japanese firms' ESG performance: "We are expecting that the use of those selected ESG indices will provide an incentive for Japanese companies to enhance responses to ESG issues to lead to the improvement of their corporate value in the long term" (GPIF, 2018, p. 40). To this end, the index managers are required to "publicly disclose how they conduct ESG evaluation and how they develop indices, and to proactively engage with companies" (GPIF, 2018, p. 41). In addition to the GPIF's efforts to enhance their stewardship role in Japanese companies, the Japanese government under President Abe pushed for empowering women to revitalize the Japanese economy.[12]

4.2.2 Gender Gap in Labor Contracts in Japan

Despite ranking close to the top decile for human development, Japan ranks in the bottom quartile of all nations in gender equality.[13] The reasons behind Japan's low scores on gender equality are complex. For instance, the post-war emphasis by corporate Japan on employment protection and seniority-based compensation meant that employees who would decide to leave the workforce, even temporarily, were *de facto* penalized in their careers (Crawford, 2021). Employees who chose to leave the regular work track faced the risk of losing their position in the lifetime employment track and often returned to non-regular work with lower pay and less protection from job termination. While this inefficiency could be overlooked in the post-war economic boom, the shifting demographics in the 1990s meant that the social cost of Japan's gender inequality started to rise. Crawford (2021) notes that the term *Womenomics* made its debut in 1999,[14] with a clear mandate of closing the economic gender gap and boosting productivity.[15] Politicians were quick to note the economic

[12] "Shinzo Abe: Unleashing the Power of '*Womenomics*'," Wall Street Journal, September 25, 2013.

[13] In 2020, Japan ranked 121 out of 153 countries on the Global Gender Gap Index compiled by the World Economic Forum. This was worse than its ranking in 2017 when it ranked 114 out 144 countries included in the survey. Japan's rank in the Human Development Index compiled by UN Development Programme was 20 out of the 183 countries in the survey in 2020.

[14] Kathy Matsui at Goldman Sachs Japan is credited with coining the term.

[15] The Japanese government's efforts to promote female labor participation in the past few decades can be traced back further to the enactment of the Equal Employment Opportunity Law (EEOL) in 1985. Although the original EEOL stated that firms should "endeavor" to avoid gender discrimination, the 1997 revision prohibited it. It was in

benefit of workplace gender equality, with ex-Prime Minister Shinzo Abe making Womenomics a key part of his pro-growth agenda. Womenomics was introduced in 1999; since then, the gains in workplace gender equality have been unimpressive (Crawford, 2021). Part of the problem is that there was little teeth in the recommendations to improve workplace gender equality. Even when the government announced the availability of small grants to encourage hiring and promotion of women, there were few takers.

We believe that the WIN offers a unique opportunity to realize the recommendations by offering the chance of a reward for the included firms, and equally importantly, an opportunity loss for the firms that fail to make the index. The quasi-tournament like structure of the index where only half of the firms in specific industry groups can gain membership precludes a box-checking response where everyone can claim victory by suitably defining their firm's gender equality scores. The emphasis on relative performance means that firms near the median find it incentive compatible to expend effort to improve their gender scores, assuming of course that WIN membership provides tangible benefits to the firm and its key stakeholders such as top executives and shareholders.

To the extent that taking time off from work for child raising presented a determinant to women's career choices, we want to examine the impact of the WIN on slow moving cultural shifts. In particular, we want to examine if the uptake of maternity and parental leaves changed in the wake of the WIN creation. Both represent hurdles in the path of women's career advancement as noted above, and a higher uptake of either policy would support the role of the WIN in nudging corporate culture in the direction of gender equality.

Finally, our interest is in determining if the creation of the WIN led to a shift in the nature of work for women. While women's participation in the labor force in Japan has indeed advanced remarkably over the last two decades to the extent that today it surpasses that in the U.S., critics have noted that women represent a majority of irregular work (Crawford, 2021). Regular jobs that provide employment stability and salary bonuses are disproportionately held by men. What is troubling is that regular employment represents a lower fraction of jobs held by women today than it did two decades ago (Crawford, 2021). In many ways, women have provided a cushion to

1999 that the revised EEOL came into effect.

corporate Japan's ability to withstand economic shocks over this period.

The government of Japan's own report on *Abenomics*[16] admits that less than 10% of managerial positions are filled by women—this fraction is considerably lower when the managerial positions are restricted to the C-suite.[17] Since the WIN explicitly takes into account the representation of women in managerial ranks, we are in a position to provide a report card on progress made in this dimension by treatment and control firms. We also extend the analysis to board representation.

In 2017, women held less than one out of 15 senior management positions in corporate Japan. The numbers improve at the middle and lower management cadres, but still remain at less than one out 10 and one out of five positions, respectively (see Table 4.2, Panel A). Part of the problem is structural and goes back to post-war economic boom in Japan where lifetime employment for regular workers was encouraged by firms. A downside of this practice was that employees who chose to leave the regular work track often lost their position in the lifetime employment track and often returned to non-regular work with lower pay and less protection from job termination.

We believe that the creation of a purposefully designed index with gender empowerment scores as the basis for qualification has the potential to bring out faster change than government edicts and exhortations that in practice lack the bite of enforcement penalties, relying instead on the softer principle of comply or explain. The quasi-tournament like structure of the WIN by design rules out an expansion of membership to all and sundry and supports the old adage that efforts to make everyone special ensures no one is.

4.2.3 The WIN and Empirical Strategy

To examine whether indexation leads to improvements in firms' social performance, we take advantage of the MSCI Empowering Women Index (WIN). As discussed above, the WIN was among the

[16]See Abenomics, Section 2, published by the government of Japan: `https://www.japan.go.jp/abenomics/index.html`.

[17]According to the 2017 Annual Report of the Ministry of Health, Labour, and Welfare, women held just 6.6% of senior management positions (department director or higher); 9.3% of middle management (section heads); and 18.6% of lower management (e.g., task unit supervisor) positions. Crawford (2021) writes that the issue of job status and pay gap for women in Japan is "[. . .] deeply structural and highly resistant to ordinary incremental inducements and exhortations (p. 7)."

three indices launched in July 2017 with the support of the GPIF. Each May and November, all firms included in the MSCI Japan IMI Top 700 Index are ranked within their industry sector (there are 11 MSCI GICS sector groups) based on their MSCI Gender Diversity Score. Companies with scores of zero or missing scores (unrated companies), REITs, and companies with ESG controversies are excluded from the universe of eligible companies. Using the eligible universe, companies with a gender diversity score equal or greater than the median score of all companies in the same industry sector are included in the WIN (MSCI, 2019a).

The MSCI Gender Diversity Score was first introduced by MSCI in July 2017 (MSCI, 2019b). The score aims to assess a company's overall practices to improve gender diversity and performance and ranges between zero and ten. It is calculated as the weighted average of the gender diversity performance score (with a weight of 75%) and the gender diversity practices score (weight of 25%). The performance score comprises performance metrics of a firm's attraction (% women employees among new hires and % women employees in total workforce), retention (% difference in average employment years for women and men employees), and promotion (% women in senior management and % women directors on the board) of women in the workforce. The performance score is the average of the five measures (each ranging from zero to ten). To account for firms' selective data disclosure, MSCI applies a disclosure discount—for example, if a firm only discloses three out of the five performance metrics, a 10% discount is applied. The gender diversity performance score ranges from zero to ten, with a greater score indicating better performance. To measure a firm's effort and intent to improve employment practices, MSCI calculates the MSCI Gender Diversity Practices Score that assesses the existence of workforce diversity policies, management oversight responsibility for diversity, and targets for improvement in women's representation and programs that make it easier for women work (e.g., flexible working arrangements, subsidized childcare, and parental leave). The score is calculated as the average of two policy variables (each ranging from zero to ten)——workforce diversity policy and management oversight and programs to increase workforce diversity. The gender diversity practices score ranges from zero to ten with higher values indicating better diversity practices.

In our main empirical tests to examine whether indexation based on gender diversity can bring about real changes to corporate social behaviour for the marginal firm, we compare changes in firms' gender diversity for firms that gain inclusion in the index or just misses it vis-à-vis firms that rank sufficiently low that the exclusion from the index is a *fait accompli*.

Given the tournament-like structure of the WIN, the identification assumption is that firms around the inclusion threshold have an incentive to improve their gender diversity performance to remain or to move into the WIN at the next rebalancing date, whereas the firms farther away from the threshold have little chance of being included. Because MSCI only discloses inclusion in the WIN as a binary variable, we obtain the workforce gender diversity data from MSCI and re-create the gender diversity scores used by MSCI for each of the IMI 700 firms and assign ordinal ranks based on these scores for each firm in a given industry (MSCI GICS sector) and year. Since the outcome variables we use are annual measures, we rank firms annually at the end of May. We then identify 'treated' firms as those that rank in the vicinity of the WIN inclusion threshold (ranked between the 40th to 60th percentile; the threshold is the median), and 'control' firms as those with a much lower likelihood of gaining inclusion (ranked between the 40th to 10th percentile). Our empirical approach relies on a difference-in-differences specification that compares the differences of various workforce diversity measures in these two groups between the years before the inauguration of the WIN in July 2017 and the years after 2017.

4.3 Sample and Summary Statistics

4.3.1 Toyo Keizai CSR Workforce Database

We obtain detailed workforce data for Japanese companies from the Toyo Keizai CSR Workforce database for the years 2013 to 2020 surrounding the launch of the WIN. The database covers a wide range of workforce-related line items. Toyo Keizai compiles these data based on annual CSR surveys of listed companies and some large private Japanese companies since 2005.[18] Each June/July, Toyo Keizai sends out a questionnaire to the surveyed companies requiring their response by August. To

[18] See Toyo Keizai's survey manual at http://www.toyokeizai.net/csr/pdf/ht_u/CSR_howtouse2018_1koyo.pdf, accessed at July 31, 2022.

facilitate the survey process, Toyo Keizai has created a specific website that explains how to answer the survey questions.[19] In addition, prior to sending out the questionnaire, usually in May, Toyo Keizai offers a workshop for all companies to provide guidelines in how to fill out the questionnaire. The response rate of all companies is quite high—for example, in 2020, 1,349 out of 3,819 companies filled out the questionnaire for an overall response rate of 35%, which increases with firm size. For non-responding firms that have responded in the prior two years and/or are included in the GPIF's ESG indices, Toyo Keizai fills the data from in-house data collections and public sources. Overall, the Toyo Keizai CSR data are widely used in Japan and highly regarded, mostly because of the high survey response rate and Toyo Keizai's century old reputation as a prominent publisher.

