The Impact of Central Bank Stock Purchases:

Evidence from Discontinuities in Policy Rules*

Masao Fukui[†] Masayuki Yagasaki[‡]

May 2026

Abstract

We trace the impact of central bank stock purchases by exploiting discontinuities in the Bank of Japan's policy rule, which triggers purchases when stock market index movements fall below a threshold. In normal times, the purchases raise the long-term interest rate while leaving no detectable impact on stock prices. After the introduction of yield curve control, which pegs the long-term yield at 0%, interest rates stop responding and stock prices rise robustly following the purchases. While these results support a model in which the stock market is inelastic, taking into account an inelastic bond market is crucial to explain our findings.

*First version: February 2023. We are especially grateful to Emi Nakamura for her continued encouragement and support. We thank Manuel Amador, Anmol Bhandari, Ricardo Caballero, Gabriel Chodorow-Reich, Michael Droste, Xavier Gabaix, Alp Simsek, Jón Steinsson, and seminar participants at UC Berkeley, Minnesota, the Bank of Japan, SED, Barcelona Summer Forum, NBER Summer Institute (ME), Hitotsubashi, and Columbia Japan Economic Seminar for helpful comments. We also thank Hannah Rhodenhiser for excellent research assistance. Fukui gratefully acknowledges the financial support from the George-Orbis Shultz Fund.

†Boston University. mfukui@bu.edu

‡Tohoku University. masayuki.yagasaki.c4@tohoku.ac.jp

1 Introduction

How does a financial flow into the stock market impact asset prices? While textbook models predict almost no effect, the inelastic stock market hypothesis recently proposed by Gabaix and Koijen (2021) argues for a large effect on stock prices. Answering this question goes a long way toward understanding the source of asset price volatility and the effectiveness of central bank asset purchase programs. In fact, as a new form of "quantitative easing," the Bank of Japan (henceforth BoJ) started to purchase stocks in 2010. At the end of 2020, the BoJ owned more than 6% of the stock market capitalization, becoming the largest owner of Japanese stocks worldwide. The BoJ explains that the primary goal of this extreme form of quantitative easing is to "reduce the risk premium," but its effectiveness is often subject to policy debates.

The reason why this question is difficult to answer is the endogeneity of financial flows. No hedge fund randomizes when trading financial assets, and the central bank intervenes in the financial market for a reason. For example, a hedge fund might put money into the stock market when it believes the stock market will perform well, or the BoJ might purchase stocks when the stock market performs poorly. Using these kinds of variations in financial flows will give misleading answers to the question because of reverse causality.

In this paper, we address this challenge by exploiting discontinuities in the BoJ's policy rules. Although the BoJ has never made it public, it is widely known that the BoJ tends to purchase stocks precisely on the day when the changes in the stock market index in the morning session fall below a certain threshold. By comparing days when the movements in a stock market index fall slightly below the threshold and slightly above, the discontinuous increase in inflow into the stock market can be viewed as orthogonal to the underlying economic fundamentals. A difficulty in implementing the standard regression discontinuity design in our setup is that the threshold is not necessarily known, as the BoJ has never made the policy rule public. We overcome this problem by estimating the threshold relying on the econometric literature on regression discontinuity design with unknown cutoffs (Porter and Yu, 2015). Indeed, we find a striking discontinuity in the likelihood of the intervention around the estimated cutoffs.

Exploiting the discontinuity in the policy rules, we first show that flows into stock markets have a large causal impact on *both* the stock prices and the long-term government bond interest rates. In response to an average size of intervention, which amounts to 0.01% of stock market capitalization, our estimates indicate that the stock prices rise by around 0.1-0.4% following the intervention. Perhaps more surprisingly, the 10-year

Japanese government bond (JGB) yield also rises by around 0.5 basis points (b.p.). Although the standard errors increase, these results appear to be persistent, lasting at least several days.