We use these data to create annual measures of women's participation in the workforce. The overall employment is measured with the *Fraction of Women in the Workforce*, calculated as the number of women employees divided by the number of all employees, and the *Ratio of Women to Men*, defined as the number of women employees divided by the number of men employees. We measure turnover of women and men employees with the *Change in Women (Men) Employees*, measured by the difference in the number of women (men) employees between year t and t–1 divided by the number of women (men) employees in year t–1, and the *Turnover in Women (Men) Employees*, measured as the number of women employees (men) leaving the company for reasons other than mandatory retirement over the year t–1 to year t divided by the number of women (men) employees in year t–1.

We also calculate four variables indicating various roles of women employees in the workforce ranging from positions in the top-two management teams, namely the board of directors (*torishima yaku*) and the executive operation officer committee (*shikko yakuin*),[20] general management (*kanri shoku*), and non-management and general workers (*ippan shoku*). Specifically, these variables are measured with *Board*, calculated as the number of women directors divided by the number of all directors (*torishimari yaku*), *Executives*, defined as the number of women executive officers divided

[19] http://www.toyokeizai.net/csr/research/No17-2021.html, accessed 31 July 2022 (in Japanese).

[20] It is not uncommon that some members in the executive officer team at the same time also served as a member of the board of director. For example, Akio Toyota, the CEO of Toyota Motors, served both in the officer team and the board as of 2022.

by the number of all executives (*shikko yakuin*), *General Management*, measured as the number of women with a general management position divided by the number of all general management positions (*kanri shoku*); and *Workers*, calculated as the number of women general workers divided by the number of all general workers.

Finally, we use the Toyo Keizai data to measure firms' maternity and paternity leaves. *Women Taking Maternity Leaves*, measured as the number of women employees who take maternity leaves in year t divided by the number of all employees in year t–1, and *Men Taking Paternity Leaves*, calculated as number of men employees taking paternity leaves in year t divided by the number of all employees in year t–1. We create two additional versions for each of the two variables dividing women (men) taking maternity (paternity) leaves by employees that are less than 30 years old and alternatively, less than 40 years old to account for the fact that older employees are less likely to take maternity/paternity leaves.

Our final sample covers the 2013-2020 period and consists of 723 firms for a total of 5,613 firm-year observations. Panel A of Table 4.2 shows summary statistics for the employment characteristics of all sample firms. Women account for 18% (14%) of the average (median) firm's total workforce. Women are relatively more represented at general worker positions, where they account for 23.2% of the workforce. This fraction drops to 4.4% for general management positions, and to 0.9% for executive positions. Four percent of board members are women. The median firm in our sample has no women in executive or board positions.

The number of women and men employees show a small increase of 0.6% and 0.5% over this period. Average employee turnover for women is lower (1.0%) compared to men (2.7%). Women taking maternity leave represent 1.3% of the workforce—this fraction is higher (5.1%) for women under 30. Men on parental leave represent 0.1% of the workforce. Men under 30 are six times more likely to avail of parental leave.

4.3.2 Firm Characteristics

In our analyses, we control for time-variant firm characteristics that may affect firms' workforce performance directly. We obtain financial, accounting, and stock market data from the Thomson Reuters Worldscope and Database databases. In all our tests, we control for *Firm Size* (log of total assets in ¥ billion), *Cash Holdings* (cash and cash equivalents divided by total assets), *Tangibility* (PP&E divided by total assets), *Financial Leverage* (total debt divided by total assets), ROA (net income divided by total assets), and *Tobin's q* (measured as the market value of equity plus the book value of debt divided by total assets). *Firm Size* may be related to gender scores because larger firms may be better positioned to attract women workers—–perhaps because smaller firms represent regional firms with different culture vis-à-vis national firms. *Cash Holdings* and *Leverage* are included as control variables because financial constraints may prevent firms from prioritizing gender scores. On the flip side, *ROA* and *Tobin's q* are included because profitable and higher-valued firms may have an advantage in attracting more talented employees, including women who may have a preference to work at such firms. Since these variables are not directly affected by the WIN inclusion, we feel safe in including them as controls in our regression models.

Panel B of Table 4.2 reports summary statistics for the full sample. We study large Japanese firms included in the IMI 700 index (the parent index from which the WIN draws its constituents) with an average (median) total assets of ¥2,871 billion (¥625 billion), corresponding to about \$27.8 (\$6.1) billion using the 2020 end-of-year exchange rate. The firms we study have average cash holdings of 16.3%, tangibility of 28.4%, and a financial leverage of 20.2%. The mean ROA is 4.2%. Average (median) institutional ownership is 48% (47.5%), which is lower than the average institutional ownership of more than 80% in S&P 500 firms. We note that the average institutional ownership we document is consistent with significant cross-shareholdings in the Japanese market, and with the high levels of insider ownership (including that of banks) documented in Franks et al. (2014).

4.4 Results

4.4.1 Did the WIN Increase Women's Participation in the Workforce?

We begin with plotting the fraction of women in the workplace from 2015 to 2019 for treated and control firms in Panel A of Figure 4.1. All firms are ranked based on their MSCI Gender Diversity Score (GDS) within their respective industry codes (see Table 4.3). Recall that only 50% of firms from each industry grouping are included in the WIN. We define treated firms as those that rank between the 40th and 60th GDS percentiles. Control firms are defined as those ranked in the 10th to the 40th GDS percentiles. Control firms have little chance of being included in the WIN—treated firms, by contrast, have a higher likelihood of being included in the WIN, depending on relative improvements in their GDS.

Prior to 2017, both groups show a similar GDS — the mean fraction of women employees is 18.1% and 17.9% for the treated and control firms. In 2015 too, the fraction of women employees is similar, albeit lower, for both groups indicating a rising trend for women participation in the workplace, but little difference in GDS for the two groups. Following the inception of the WIN, we see a remarkable divergent trend for the treatment and control groups. While the fraction of women employees remains at the pre-WIN level for the control group, it increases secularly over the two years following the WIN creation and exceeds 20.2% for the treatment group by 2019, while it is 17.7% for the control group. The difference between the treatment and control groups, while statistically indistinguishable in the pre-WIN period, shows a remarkable divergence in the post-WIN period. A similar pattern is observed when we plot the ratio of women to men employees from 2015 to 2018 in Panel B. Overall, Figure 4.1 provides a strong basis for the parallel trends condition for the difference-in-difference tests that follow in the subsequent tables regarding the effect of the WIN creation on the treatment vs. control sample of firms.

We discuss our baseline regression results shown in Table 4.4. The main dependent variable in these regressions is the fraction of women in the firm's workforce. We also re-estimate the regressions with the ratio of women to men employees as a dependent variable. The first four columns represent coefficient estimates from a panel regression. We also provide event year regressions in columns 5 and 6 where year dummies are measured relative to the year in which the treatment firm is added to the WIN. In all regressions, we include firm and time (year) fixed effects and the standard errors are double clustered by firm and year.

In columns 1 through 4 of Table 4.4, the main variable of interest is the treated firm interacted with the post-WIN inclusion period (years 2018 and 2019), *Treated* $\times$ *Post*. In column 1, we find that the interacted coefficient is positive and significant at the 5% level. We interpret this result as showing that treated firms are significantly more likely to increase the fraction of women workers relative to control firms in the two years following the creation of the WIN in 2017. We exclude 2017 from the analysis. In the next column, we re-estimate the regression with a slew of additional control variables such as firm size (total assets), liquidity position (cash), asset tangibility (property plant and equipment scaled by assets), leverage, ROA, and Tobin's q. The main variable of interest, *Treated* $\times$ *Post*, remains positive and significant. In columns 3 and 4, we repeat the regressions using the ratio of women to men in the workforce as the dependent variable. *Treated* $\times$ *Post* remains positive and significant in both models. We also note that the variable *Treated* by itself is not significant in any of the models, indicating that prior to the WIN creation, the treatment and control group firms have similar levels of women in the workforce.

In columns 5 and 6, we show regressions in event time to examine the time dynamics of the effects. Specifically, we include five additional time indicator variables denoting calendar years before and after the creation of the WIN in 2017 and interact these time indicators with the treated variable. We find that the time indicators interactions are significant only in the post-WIN inception years, and not before, confirming the findings in the first four columns. In other words, only the treated firms show an increase in the fraction of women employed at the firm in the aftermath of the creation of the WIN. Prior to the creation of the WIN, treated firms do not display any variance with the control group of firms.

Overall Table 4.4 shows the following: the fraction of a firm's workforce that is represented by women is similar for the treated and control firms prior to the creation of the WIN, but undergoes a shift in the post-WIN era. Specifically, treated firms display an incrementally higher fraction of women employment relative to the pre-WIN period. Importantly, this effect survives the contemporaneous change in the fraction of women in the workforce for the control group of firms.

4.4.2 Decomposing the Increase in Women's Participation in the Workforce

We next turn to examining whether the increase in the fraction of women in the workforce is associated with fewer men employees, or is due to hiring more women, or due to higher turnover among men, or some combination of all three. In Table 4.5, we replace the dependent variable in Table 4.4 with four separate dependent variables. These are the change in the number of women employees, the change in the number of men employees, the turnover in women employees, and the turnover in men employees. The variable of interest is again *Treated* × *Post* in each regression. In other words, we are interested in examining whether treated firms experience a change in any of the four dependent variables noted above.

Our results show that only the change in women employees has a significantly positive coefficient, confirming that treated firms experienced an increase in the number of women employees at the firm. The remaining three dependent variables are also positive, but not statistically significant, indicating that there is no significant change in the number of men employees or in the turnover of men or women employees. These results are inconsistent with the hypothesis that the increase in women employees comes at the expense of laying off men workers. Instead, we interpret the results as being consistent with treated firms actively pursuing policies that improve their gender diversity score in an effort to gain WIN membership. The effect survives any secular trends that may have also affected the control group of firms, and provides strong evidence that treated firms employ deliberate tactics to improve their odds of being included in the WIN.

4.4.3 Did Women Gain Job Status Following the Creation of the WIN?

We next turn to examining the breakdown of the increase in women employees documented above. We break down the workforce into various positions: workers, general management, executive officers, and board of directors. Again, the question we ask is whether treated firms seek to improve their gender diversity scores by actively increasing the fraction of women in senior positions. Table 4.6 presents the results.

Each of the four dependent variables corresponds to a specific workplace position. We begin with workers without any management ranks in column 1. Treated firms do not appear to experience a significant change in the number of non-management women workers. When we look at managerial positions and executives in columns 2 and 3, respectively, we get a different picture. The coefficient estimate on *Treated* $\times$ *Post* is positive and significant in both regressions, indicating that treated firms increase the number of women in these positions more so than control firms do. When we examine board positions, the result is marginally significant at the 10% level. We interpret these results as indicating a desire by treated firms to improve all aspects of their gender diversity score, including the promotion of women to senior management and board positions in order to improve their odds of being included in the WIN. Table 4.6 shows that the increase in the number of women in the workforce is driven by supervisory and higher ranked workers. Overall, the results in Table 4.6 show that the WIN effects are not driven by augmenting women participation in non-regular workforce; rather, the treatment firms appear to make meaningful changes to the representation of women in regular, permanent jobs at senior management levels. Treatment firms are not seen to make temporary adjustments to game WIN membership. We view the increase in women representation at senior levels as a deliberate attempt by treated firms to gain access to the WIN. The permanence of the senior positions gained by women also suggests that treated firms are genuinely striving for a change in culture in tune with their desire to improve their gender diversity score. In the next section, we gauge one metric of a cultural change by examining parental leave taken by men at treated vs. control firms. A more favorable culture for gender diversity would be consistent with an increase in the uptake of parental leaves by men employees.

4.4.4 Did the WIN Influence Attitudes Towards Parental Leave?

We next turn to the question of workplace disruptions due to childbirth and examine whether treated firms experience a change in the take-up rate for maternity or paternity leaves. Our priors are that treated firms are more likely to encourage such parental leaves under the assumption that they actively seek a change in culture in line with their desire to improve their GDS and gain access to the WIN. We find mixed results and present them in Table 4.7. The first three columns use maternity leave as the dependent variable. The variable of interest, *Treated* $\times$ *Post*, is not significant in any of the three regressions. In contrast, in columns 4 through 6, we find that men employees are significantly more likely to go on parental leave in treated firms in the post-WIN period compared to control firms. We submit that the higher uptake of parental leaves among men employees in treated firms marks a shift in corporate culture that is an outcome of the desire by treated firms to be included in the WIN. We contend that this is the first documented instance of a change in index design leading to a deliberate change in corporate culture. In the next section we investigate if institutional investors are influenced by the change in corporate culture at treated firms and respond by increasing their investment in treated firms relative to control firms. A preferential treatment by institutional investors would provide a clear incentive for treated firms to improve their gender diversity performance in order to gain membership on the WIN.

4.4.5 Changes in Institutional Ownership, Profitability, and Firm Value

Table 4.8 shows that institutional ownership of treated firms increased in the post period, confirming one of our discussed channels why firms may care. However, there is no significant change in the treated firms' market value (based on Tobin's q), and there is no significant impact on annual returns. Table 4.8 also shows that adding more women to the workforce did not come at the expense of profitability——EBITDA is unchanged for treated firms. Instead, our findings indicate that institutional investors display an incremental demand for shares in treated firms, consistent with their preference for investing in WIN firms. Moreover, the improvement in gender diversity scores at treated firm is not coming at the expense of shareholder value. Indeed, investors appear to increase their ownership in treated firms.

4.5 Conclusions

We use a unique experiment in Japan to examine whether equity market indexation can bring about real changes to corporate social performance. We examine the years surrounding the launch of the MSCI Empowering Women Index (WIN) that includes Japanese firms based on their workforce gender diversity performance. Using a difference-in-differences framework, we find that firms ranked around the index inclusion threshold improve their fraction of women in the workforce after the launch of the WIN compared to control firms. These results are not driven by firms' reducing the number of men employees.

We further document that the improvement of the fraction of women in the workforce happens in management and executive positions, while general worker positions do not show a significant change. We also show a potential shift in corporate culture in that men employees are more likely to take up parent leaves following the launch of the WIN. Finally, firms included in the WIN experience an increase in institutional ownership, but do not exhibit lower operating performance, which is often used as an argument against investments in firms' social performance.

Our results highlight an important capital markets channel through which changes in corporate social behaviour can be achieved. This finding is potentially important for regulators and sustainability-minded investors who are interested in improving the social performance of companies.

Figures

Figure 4.1: Women in the Workforce for Treated and Control Firms

(a) **Panel A: Fraction of Women**

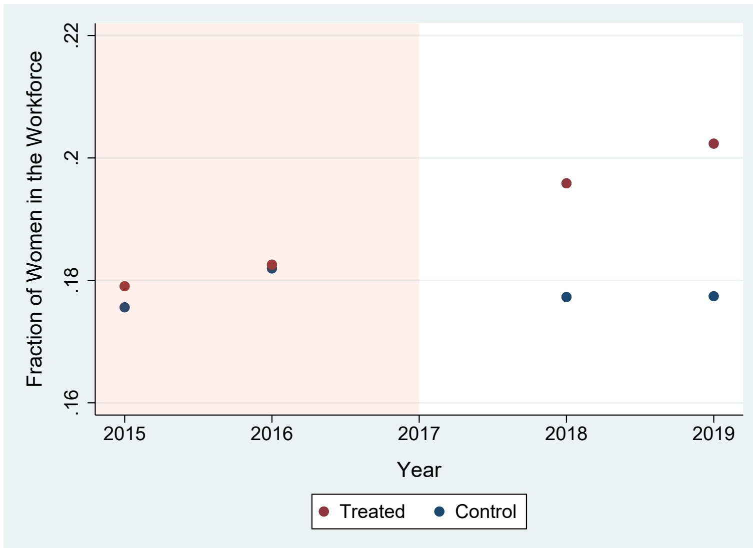

(b) **Panel B: Ratio of Women to Men**

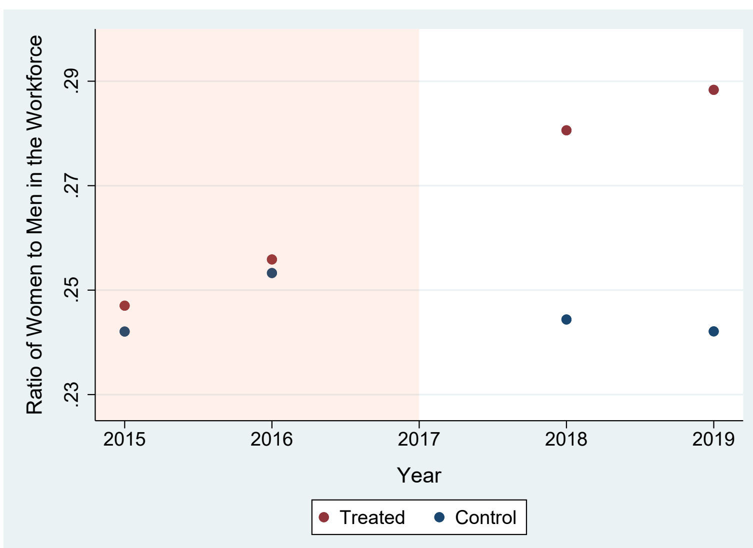

This figure shows the fraction of women in the workforce and the ratio of women to men in the workforce for the two years surrounding the inception of the MSCI Japan Empowering Women Index (WIN). The figure plots averages for treated and control firms. We follow MSCI's WIN creation methodology, for each year after the inception of the WIN, we rank all firms within their respective GICS sector by their MSCI Gender Diversity Score. Treated firms are those that rank around the WIN inclusion cutoff (the median of the Gender Diversity Score in each GICS sector) and fall within the 40th and 60th percentile of the distribution of the Gender Diversity Score. Control firms are those that rank below the 40th and above (equal) the 10th percentile of the distribution of the Gender Diversity Score. In the years before the inception of the WIN (2014-2016), a firm is treated if it was treated in any of the years after the inception of the WIN, and a firm is a control firm if it was a control firm in any of the years after the inception of the WIN. The WIN inception year (2017) is excluded from the analysis. All variables are defined in Table 4.1.

Tables

Table 4.1: Variable Descriptions

Variables	Description
Panel A: Employment Characteristics	
Fraction of Women in the Workforce	The number of women employees divided by the number of total men employees. *Source*: Toyo Keizai.
Ratio of Women to Men	The number of women employees divided by the number of men employees. *Source*: Toyo Keizai.
Change in Women Employees	Change in the number of women employees divided by the lagged number of women employees. *Source*: Toyo Keizai.
Change in Men Employees	Change in the number of men employees divided by the lagged number of men employees. *Source*: Toyo Keizai.
Turnover in Women Employees	The number of women employees leaving the company for reasons other than mandatory retirement divided by the lagged number of women employees. *Source*: Toyo Keizai.
Turnover in Men Employees	The number of men employees leaving the company for reasons other than mandatory retirement divided by the lagged number of men employees. *Source*: Toyo Keizai.
Fraction of Women in Workers	The number of women general workers in non-supervisory and non-management positions divided by the number of all workers. *Source*: Toyo Keizai.
Fraction of Women in General Management	The number of women with a general management position divided by the number of all general managers (*kanri-shoku*). *Source*: Toyo Keizai.
Fraction of Women in Executives	The number of women executive officers divided by the number of all executives (*shikko-yakuin*). *Source*: Toyo Keizai.
Fraction of Women on the Board	The number of women directors on the board of directors/audit & supervisory board divided by the number of all members of the board of directors/audit & supervisory board (*torishimari yaku* & *kansa yaku*). *Source*: Toyo Keizai.
Women Taking Maternity Leaves (All)	Number of women employees taking maternity leaves divided by the lagged number of all employees. *Source*: Toyo Keizai.
Women Taking Maternity Leaves (Under 30)	Number of women employees taking maternity leaves divided by the lagged number of all employees aged under 30. *Source*: Toyo Keizai.
Women Taking Maternity Leaves (Under 40)	Number of women employees taking maternity leaves divided by the lagged number of all employees aged under 40. *Source*: Toyo Keizai.
Men Taking Paternity Leaves (All)	Number of men employees taking paternity leaves divided by the lagged number of all employees. *Source*: Toyo Keizai.
Men Taking Paternity Leaves (Under 30)	Number of men employees taking paternity leaves divided by the lagged number of all employees aged under 30. *Source*: Toyo Keizai.
Men Taking Paternity Leaves (Under 40)	Number of men employees taking paternity leaves divided by the lagged number of all employees aged under 40. *Source*: Toyo Keizai.
Panel B: Firm Characteristics	
Total Assets	Total assets. *Source*: Worldscope.
Log (Total Assets)	Natural logarithm of total assets. *Source*: Worldscope.
Cash	Cash and short-term investments divided by total assets. *Source*: Worldscope.
Tangibility	Net property, plant, and equipment divided by total assets. *Source*: Worldscope.
Leverage	Total debt divided by total assets. *Source*: Worldscope.
ROA	Return on assets. *Source*: Worldscope.
Tobin's q	Calculated as (market value of equity plus book value of debt) divided by total book assets. *Source*: Worldscope.
Institutional Ownership	Fraction of total shares held by institutional owners. *Source*: Bloomberg.
Annual Stock Return	Fiscal year stock return measured as the total return index (RI) at the end of the fiscal year divided by the total return index at the end of the previous fiscal year minus one. *Source*: Datastream.
EBITDA	Earnings before interest, taxes, and depreciation divided by total assets. *Source*: Worldscope.