We then argue that the above results mask a stark underlying heterogeneity that depends on the presence of another unconventional monetary policy, yield curve control (YCC). In the middle of 2016, the BoJ pegged the 10-year JGB yield at 0%. Since then, long-term interest rates have stabilized at around 0%, as Figure 1C shows. Given this, it is natural to expect the response of long-term interest rates to the financial flows to be different before and after the introduction of the YCC, and we find that this is indeed the case.

Specifically, before the introduction of the YCC, in response to the BoJ's stock market purchases, the long-term interest rates rise sharply and persistently following the intervention, while leaving virtually no detectable impact on the stock price. Quantitatively, we find that long-term interest rates rise by 1.0-1.5 b.p. in response to the average size of the intervention. In contrast, we cannot reject the null that the BoJ's stock purchases left no impact on stock prices in the following days, although the standard errors are large.

After the introduction of the YCC, the long-term interest rate stops responding entirely, and instead, stock prices rise sharply and persistently in response to the BoJ's stock purchases. The effect on the long-term interest rate is precisely estimated at zero, suggesting that yield curve control successfully stabilizes the long-term interest rate. The stock price responds by around 0.2-0.4% in response to the typical size of the intervention, which persists for several days after the intervention.

In the final part of the paper, we explain these findings with a simple model of inelastic financial markets. Existing models in which only the stock market is inelastic (Gabaix and Koijen, 2021) predict that a central bank stock purchase raises stock prices but leaves bond rates unchanged, as the bond market is perfectly elastic.[1] This class of models, therefore, cannot explain the strong interest rate response in the data. We instead allow not only the stock market, but also the bond market to be inelastic. In this environment, a flow from bonds to stocks puts upward pressure on both stock prices and bond interest rates, with the latter dampening or even reversing the former. Yield curve control shuts down this offsetting force, so the same flow is absorbed entirely by stock prices.

Our results are crucial in understanding the consequences of unconventional monetary policy. As unconventional monetary policies become increasingly important, other central banks may follow in the footsteps of the Bank of Japan in purchasing not only bonds but also stocks. Since most central banks do not intervene in the stock market, evidence of the

[1]We do not make a distinction between money and bonds in the model. Bonds include all money-like assets that are liquid and risk-free.

effects of such policies is very limited. Our results suggest that this policy has far-reaching impacts on the financial markets.

Related Literature

We build most directly on the pioneering work by Gabaix and Koijen (2021) in assessing how a flow from bonds to stocks impacts the financial market. Gabaix and Koijen (2021) and a growing number of studies (e.g., Da et al., 2018; Hartzmark and Solomon, 2023; Li et al., 2021; Parker et al., 2023) estimate how flows from bonds to stocks affect stock prices using various identification strategies. Our contributions to this literature are twofold. First, we argue that from both empirical and theoretical perspectives, it is important to consider how bond prices are affected by such flows.[2] In this regard, our results provide empirical support for the theoretical arguments in Fuchs et al. (2023) that taking into account cross-asset spillovers plays a crucial role in asset demand estimation. Second, we propose a novel identification strategy relying on discontinuity in the Bank of Japan's policy rules. An advantage of our approach is that it transparently points toward the source of identification and the underlying assumptions.

To the best of our knowledge, our paper is the first to identify the aggregate effect of the Bank of Japan's stock purchases. Many studies exploit the difference in weights in the BoJ's purchase basket to identify relative price impacts on stocks with larger versus smaller weights (Barbon and Gianinazzi, 2019; Charoenwong et al., 2021; Harada and Okimoto, 2021; Adachi et al., 2021; Katagiri et al., 2022). In contrast, our empirical strategy allows us to focus on the aggregate effect. In this regard, various studies (Shirota, 2018; Fukuda and Tanaka, 2022; Chung, 2020; Hattori and Yoshida, 2023; Ichiue, 2024) assume selection on observables and use the unexplained policy variation that remains after conditioning on observables to identify the aggregate effect. However, it is difficult to rule out the presence of unobservables that simultaneously affect the financial market performance and the likelihood of BoJ intervention. Our identification assumptions, which only require continuity of economic fundamentals with respect to stock price changes in the morning session, are substantially weaker than those of these studies.