This table shows variable descriptions and lists the data sources.

Table 4.2: Summary Statistics

Variable	Mean	Median	Min	Max	SD
Panel A: Employment Characteristics					
Fraction of Women in the Workforce	0.180	0.140	0.055	0.686	0.112
Ratio of Women to Men	0.250	0.162	0.059	2.187	0.238
Change in Women Employees	0.006	0.003	-0.081	0.188	0.021
Change in Men Employees	0.005	0.003	-0.194	0.318	0.049
Turnover in Women Employees	0.010	0.004	0.000	0.292	0.029
Turnover in Men Employees	0.027	0.015	0.002	0.539	0.063
Fraction of Women in Workers	0.232	0.179	0.066	0.823	0.160
Fraction of Women in General Management	0.044	0.026	0.000	0.331	0.052
Fraction of Women in Executives	0.009	0.000	0.000	0.222	0.027
Fraction of Women on the Board	0.040	0.000	0.000	0.333	0.063
Women Taking Maternity Leaves (All)	0.013	0.007	0.001	0.330	0.033
Women Taking Maternity Leaves (Under 30)	0.051	0.036	0.005	0.348	0.047
Women Taking Maternity Leaves (Under 40)	0.021	0.015	0.002	0.230	0.023
Men Taking Paternity Leaves (All)	0.001	0.000	0.000	0.020	0.003
Men Taking Paternity Leaves (Under 30)	0.006	0.002	0.000	0.109	0.015
Men Taking Paternity Leaves (Under 40)	0.002	0.001	0.000	0.040	0.005
Panel B: Firm Characteristics					
Total Assets	2,871	625	6	204,000	13,337
Log (Total Assets)	6.594	6.438	1.792	12.226	1.336
Cash	0.163	0.137	0.014	0.706	0.122
Tangibility	0.284	0.273	0.003	0.851	0.175
Leverage	0.202	0.165	0.000	0.705	0.174
ROA	0.042	0.040	-0.096	0.201	0.033
Tobin's q	0.911	0.781	0.161	2.733	0.514
Institutional Ownership	0.480	0.475	0.001	0.996	0.169
Annual Stock Return	0.144	0.110	-0.534	1.946	0.288
EBITDA	0.097	0.096	-0.051	0.285	0.048

This table shows summary statistics. The variables are described in Table 4.1. The sample covers the 2013-2020 period and consists of 723 firms and 5,613 firm-year observations.

Table 4.3: Predicting WIN Membership with MSCI Data

	(1)	(2)	(3)	(4)
Ranked WIN	0.925***	0.936***		
	(88.72)	(106.78)		
Quartile 2			0.016**	0.025***
			(2.87)	(4.12)
Quartile 3			0.933***	0.944***
			(76.19)	(72.19)
Quartile 4			0.933***	0.954***
			(70.58)	(101.90)
Log (Total Assets)		-0.023***		-0.025***
		(-3.53)		(-3.84)
Cash		-0.042		-0.038
		(-0.90)		(-0.81)
Tangibility		-0.002		-0.003
		(-0.03)		(-0.05)
Leverage		-0.004		-0.005
		(-0.09)		(-0.11)
ROA		0.154		0.164
		(1.76)		(1.75)
Tobin's q		0.002		0.001
		(0.20)		(0.07)
Industry FE	Yes	Yes	Yes	Yes
N	1,718	1,718	1,718	1,718
Adjusted R^2	0.865	0.868	0.865	0.868

This table shows regression estimates of MSCI Japan Empowering Women (WIN) index membership, a dummy variable equal to one if a firm is a WIN member, and zero otherwise, on our own measures of whether a firm is a member of the WIN and control variables. We follow the MSCI's WIN creation methodology, and for each year after the inception of the WIN, we rank all firms within their respective GICS sector by their MSCI Gender Diversity Score. Ranked WIN is an indicator variable that is equal to one if a firm has an above-median Gender Diversity Score compared to firms in the same GICS sector and year, and zero otherwise. Quartiles 2, 3, and 4, are indicator variables equal to one if a firm falls within the 2, 3, and 4, quartile, respectively, of the Gender Diversity Score in a given GICS sector and year. The WIN inception year (2017) is excluded from the analysis. All other variables are defined in Table 4.1. The sample period is 2018 to 2020. The WIN inception year (2017) is excluded from the analysis. All variables are winsorized at the 1st and 99th percentiles. All right-hand side variables are lagged by one year. Standard errors are double clustered by firm and time and t-statistics are reported in parentheses. ***, **, * denote statistical significance at the 1%, 5%, and 10% level, respectively.

Table 4.4: Treatment Effects on Women in the Workforce

	Fraction of Women in the Workforce		Ratio of Women to Men in the Workforce		Effects in Event Time: Fraction of Women in the Workforce	Effects in Event Time: Ratio of Women to Men in the Workforce
	(1)	(2)	(3)	(4)	(5)	(6)
Treated × Post	0.007**	0.006**	0.020**	0.020**		
	(2.69)	(2.50)	(2.52)	(2.25)		
Treated	0.001	0.001	-0.002	-0.001		
	(0.37)	(0.49)	(-0.25)	(-0.11)		
Log (Total Assets)		0.001		-0.003	0.001	-0.004
		(0.10)		(-0.18)	(0.12)	(-0.19)
Cash		0.034		0.180	0.034	0.181
		(0.68)		(1.39)	(0.69)	(1.40)
Tangibility		0.027		0.135	0.027	0.140
		(0.77)		(1.59)	(0.77)	(1.59)
Leverage		0.015		0.057*	0.017	0.066**
		(1.33)		(2.03)	(1.47)	(2.33)
ROA		0.029		0.097	0.025	0.089
		(0.95)		(0.79)	(0.78)	(0.71)
Tobin's q		-0.009		-0.047	-0.008	-0.046
		(-0.93)		(-0.99)	(-0.90)	(-0.98)
Event × Year dummy						
Treated × D_{2014}					-0.000	-0.003
					(-0.06)	(-0.47)
Treated × D_{2015}					0.004	0.017
					(1.57)	(1.14)
Treated × D_{2016}					0.003	0.016
					(1.45)	(1.46)
Treated × D_{2018}					0.007**	0.023**
					(2.84)	(2.47)
Treated × D_{2019}					0.008**	0.022**
					(2.47)	(2.32)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
N	1,335	1,324	1,335	1,324	1,335	1,324
Adjusted R^2	0.973	0.973	0.932	0.934	0.973	0.934

This table shows regression estimates of the fraction of women in the workforce and the ratio of women to men in the workforce on an interaction term between Treated and Post and control variables. We follow MSCI's WIN creation methodology, for each year after the inception of the WIN, we rank all firms within their respective GICS sector by their MSCI Gender Diversity Score. Treated firms are those that rank around the WIN inclusion cutoff (the median of the Gender Diversity Score in each GICS sector) and fall within the 40th and 60th percentile of the distribution of the Gender Diversity Score. Control firms are those that rank below the 40th and above (equal) the 10th percentile of the distribution of the Gender Diversity Score. In the years before the inception of the WIN launch (2013-2016), a firm is treated if it was treated in any of the years after the inception of the WIN, and a firm is a control firm if it was a control firm in any of the years after the inception of the WIN. Post is a dummy variable equal to one for years 2018 and 2019, and zero otherwise. The WIN inception year (2017) is excluded from the analysis. Columns 4 and 5 report regression results in event time. For each firm, a set of time indicators variables is created for the four years before 2017 (the inception year of the WIN) and the two years after 2017. The indicator variable for the year 2013 (4 years before the inception of the WIN) is omitted from the regressions because of collinearity. Once a firm is treated after the WIN inception, all subsequent time indicators variables are set to one. All other variables are defined in Table 4.1. The sample period is 2013 to 2020. All variables are winsorized at the 1st and 99th percentiles. All right-hand side variables are lagged by one year. Standard errors are double clustered by firm and time and t-statistics are reported in parentheses. ***, **, * denote statistical significance at the 1%, 5%, and 10% level, respectively.

Table 4.5: Changes and Turnover in Women and Men Employees

	Change in Women Employees (1)	Change in Men Employees (2)	Turnover in Women Employees (3)	Turnover in Men Employees (4)
Treated × Post	0.007**	0.010	0.002	0.001
	(2.13)	(1.17)	(0.94)	(0.32)
Treated	-0.004	-0.010	-0.001	-0.003
	(-1.49)	(-1.31)	(-0.61)	(-0.73)
Log (Total Assets)	-0.020*	-0.043	0.003	0.010
	(-1.77)	(-1.25)	(0.66)	(1.14)
Cash	0.057**	0.081	-0.005	0.037
	(2.16)	(1.65)	(-0.44)	(1.01)
Tangibility	0.020	0.004	0.004	0.064*
	(0.96)	(0.06)	(0.23)	(1.77)
Leverage	0.007	-0.040	0.005	0.012
	(0.58)	(-1.25)	(0.46)	(0.55)
ROA	0.098*	0.159	-0.006	-0.117***
	(1.92)	(1.39)	(-0.58)	(-3.38)
Tobin's q	-0.002	-0.004	-0.003*	-0.003
	(-0.31)	(-0.53)	(-1.94)	(-1.01)
Firm FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
N	1,227	1,227	1,175	1,175
Adjusted R^2	0.110	0.144	0.809	0.803

This table shows regression estimates of measures of changes of women and men in the workforce and turnover of women and men employees on an interaction term between Treated and Post and control variables. We follow MSCI's WIN creation methodology, for each year after the inception of the WIN, we rank all firms within their respective GICS sector by their MSCI Gender Diversity Score. Treated firms are those that rank around the WIN inclusion cutoff (the median of the Gender Diversity Score in each GICS sector) and fall within the 40th and 60th percentile of the distribution of the Gender Diversity Score. Control firms are those that rank below the 40th and above (equal) the 10th percentile of the distribution of the Gender Diversity Score. In the years before the inception of the WIN launch (2013-2016), a firm is treated if it was treated in any of the years after the inception of the WIN, and a firm is a control firm if it was a control firm in any of the years after the inception of the WIN. Post is a dummy variable equal to one for years 2018 and 2019, and zero otherwise. The WIN inception year (2017) is excluded from the analysis. All other variables are defined in Table 4.1. The sample period is 2013 to 2020. The WIN inception year (2017) is excluded from the analysis. All variables are winsorized at the 1st and 99th percentiles. All right-hand side variables are lagged by one year. Standard errors are double clustered by firm and time and t-statistics are reported in parentheses. ***, **, * denote statistical significance at the 1%, 5%, and 10% level, respectively.

Table 4.6: Women in the Workforce by Position

	Fraction of Women by Position in the Workforce			
	Workers (1)	General Management (2)	Executives (3)	Board (4)
Treated × Post	-0.000	0.008**	0.013***	0.015*
	(-0.02)	(2.11)	(4.27)	(2.06)
Treated	0.005*	-0.003	-0.003	-0.004
	(1.96)	(-1.06)	(-0.81)	(-0.49)
Log (Total Assets)	0.013	0.001	-0.004	0.005
	(1.08)	(0.40)	(-0.45)	(0.30)
Cash	-0.017	0.057**	-0.018	-0.064
	(-0.47)	(2.69)	(-0.79)	(-1.66)
Tangibility	0.015	0.035	0.003	0.080
	(0.42)	(1.32)	(0.11)	(1.67)
Leverage	-0.046*	0.032**	0.015	-0.071**
	(-1.95)	(2.14)	(0.74)	(-2.15)
ROA	0.008	0.003	-0.020	-0.156**
	(0.14)	(0.11)	(-0.59)	(-2.26)
Tobin's q	-0.011	-0.006	0.001	-0.004
	(-1.07)	(-0.89)	(0.41)	(-0.51)
Firm FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
N	1,224	1,302	1,164	1,202
Adjusted R^2	0.956	0.912	0.567	0.567

This table shows regression estimates of the fraction of women by position in the workforce on an interaction term between Treated and Post and control variables. We follow MSCI's WIN creation methodology, for each year after the inception of the WIN, we rank all firms within their respective GICS sector by their MSCI Gender Diversity Score. Treated firms are those that rank around the WIN inclusion cutoff (the median of the Gender Diversity Score in each GICS sector) and fall within the 40th and 60th percentile of the distribution of the Gender Diversity Score. Control firms are those that rank below the 40th and above (equal) the 10th percentile of the distribution of the Gender Diversity Score. In the years before the inception of the WIN launch (2013-2016), a firm is treated if it was treated in any of the years after the inception of the WIN, and a firm is a control firm if it was a control firm in any of the years after the inception of the WIN. Post is a dummy variable equal to one for years 2018 and 2019, and zero otherwise. The WIN inception year (2017) is excluded from the analysis. All other variables are defined in Table 4.1. The sample period is 2013 to 2020. All variables are winsorized at the 1st and 99th percentiles. All right-hand side variables are lagged by one year. Standard errors are double clustered by firm and time and t-statistics are reported in parentheses. ***, **, * denote statistical significance at the 1%, 5%, and 10% level, respectively.

Table 4.7: Real Effects: Maternity and Parental Leaves

	Women Taking Maternity Leaves			Men Taking Paternity Leaves		
	All	Under 30	Under 40	All	Under 30	Under 40
	(1)	(2)	(3)	(4)	(5)	(6)
Treated × Post	0.002	-0.004	0.004	0.001**	0.007***	0.002***
	(0.91)	(-1.27)	(1.23)	(2.56)	(3.85)	(3.35)
Treated	-0.002	-0.003	-0.005	-0.001*	-0.004*	-0.001
	(-0.96)	(-0.86)	(-1.67)	(-1.79)	(-2.00)	(-1.37)
Log (Total Assets)	-0.003	-0.014	-0.010	-0.000	-0.005	-0.002
	(-0.98)	(-1.30)	(-1.52)	(-0.63)	(-1.37)	(-1.43)
Cash	-0.054	0.013	0.005	-0.000	0.004	0.000
	(-1.08)	(0.42)	(0.57)	(-0.20)	(0.27)	(0.01)
Tangibility	-0.025	-0.021	-0.002	0.002	-0.002	0.003
	(-1.12)	(-0.68)	(-0.12)	(0.52)	(-0.12)	(0.42)
Leverage	0.013*	0.049	0.023**	0.002	0.025*	0.007*
	(1.84)	(1.65)	(2.27)	(1.57)	(1.99)	(1.87)
ROA	0.009	0.131	0.046	0.004	0.046	0.010
	(0.79)	(1.12)	(1.09)	(0.95)	(1.39)	(1.15)
Tobin's q	0.000	0.009*	-0.001	0.001*	0.003	0.001
	(0.21)	(1.97)	(-0.26)	(1.85)	(1.40)	(1.59)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
N	1,202	1,081	1,098	1,130	1,023	1,031
Adjusted R^2	0.930	0.808	0.756	0.611	0.575	0.586

This table shows regression estimates of the fraction of women who take maternity leaves and the fraction of men who take paternity leaves on an interaction term between Treated and Post and control variables. We follow MSCI's WIN creation methodology, for each year after the inception of the WIN, we rank all firms within their respective GICS sector by their MSCI Gender Diversity Score. Treated firms are those that rank around the WIN inclusion cutoff (the median of the Gender Diversity Score in each GICS sector) and fall within the 40th and 60th percentile of the distribution of the Gender Diversity Score. Control firms are those that rank below the 40th and above (equal) the 10th percentile of the distribution of the Gender Diversity Score. In the years before the inception of the WIN launch (2013-2016), a firm is treated if it was treated in any of the years after the inception of the WIN, and a firm is a control firm if it was a control firm in any of the years after the inception of the WIN. Post is a dummy variable equal to one for years 2018 and 2019, and zero otherwise. The WIN inception year (2017) is excluded from the analysis. All other variables are defined in Table 4.1. The sample period is 2013 to 2020. All variables are winsorized at the 1st and 99th percentiles. All right-hand side variables are lagged by one year. Standard errors are double clustered by firm and time and t-statistics are reported in parentheses. ***, **, * denote statistical significance at the 1%, 5%, and 10% level, respectively.

Table 4.8: Real Effects: Institutional Ownership and Performance

	Institutional Ownership (1)	Tobin's q (2)	Annual Stock Return (3)	EBITDA (4)
WIN × Post	2.908**	-0.040	-0.013	-0.003
	(2.67)	(-0.76)	(-0.63)	(-1.82)
Log (Total Assets)	2.688*	-0.032	-0.080**	-0.013**
	(2.19)	(-0.27)	(-3.19)	(-3.06)
Cash	-0.167	-0.771	-0.157	-0.004
	(-0.05)	(-1.13)	(-1.56)	(-0.30)
Tangibility	-20.694**	0.195	-0.327**	-0.009
	(-3.38)	(0.41)	(-2.55)	(-0.82)
Leverage	-7.759	0.129	0.092	-0.031
	(-1.86)	(0.50)	(0.92)	(-1.61)
ROA	-11.449	2.086	-1.237*	
	(-0.88)	(0.99)	(-2.30)	
Tobin's q	1.866**			0.037***
	(2.95)			(10.87)
Firm FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
N	4,106	4,172	4,157	4,009
Adjusted R^2	0.802	0.850	0.185	0.823

This table shows regression estimates of measures of institutional ownership and performance on an interaction term between WIN and Post and control variables. WIN is a dummy variable equal to one if the firm was included in the WIN in any year, and zero otherwise. Post is a dummy variable equal to one for years 2018 and 2019, and zero otherwise. The WIN inception year (2017) is excluded from the analysis. All other variables are defined in Table 4.1. Columns 2 and 3 exclude lagged Tobin's q and column 4 excludes ROA as explanatory variables. The sample period is 2013 to 2020. All variables are winsorized at the 1st and 99th percentiles. All right-hand side variables are lagged by one year. Standard errors are double clustered by firm and time and t-statistics are reported in parentheses. ***, **, * denote statistical significance at the 1%, 5%, and 10% level, respectively.

References

Abidi, N., Miquel-Flores, I., 2018. Who benefits from the corporate QE? A regression discontinuity design approach. ECB Working Paper, No. 2145.

Acharya, V. V., Steffen, S., 2020. The risk of being a fallen angel and the corporate dash for cash in the midst of COVID. The Review of Corporate Finance Studies 9, 430–471.

Adams, R. B., Ferreira, D., 2009. Women in the boardroom and their impact on governance and performance. Journal of Financial Economics 94, 291–309.

Adams, R. B., Funk, P., 2012. Beyond the glass ceiling: Does gender matter. Management Science 58, 219–235.

Adhikari, B., Agrawal, A., Malm, J., 2019. Do women managers keep firms out of trouble? evidence from corporate litigation and policies. Journal of Accounting and Economics 67, 202–225.

Ahern, K. R., Dittmar, A., 2012. The changing of the boards: The impact on firm valuation of mandated female board representation. The Quarterly Journal of Economics 127, 137–197.

Akepanidtaworn, K., Di Mascio, R., Imas, A., Schmidt, L., 2021. Selling fast and buying slow: Heuristics and trading performance of institutional investors. NBER Working Paper, No. 29076.

Albert, A., Anderson, J. A., 1984. On the existence of maximum likelihood estimates in logistic regression models. Biometrika 71, 1–10.