More broadly, we contribute to the large literature studying the effect of central bank asset purchases, so-called "quantitative easing" (e.g., Krishnamurthy and Vissing-Jorgensen, 2011; Chodorow-Reich, 2014; Droste et al., 2021; Selgrad, 2023). These studies have focused on central bank purchases of long-term government bonds or mortgage-backed securities,

[2]Recent work by Caballero et al. (2024) also shows that a flow from bonds to stocks has a positive impact on interest rates.

which are swaps between one type of bond (e.g., long-term bonds) and another (e.g., reserves). Our focus is conceptually different from these because central bank stock purchases are swaps between stocks and bonds.

While there are many studies that isolate quasi-experimental variation in monetary policy, our approach is unique in exploiting the discontinuity in the policy rule. The closest to our approach is the one in Kuersteiner, Phillips, and Villamizar-Villegas (2018), who also use a discontinuous policy rule to investigate the effectiveness of sterilized foreign exchange interventions in Colombia. Our approach differs not only in terms of the empirical context but also in methodology, as we use the technique of regression discontinuity with unknown discontinuity points (Porter and Yu, 2015).

2 Data

Our primary goal is to measure the impact of an inflow into the stock market induced by central bank stock purchases. We focus on the period from October 2010, the start of the BoJ stock market intervention program, to the end of 2020. The BoJ purchases Exchange Traded Funds (ETFs) indexed to stock market indexes in Japan. We do not distinguish between ETFs and stocks and use the terms interchangeably. We obtain the dates and amounts of stock purchases for each of the BoJ's interventions from the BoJ website. Figure 1A shows the cumulative amount of ETF purchases over time. The BoJ started the stock market purchases in December 2010 as a new form of quantitative easing. By the end of 2020, the BoJ held over 6% of the stock market capitalization in Japan. The average size of each intervention is roughly 500 million USD, corresponding to 0.01% of the stock market capitalization. The total number of interventions during our sample period is 668. The amount of ETF purchases is normalized by the stock market capitalization, which we obtained from the Japan Exchange Group Data Cloud.

The BoJ publishes the amount of ETF purchases on the morning of the day following the intervention. Based on the trading volume, it is widely believed that the BoJ submits an order at lunchtime, although the BoJ has never made this practice public (Harada and Okimoto, 2021). Therefore, investors potentially face uncertainty about whether the large inflow into the stock market reflects the BoJ's intervention or other factors on the day of the intervention. For this reason, we prefer to use our empirical estimates on the next day as our benchmark estimates.

To measure the response of the stock market, we use tick-by-tick data on the Tokyo Stock Price Index (henceforth, TOPIX), which is the index of the Tokyo Stock Exchange in Japan, tracking all domestic companies of the exchange's first section. We obtain these