Allen, J., Kastl, J., Wittwer, M., 2020. Primary dealers and the demand for government debt. Bank of Canada Staff Working Paper, No. 2020-29.

Almeida, H., Campello, M., Laranjeira, B., Weisbenner, S., 2011. Corporate debt maturity and the real effects of the 2007 credit crisis. Critical Finance Review 1, 3–58.

An, Y., Song, Z., forthcoming. Does the Federal Reserve Obtain Competitive and Appropriate Prices in Monetary Policy Implementation? The Review of Financial Studies, forthcoming .

Ap Gwilym, O., Clare, A., Thomas, S., 1998a. Extreme price clustering in the London equity index futures and options markets. Journal of Banking & Finance 22, 1193–1206.

Ap Gwilym, O., Clare, A., Thomas, S., 1998b. Price clustering and bid-ask spreads in international bond futures. Journal of International Financial Markets, Institutions and Money 8, 377–391.

Arce, Ó., Mayordomo, S., Gimeno, R., 2021. Making room for the needy: the credit-reallocation effects of the ECB's corporate QE. Review of Finance 25, 43–84.

Arrow, K. J., 1987. The demand for information and the distribution of income. Probability in the Engineering and Informational Sciences 1, 3–13.

Avraham, D., Selvaggi, P., Vickery, J. I., 2012. A structural view of US bank holding companies. Economic Policy Review 18, 65–81.

Badoer, D. C., James, C. M., 2016. The determinants of long-term corporate debt issuances. The Journal of Finance 71, 457–492.

Baker, M., Greenwood, R., Wurgler, J., 2009. Catering through nominal share prices. The Journal of Finance 64, 2559–2590.

Ball, C. A., Torous, W. N., Tschoegl, A. E., 1985. The degree of price resolution: The case of the gold market. Journal of Futures Markets 5, 29–43.

Bank of Japan, 2009. Outright purchases of corporate financing instruments. Announcement on January 22, 2009.

Bank of Japan, 2020a. Enhancement of monetary easing. Announcement on April 27, 2020. Retrieved from https://www.boj.or.jp/en/announcements/release_2020/k200427a.pdf.

Bank of Japan, 2020b. Enhancement of monetary easing in light of the impact of the outbreak of the novel coronavirus (COVID-19). Announcement on May 16, 2020. Retrieved from https://www.boj.or.jp/en/announcements/release_2020/k200316b.pdf.

Bank of Japan, 2020c. Statement on monetary policy. Announcement on May 22, 2020. Retrieved from https://www.boj.or.jp/en/announcements/release_2020/k200522b.pdf.

Bansal, R., Wu, D., Yaron, A., 2021. Socially responsible investing in good and bad times. The Review of Financial Studies 35, 2067–2099.

Barbon, A., Gianinazzi, V., 2019. Quantitative easing and equity prices: Evidence from the ETF program of the Bank of Japan. The Review of Asset Pricing Studies 9, 210–255.

Barclay, M. J., Hendershott, T., Kotz, K., 2006. Automation versus intermediation: Evidence from Treasuries going off the run. The Journal of Finance 61, 2395–2414.

Bartus, T., 2005. Estimation of marginal effects using margeff. The Stata Journal 5, 309–329.

Bebchuk, L. A., Hirst, S., 2019. Index funds and the future of corporate governance: Theory, evidence, and policy. Columbia Law Review 119, 2029–2146.

Becker, B., Benmelech, E., 2021. The resilience of the US corporate bond market during financial crises. NBER Working Paper, No. 28868.

Berk, J. B., van Binsbergen, J. H., 2021. The impact of impact investing. Available at SSRN 3909166.

Bernanke, B., 2022. 21st Century Monetary Policy: The Federal Reserve from the Great Inflation to COVID-19. W. W. Norton & Company, New York, New York.

Bertrand, M., Black, S. E., Jensen, S., Lleras-Muney, A., 2019. Breaking the glass ceiling? the effect of board quotas on female labor market outcomes in norway. The Review of Economic Studies 86, 191–239.

Bessembinder, H., Spatt, C., Venkataraman, K., 2020. A survey of the microstructure of fixed-income markets. Journal of Financial and Quantitative Analysis 55, 1–45.

Bhargava, S., Loewenstein, G., 2015. Behavioral economics and public policy 102: Beyond nudging. American Economic Review 105, 396–401.

Bhattacharya, U., Holden, C. W., Jacobsen, S., 2012. Penny wise, dollar foolish: Buy–sell imbalances on and around round numbers. Management Science 58, 413–431.

Bhattacharya, U., Kuo, W.-Y., Lin, T.-C., Zhao, J., 2018. Do superstitious traders lose money? Management Science 64, 3772–3791.

Blanchard, O. J., Summers, L. H., 2020. Automatic stabilizers in a low-rate environment. AEA Papers and Proceedings 110, 125–30.

Boneva, L., Elliott, D., Kaminska, I., Linton, O., McLaren, N., Morley, B., 2022. The impact of corporate QE on liquidity: Evidence from the UK. The Economic Journal 132, 2615–2643.

Boneva, L., Kastl, J., Zikes, F., 2020. Dealer balance sheets and bidding behavior in the UK QE reverse auctions. Working paper. https://www.princeton.edu/~jkastl/draft_qe.pdf.

Boyarchenko, N., Cox, C., Crump, R. K., Danzig, A., Kovner, A., Shachar, O., Steiner, P., 2021. COVID response: The primary and secondary corporate credit facilities. Economic Policy Review.

Boyarchenko, N., Kovner, A., Shachar, O., 2022. It's what you say and what you buy: A holistic evaluation of the corporate credit facilities. Journal of Financial Economics 144, 695–731.

Brainard, L., 2018. The structure of the Treasury market: What are we learning? Speech at "The Evolving Structure of the US Treasury Market," Fourth Annual Conference (December 2015). Federal Reserve Bank of New York, New York City, Federal Reserve Bank of New York.

Breedon, F., 2018. On the transactions costs of UK quantitative easing. Journal of Banking & Finance 88, 347–356.

Brunnermeier, M., Farhi, E., Koijen, R. S. J., Krishnamurthy, A., Ludvigson, S. C., Lustig, H., Nagel, S., Piazzesi, M., 2021. Review Article: Perspectives on the Future of Asset Pricing. The Review of Financial Studies 34, 2126–2160.

Butler, A. W., Gao, X., Uzmanoglu, C., 2023. Maturity clienteles and corporate bond maturities. Journal of Financial and Quantitative Analysis 58, 1263–1294.

Caballero, R. J., Krishnamurthy, A., 2008. Collective risk management in a flight to quality episode. The Journal of Finance 63, 2195–2230.

Calomiris, C. W., Larrain, M., Schmukler, S. L., Williams, T., 2022. Large international corporate bonds: Investor behavior and firm responses. Journal of International Economics 137, 103624.

Camerer, C. F., Issacharoff, S., Loewenstein, G., O'Donoghue, T., Rabin, M., 2003. Regulation for conservatives: Behavioral economics and the case for 'asymmetric paternalism'. University of Pennsylvania Law Review 151, 1211–1254.

Charoenwong, B., Morck, R., Wiwattanakantang, Y., 2021. Bank of Japan equity purchases: The (non-)effects of extreme quantitative easing. Review of Finance 25, 713–743.

Chattopadhyaya, A., Shaffer, M. D., Wang, C. C., 2020. Governance through shame and aspiration: Index creation and corporate behavior. Journal of Financial Economics 135, 704–724.

Chen, H., Xu, Y., Yang, J., 2021. Systematic risk, debt maturity, and the term structure of credit spreads. Journal of Financial Economics 139, 770–799.

Chernenko, S., Faulkender, M., 2011. The two sides of derivatives usage: Hedging and speculating with interest rate swaps. Journal of Financial and Quantitative Analysis 46, 1727–1754.

Choi, J., Hackbarth, D., Zechner, J., 2018. Corporate debt maturity profiles. Journal of Financial Economics 130, 484–502.

Choi, J., Hackbarth, D., Zechner, J., 2021. Granularity of corporate debt. Journal of Financial and Quantitative Analysis 56, 1127–1162.

Christie, W. G., Harris, J. H., Schultz, P. H., 1994. Why did Nasdaq market makers stop avoiding odd-eighth quotes? The Journal of Finance 49, 1841–1860.

Christie, W. G., Schultz, P. H., 1994. Why do NASDAQ market makers avoid odd-eighth quotes? The Journal of Finance 49, 1813–1840.

Christie, W. G., Schultz, P. H., 1999. The initiation and withdrawal of odd-eighth quotes among Nasdaq stocks: An empirical analysis. Journal of Financial Economics 52, 409–442.

Crawford, M., 2021. Abe's womenomics policy, 2013-2020: Tokenism, gradualism, or failed strategy? Asia-Pacific Journal 19, 1–16.

Cronqvist, H., Yu, F., 2017. Shaped by their daughters: Executives, female socialization, and corporate social responsibility. Journal of Financial Economics 126, 543–562.

Dai, R., Liang, H., Ng, L., 2021. Socially responsible corporate customers. Journal of Financial Economics 142, 598–626.

D'Amico, S., King, T. B., 2013. Flow and stock effects of large-scale treasury purchases: Evidence on the importance of local supply. Journal of Financial Economics 108, 425–448.

Dathan, M., Davydenko, S. A., 2020. Debt issuance in the era of passive investment. Available at SSRN 3152612.

De Santis, R. A., Zaghini, A., 2021. Unconventional monetary policy and corporate bond issuance. European Economic Review 135, 103727.

Demirgüç-Kunt, A., Horváth, B. L., Huizinga, H., 2023. Corporate QE in Europe during the COVID-19 crisis and debt overhang. Journal of International Money and Finance 135, 102855.

Demsetz, H., 1997. Limit orders and the alleged Nasdaq collusion. Journal of Financial Economics 45, 91–95.

Diamond, D. W., 1991. Debt maturity structure and liquidity risk. The Quarterly Journal of Economics 106, 709–737.

Dimson, E., Karakaş, O., Li, X., 2015. Active ownership. The Review of Financial Studies 28, 3225–3268.

Duffie, D., 2020. Still the world's safe haven? Redesigning the US Treasury market after the COVID-19 crisis. *Brookings Institution* Hutchins Center Working Paper No. 62.