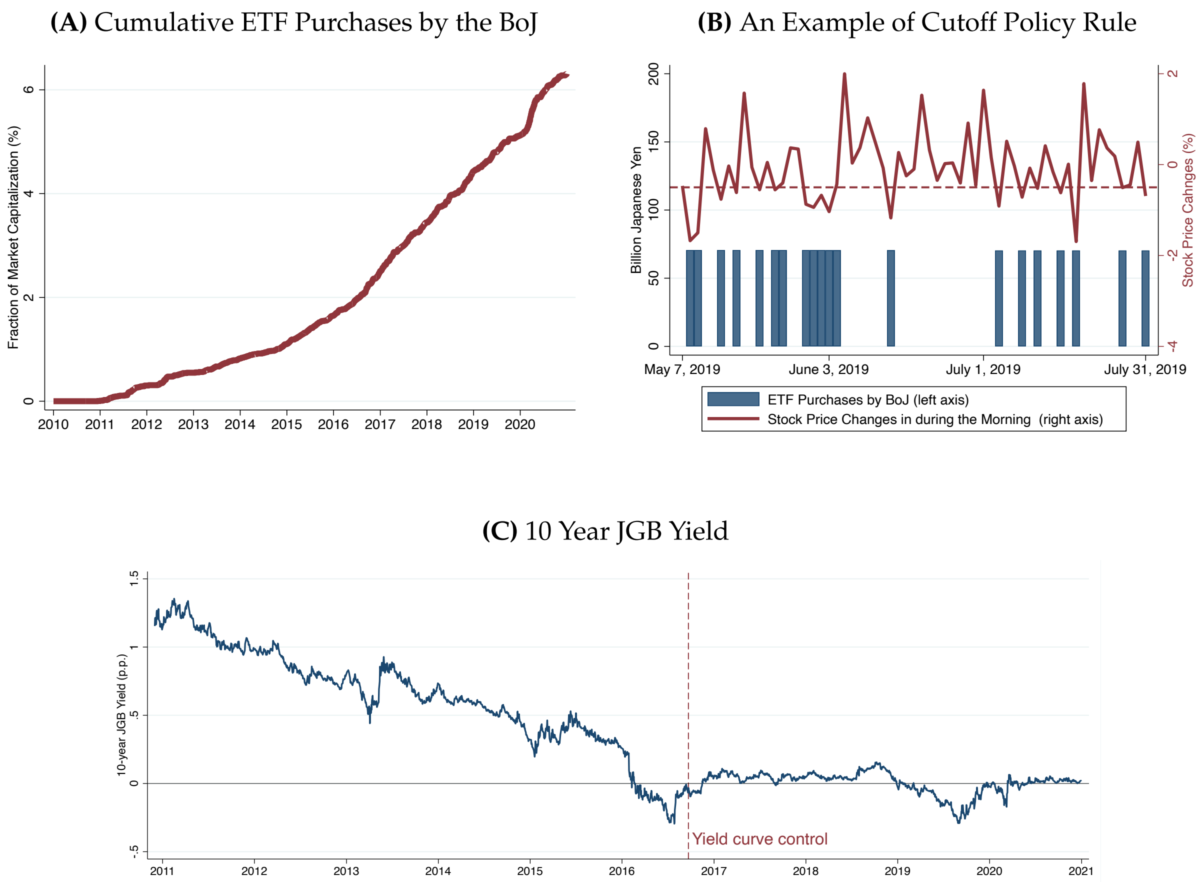

Figure 1: Unconventional Monetary Policies by the Bank of Japan

Notes: Panel A plots the cumulative amount of stock purchases by the BoJ from 2010 to 2020 as a fraction of stock market capitalization. Panel B illustrates the cut-off policy rule by showing the percentage TOPIX changes and the BoJ (Bank of Japan) purchase amount for each day from May 2019 to July 2019. The solid red line shows the TOPIX changes in the morning session, and the dashed red line is the estimated cutoff of 0.5%. The bar shows the amount of purchases for each intervention in billions of Japanese Yen (approximately 10 million US dollars). Panel C shows the path of 10-year JGB (Japanese Government Bond) yield over time, where the red vertical dashed line (September 21, 2016) denotes the start of the yield curve control.

data from the Japan Exchange Group Data Cloud. To measure the response of the long-term interest rate, we use tick-by-tick Japanese Government Bond yield data, which we obtained from Refinitiv Japan. Since some observations are missing in the Refinitiv data, we supplement those missing observations with the tick-by-tick data from Bloomberg.