Duval, R., Hong, G. H., Timmer, Y., 2020. Financial frictions and the great productivity slowdown. The Review of Financial Studies 33, 475–503.

Dyck, A., Lins, K. V., Roth, L., Towner, M., Wagner, H. F., 2023. Renewable governance: Good for the environment? Journal of Accounting Research 61, 279–327.

Dyck, A., Lins, K. V., Roth, L., Wagner, H. F., 2019. Do institutional investors drive corporate social responsibility? international evidence. Journal of Financial Economics 131, 693–714.

Erel, I., Julio, B., Kim, W., Weisbach, M. S., 2012. Macroeconomic conditions and capital raising. The Review of Financial Studies 25, 341–376.

Faccio, M., Marchica, M.-T., Mura, R., 2016. CEO gender, corporate risk-taking, and the efficiency of capital allocation. Journal of Corporate Finance 39, 193–209.

Falato, A., Goldstein, I., Hortaçsu, A., 2021. Financial fragility in the COVID-19 crisis: The case of investment funds in corporate bond markets. Journal of Monetary Economics 123, 35–52.

Faulkender, M., 2005. Hedging or market timing? Selecting the interest rate exposure of corporate debt. The Journal of Finance 60, 931–962.

Firth, D., 1993. Bias reduction of maximum likelihood estimates. Biometrika 80, 27–38.

Flanagan, T., Purnanandam, A., 2020. Corporate bond purchases after COVID-19: Who did the fed buy and how did the markets respond? Available at SSRN: http://dx.doi.org/10.2139/ssrn.3668342.

Flannery, M. J., 1986. Asymmetric information and risky debt maturity choice. The Journal of Finance 41, 19–37.

Fleming, M., Nguyen, G., Ruela, F., forthcoming. Tick size, competition for liquidity provision, and price discovery: Evidence from the us treasury market. Management Science, forthcoming .

Fleming, M. J., Garbade, K. D., Keane, F., 2005. Anomalous bidding in short-term Treasury bill auctions. Journal of Financial Research 28, 165–176.

Franks, J. R., Mayer, C., Miyajima, H., 2014. The ownership of Japanese corporations in the 20th century. The Review of Financial Studies 27, 2580–2625.

Freund, S., Nguyen, N. H., Phan, H. V., 2021. Shareholder litigation and corporate social responsibility. Journal of Financial and Quantitative Analysis 56, 1161–1194.

Galema, R., Lugo, S., 2021. When central banks buy corporate bonds: Target selection and impact of the European Corporate Sector Purchase Program. Journal of Financial Stability 54, 100881.

Gennaioli, N., Shleifer, A., 2018. A Crisis of Beliefs: Investor Psychology and Financial Fragility. Princeton University Press, Princeton.

Giannone, D. A., Robotti, C., 2022. Are the primary dealers of the New York Fed really special? Available at SSRN: http://dx.doi.org/10.2139/ssrn.4010787.

Giglio, S., Kelly, B., Stroebel, J., 2021. Climate finance. Annual Review of Financial Economics 13, 15–36.

Gilchrist, S., Wei, B., Yue, V. Z., Zakrajšek, E., 2020. The Fed takes on corporate credit risk: An analysis of the efficacy of the SMCCF. NBER Working Paper, No. 27809.

Glosserman, B., 2019. Peak Japan: The End of Great Ambitions. Georgetown University Press, Washington, DC.

Goldberg, J., 2020. Liquidity supply by broker-dealers and real activity. Journal of Financial Economics 136, 806–827.

Goldreich, D., 2015. Bounded rationality of dealers in US Treasury auctions. Available at: http://www-2.rotman.utoronto.ca/david.goldreich/bounded.pdf.

GPIF, 2018. Annual report fiscal year 2017. https://www.gpif.go.jp/en/annual_report_fiscal_year_2017.pdf.

Graham, J. R., Harvey, C. R., 2001. The theory and practice of corporate finance: Evidence from the field. Journal of Financial Economics 60, 187–243.

Greenwood, R., Hanson, S., Stein, J. C., 2010. A gap-filling theory of corporate debt maturity choice. The Journal of Finance 65, 993–1028.

Griffin, J. M., Harris, J. H., Topaloglu, S., 2003. The dynamics of institutional and individual trading. The Journal of Finance 58, 2285–2320.

Griffin, J. M., Hirschey, N., Kruger, S., 2023. Do municipal bond dealers give their customers "fair and reasonable" pricing? The Journal of Finance 78, 887–934.

Grosse-Rueschkamp, B., Steffen, S., Streitz, D., 2019. A capital structure channel of monetary policy. Journal of Financial Economics 133, 357–378.

Grossman, S. J., Miller, M. H., Cone, K. R., Fischel, D. R., Ross, D. J., 1997. Clustering and competition in asset markets. The Journal of Law and Economics 40, 23–60.

Grossman, S. J., Stiglitz, J. E., 1980. On the impossibility of informationally efficient markets. American Economic Review 70, 393–408.

Group of Thirty, 2021. U.S. Treasury markets: Steps toward increased resilience. Group of Thirty Working Group on Treasury Market Liquidity. July 2021. Washington, D.C.

Guedes, J., Opler, T., 1996. The determinants of the maturity of corporate debt issues. The Journal of Finance 51, 1809–1833.

Gupta, A., 2022. The internal capital markets of global dealer banks. Journal of Financial Crises 4, 165–188.

Haddad, V., Moreira, A., Muir, T., 2021. When selling becomes viral: Disruptions in debt markets in the COVID-19 crisis and the Fed's response. The Review of Financial Studies 34, 5309–5351.

Haddad, V., Moreira, A., Muir, T., 2023. Whatever it takes? the impact of conditional policy promises. NBER Working Paper, No. 31259.

Halling, M., Yu, J., Zechner, J., 2020. How did COVID-19 affect firms' access to public capital markets? The Review of Corporate Finance Studies 9, 501–533.

Han, S. H., Pagano, M. S., Shin, Y. S., 2012. Rating agency reputation, the global financial crisis, and the cost of debt. Financial Management 41, 849–884.

Harkrader, J., Puglia, M., 2020. Price discovery in the US Treasury cash market: On principal trading firms and dealers. Finance and Economics Discussion Series (FEDS) Working Paper No. 2020-96.

Harris, L., 1991. Stock price clustering and discreteness. The Review of Financial Studies 4, 389–415.

Hattori, T., 2019. Do liquidity enhancement auctions improve the market liquidity in the JGB market? Economics Letters 183, 108516.

He, Z., Kelly, B., Manela, A., 2017. Intermediary asset pricing: New evidence from many asset classes. Journal of Financial Economics 126, 1–35.

He, Z., Nagel, S., Song, Z., 2022. Treasury inconvenience yields during the COVID-19 crisis. Journal of Financial Economics 143, 57–79.

Heinze, G., Schemper, M., 2002. A solution to the problem of separation in logistic regression. Statistics in Medicine 21, 2409–2419.

Helwege, J., Turner, C. M., 1999. The slope of the credit yield curve for speculative-grade issuers. The Journal of Finance 54, 1869–1884.

Hortaçsu, A., Kastl, J., 2012. Valuing dealers' informational advantage: A study of Canadian Treasury auctions. Econometrica 80, 2511–2542.

Hortaçsu, A., Kastl, J., Zhang, A., 2018. Bid shading and bidder surplus in the US Treasury auction system. American Economic Review 108, 147–69.

Houston, J. F., Venkataraman, S., 1994. Optimal maturity structure with multiple debt claims. Journal of Financial and Quantitative Analysis pp. 179–197.

Huang, J., Kisgen, D. J., 2013. Gender and corporate finance: Are male executives overconfident relative to female executives. Journal of Financial Economics 108, 822–839.

Ikenberry, D. L., Weston, J. P., 2008. Clustering in US stock prices after decimalisation. European Financial Management 14, 30–54.

Iliev, P., Roth, L., forthcoming. Director expertise and corporate sustainability. Review of Finance, forthcoming .

Ioannou, I., Serafeim, G., 2012. What drives corporate social performance? the role of nation-level institutions. Journal of International Business Studies 43, 834–864.

Joint Staff Report, 2015. Joint staff report: The U.S. Treasury market on October 15, 2014. U.S. Department of the Treasury, Board of Governors of the Federal Reserve System, Federal Reserve Bank of New York, U.S. Securities and Exchange Commission, U.S. Commodity Futures Trading Commission.

Kalemli-Özcan, Ş., Laeven, L., Moreno, D., 2022. Debt overhang, rollover risk, and corporate investment: Evidence from the European crisis. Journal of the European Economic Association 20, 2353–2395.

Kandel, E., Marx, L. M., 1997. Nasdaq market structure and spread patterns. Journal of Financial Economics 45, 61–89.

Kargar, M., Lester, B., Lindsay, D., Liu, S., Weill, P.-O., Zúñiga, D., 2021. Corporate bond liquidity during the COVID-19 crisis. The Review of Financial Studies 34, 5352–5401.

Kaul, A., Mehrotra, V., Morck, R., 2000. Demand curves for stocks do slope down: New evidence from an index weights adjustment. The Journal of Finance 55, 893–912.

Kawaguchi, M., 2015. GPIF no kokuren sekinin toushi gensoku (RPI) shomei no inpakuto. https://www.dir.co.jp/report/column/20151002_010176.html.

Kim, D., Starks, L. T., 2016. Gender diversity on corporate boards: Do women contribute unique skills? American Economic Review 106, 267–271.

Kindleberger, C. P., 1978. Manias, Panics and Crashes: A History of Financial Crises. Basic Books, New York.

Kosmidis, I., 2020. brglm2: bias reduction in generalized linear models. R package version 0.7.0.

Kosmidis, I., Firth, D., 2009. Bias reduction in exponential family nonlinear models. Biometrika 96, 793–804.

Kosmidis, I., Firth, D., 2011. Multinomial logit bias reduction via the Poisson log-linear model. Biometrika 98, 755–759.

Krueger, P., Sautner, Z., Starks, L., 2020. The importance of climate risks for institutional investors. The Review of Financial Studies 33, 1067–1111.

Kunze, A., Miller, A. R., 2017. Women helping women? evidence from private sector data on workplace hierarchies. The Review of Economics and Statistics 99, 769–775.

Kuo, W.-Y., Lin, T.-C., Zhao, J., 2015. Cognitive limitation and investment performance: Evidence from limit order clustering. The Review of Financial Studies 28, 838–875.

Kuroda, H., 2020. Novel coronavirus (COVID-19): Economic and financial developments and the responses taken by the Bank of Japan. Speech via Webcast Hosted by the Naigai Josei Chosa Kai (Research Institute of Japan) in Tokyo, May 14, 2020.

Kuttner, K. N., 2018. Outside the box: Unconventional monetary policy in the great recession and beyond. Journal of Economic Perspectives 32, 121–46.

Li, D., 2007. New findings in the market microstructure of over-the-counter markets. Ph.D. Dissertation. Carnegie Mellon University.

Liang, H., Renneboog, L., 2017. On the foundations of corporate social responsibility. Journal of Finance 72, 853–910.

Linh, N. T., 2021. The impact of the Bank of Japan's exchange traded fund and corporate bond purchases on firms' capital structure. RCESR Discussion Paper Series DP21-1, Research Center for Economic and Social Risks, Institute of Economic Research, Hitotsubashi University.

Lins, K. V., Roth, L., Servaes, H., Tamayo, A., 2022. Sexism, culture, and firm value: Evidence from the Harvey Weinstein scandal and the #MeToo movement. Available at SSRN 3458312.

Lo, A. W., 2019. Adaptive Markets: Financial Evolution at the Speed of Thought. Princeton University Press, Princeton.

Lugo, S., 2021. Short-term debt catering. Journal of Corporate Finance 66, 101817.

Lugo, S., Piccillo, G., 2019. The relation between corporate and government debt maturity in Europe. Journal of Financial and Quantitative Analysis 54, 2119–2140.

Ma, Y., Xiao, K., Zeng, Y., 2022. Mutual fund liquidity transformation and reverse flight to liquidity. The Review of Financial Studies 35, 4674–4711.

Mäkinen, T., Li, F., Mercatanti, A., Silvestrini, A., et al., 2020. Effects of eligibility for central bank purchases on corporate bond spreads. BIS Working Papers No 894.

Matos, P., 2020. Esg and responsible institutional investing around the world: A critical review. Research Foundation Literature Reviews CFA Institute.

Matsa, D. A., Miller, A. R., 2011. Chipping away at the glass ceiling: Gender spillovers in corporate leadership. American Economic Review 101, 635–639.

McPartland, K., 2018. CME's NEX acquisition a game changer – but how? Greenwich Blog, April 27, 2018. https://www.greenwich.com/blog/cmes-nex-acquisition-game-changer-how.

Mizrach, B., Neely, C. J., 2006. The transition to electronic communications networks in the secondary Treasury market. Federal Reserve Bank of St. Louis Review 88, 527–41.

Modigliani, F., Sutch, R., 1966. Innovations in interest rate policy. American Economic Review 56, 178–197.

Modigliani, F., Sutch, R., 1967. Debt management and the term structure of interest rates: An empirical analysis of recent experience. Journal of Political Economy 75, 569–589.

Mondria, J., Vives, X., Yang, L., 2022. Costly interpretation of asset prices. Management Science 68, 52–74.

MSCI, 2019a. MSCI Japan empowering women (WIN) index methodology. https://www.msci.com/documents/10199/c0ae70bb-8688-46b7-93b9-c8a607ae028a.

MSCI, 2019b. MSCI workforce gender diversity data methodology. https://www.msci.com/documents/10199/0d2a428b-ee0d-4d63-9bfb-3db4b003cc10.

Murfin, J., Petersen, M., 2016. Loans on sale: Credit market seasonality, borrower need, and lender rents. Journal of Financial Economics 121, 300–326.

Myers, S. C., 1977. Determinants of corporate borrowing. Journal of Financial Economics 5, 147–175.

Nikiforov, A., Pilotte, E., 2017a. Macroeconomic announcements and the distribution of price-endings in the US Treasury market. Financial Review 52, 69–100.

Nikiforov, A., Pilotte, E., 2017b. Quote activity and round quotes in a voice-brokered bond market. Applied Economics Letters 24, 214–218.

Nikiforov, A., Pilotte, E., 2019. Security design, market risk and round quotes in the treasury bond market. Applied Economics Letters 26, 971–977.

Nikkei, 2021. *Shasai "nichigin tredo" ihen taishogai meigara uwanose kinri ni myoumi* (Change in corporate bond "BOJ trade": Attractiveness of higher yields of non-targeted bonds) (in Japanese). March 26, 2021. Retrieved from https://www.nikkei.com/article/DGXZQOGD25C1L0V20C21A3000000/?unlock=1.

Nofsinger, J. R., Sias, R. W., 1999. Herding and feedback trading by institutional and individual investors. The Journal of Finance 54, 2263–2295.

Nozawa, Y., Qiu, Y., 2021. Corporate bond market reactions to quantitative easing during the COVID-19 pandemic. Journal of Banking & Finance 133, 106153.

Ogawa, E., 2021. The case of Japan. In: Andreosso-O'Callaghan, B., Moon, W., Sohn, W. (eds.), *Economic Policy and the Covid-19 Crisis: The Macroeconomic Response in the US, Europe and East Asia*, Routledge, pp. 165–187.

O'Hara, M., Zhou, X., 2021. Anatomy of a liquidity crisis: Corporate bonds in the COVID-19 crisis. Journal of Financial Economics 142, 46–68.

Packer, F., 2002. Credit ratings and the Japanese corporate bond market. In: Levich, R. M., Majnoni, G., Reinhart, C. M. (eds.), *Ratings, Rating Agencies and the Global Financial System*, Springer, pp. 139–158.

Pastor, L., Stambaugh, R. F., Taylor, L. A., 2021. Sustainable investing in equilibrium. Journal of Financial Economics 142, 550–571.

Pegoraro, S., Montagna, M., 2022. Demand-driven bond financing in the Euro area. Available at SSRN: https://ssrn.com/abstract=3791730.

Peress, J., 2004. Wealth, information acquisition, and portfolio choice. The Review of Financial Studies 17, 879–914.

Pettenuzzo, D., Sabbatucci, R., Timmermann, A., 2021. Outlasting the pandemic: Corporate payout and financing decisions during Covid-19. CEPR Discussion Paper No. DP16145.

Potter, S., 2013. The implementation of current asset purchases. In: Remarks at the Annual Meeting with Primary Dealers. Federal Reserve Bank of New York, New York City.

Potter, S., 2015. The Federal Reserve's counterparty framework: past, present, and future. In: Remarks at the 2015 Roundtable on Treasury Markets and Debt Management. Federal Reserve Bank of New York, New York City.

Rennison, J., McLannahan, B., 2016. Wells Fargo approved as 'primary dealer' for US Treasury debt. *Financial Times*. April 18, 2016. Retrieved from https://www.ft.com/content/130e20d4-05b6-11e6-a70d-4e39ac32c284.

Sack, B. P., 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, Federal Reserve Bank of New York.

Schiller, C., 2018. Global supply-chain networks and corporate social responsibility. Available at SSRN 3089311.

Schlepper, K., Hofer, H., Riordan, R., Schrimpf, A., 2020. The market microstructure of central bank bond purchases. Journal of Financial and Quantitative Analysis 55, 193–221.

Schrimpf, A., Shin, H. S., Sushko, V., et al., 2020. Leverage and margin spirals in fixed income markets during the Covid-19 crisis. BIS Bulletin 2.

SEC, 2020. U.S. credit markets: Interconnectedness and the effects of the COVID-19 economic shock. A staff report of the Division of Economic and Risk Analysis of the U.S. Securities and Exchange Commission, October 2020.

Shiller, R. J., 2000. Irrational Exuberance. Princeton University Press, Princeton.

Shleifer, A., 1986. Do demand curves for stocks slope down? Journal of Financial Economics 16, 253–284.

Song, Z., Zhu, H., 2018. Quantitative easing auctions of Treasury bonds. Journal of Financial Economics 128, 103–124.

Suganuma, K., Ueno, Y., 2018. The effects of the Bank of Japan's corporate and government bond purchases on credit spreads. Institute for Monetary and Economic Studies, Bank of Japan, Discussion Paper No. 2018-E-4.

Tanaka, T., 2014. Corporate governance and the cost of public debt financing: Evidence from Japan. Journal of the Japanese and International Economies 34, 315–335.

Tate, G., Yang, L., 2015. Female leadership and gender equity: Evidence from plant closure. Journal of Financial Economics 117, 77–97.

Thaler, R. H., Sunstein, C. R., 2003. Libertarian paternalism. American Economic Review 93, 175–179.

Thaler, R. H., Sunstein, C. R., 2008. Nudge: Improving decisions about health, wealth, and happiness. Penguin, New York.

Todorov, K., 2020. Quantify the quantitative easing: Impact on bonds and corporate debt issuance. Journal of Financial Economics 135, 340–358.

Tomisawa, A., Hazama, I., 2019. Record-low 0.0000000091% yield on Japan bond shows BOJ effect. *Bloomberg*. October 16, 2019. Retrieved from https://www.bnnbloomberg.ca/record-low-0-0000000091-yield-on-japan-bond-shows-boj-effect-1.1332580.

Vayanos, D., 2004. Flight to quality, flight to liquidity, and the pricing of risk. NBER Working Paper, No. 10327.

Vayanos, D., Vila, J.-L., 2021. A preferred-habitat model of the term structure of interest rates. Econometrica 89, 77–112.

Vissing-Jorgensen, A., 2021. The Treasury market in Spring 2020 and the response of the Federal Reserve. Journal of Monetary Economics 124, 19–47.

Wang, Y., 2020. A cognitive bias of traders in the stock market. Ph.D. Dissertation. Hong Kong Polytechnic University, Available at https://theses.lib.polyu.edu.hk/handle/200/10670.

Weesie, J., et al., 2000. Seemlingly unrelated estimation and the cluster-adjusted sandwich estimator. Stata Technical Bulletin 9.

Weld, W. C., Michaely, R., Thaler, R. H., Benartzi, S., 2009. The nominal share price puzzle. Journal of Economic Perspectives 23, 121–42.

Wermers, R., 1999. Mutual fund herding and the impact on stock prices. The Journal of Finance 54, 581–622.

Zaghini, A., 2019. The CSPP at work: yield heterogeneity and the portfolio rebalancing channel. Journal of Corporate Finance 56, 282–297.