3 Empirical Strategy

We consider the following econometric model:

$$\Delta y_{t+l,h} = \beta_{l,h} \times ETF_t + \Gamma'_{l,h}\mathbf{X}_t + \epsilon_{t+l,h}, \tag{1}$$

where $\Delta y_{t+l,h} \equiv y_{t+l,h} - y_{t,0}$ is the change in the variable y (e.g. the log of stock prices) from the end of the morning session ($h = 0$) on day t to time h on day $t + l$, ETF_t is the amount of stock market purchases by the BoJ relative to the stock market capitalization of Japan, $\mathbf{X}_t$ is the vector of controls, and $\epsilon_{t+l,h}$ contains the unmodeled determinants of the outcome variable. We are interested in estimating $\beta_{l,h}$, which measures the effect of the central bank's stock purchases on the outcome observed at time h, l days after the purchase. We choose this simple linear model for expositional purposes. In Appendix A.1, we consider a non-linear model of (1) and present a more technical interpretation of the estimated parameter as in Angrist and Imbens (1995). In the baseline specification, we do not include any controls $\mathbf{X}_t$, but we show robustness to adding various controls in Section 4.3.

An obvious concern with estimating equation (1) by OLS is reverse causality. For example, one might expect that the central bank is more likely to intervene when the stock market is performing poorly. This leads to a downward bias in the OLS estimates of $\beta_{l,h}$.

To solve this endogeneity problem, we propose an identification strategy based on a regression discontinuity design, which builds on the observation that the BoJ's intervention appeared to follow a cut-off rule. It has been widely argued in the media that the BoJ seemed to intervene on days when movements in the value of TOPIX fell below a certain threshold in the morning session. For example, the Financial Times writes that "the central bank has tended to step in whenever the TOPIX index has lost more than 0.5 percent in the morning session."[3] In fact, Figure 1B shows that from May to July 2019, the BoJ followed a strict rule to intervene when the stock market index fell by more than 0.5%

[3] Financial Times, "Bank of Japan backs away from ETF buying scheme," (March 23, 2021), https://www.ft.com/content/a654d1c9-7126-4587-8de6-ed15f567455f.

in the morning session. The BoJ intervenes when the index falls slightly below the 0.5% threshold, while it does not intervene when the index falls slightly above the threshold.

Suppose for the moment that such a cut-off is known. Then, we can apply a standard regression discontinuity design. Formally, we assume that the policy rule has the following form:

$$ETF_t = ETF_{-,t}(\Delta p_t)\mathbb{I}(\Delta p_t < c_t) + ETF_{+,t}(\Delta p_t)\mathbb{I}(\Delta p_t \geq c_t), \tag{2}$$

where Δp_t is the log-change in the TOPIX value in the morning, c_t is the cut-off, and $ETF_{-,t}$ and $ETF_{+,t}$ are some random functions of Δp_t that represent different policy rules depending on whether Δp_t is above or below the cutoff. We assume (i) $\mathbb{E}[\epsilon_{t+l,h}|\Delta p_t, \mathbf{X_t}]$ is continuous at $\Delta p_t = c_t$, (ii) $\lim_{\Delta p \uparrow c_t} \mathbb{E}[ETF_t|\Delta p_t = \Delta p, \mathbf{X_t}]$ and $\lim_{\Delta p \downarrow c_t} \mathbb{E}[ETF_t|\Delta p_t = \Delta p, \mathbf{X_t}]$ exist, and (iii) $\lim_{\Delta p \uparrow c_t} \mathbb{E}[ETF_t|\Delta p_t = \Delta p, \mathbf{X_t}] \neq \lim_{\Delta p \downarrow c_t} \mathbb{E}[ETF_t|\Delta p_t = \Delta p, \mathbf{X_t}]$. Under these assumptions, it follows that

$$\frac{\lim_{\Delta p \uparrow c_t} \mathbb{E}[\Delta y_{t+l,h}|\Delta p_t = \Delta p, \mathbf{X_t}] - \lim_{\Delta p \downarrow c_t} \mathbb{E}[\Delta y_{t+l,h}|\Delta p_t = \Delta p, \mathbf{X_t}]}{\lim_{\Delta p \uparrow c_t} \mathbb{E}[ETF_t|\Delta p_t = \Delta p, \mathbf{X_t}] - \lim_{\Delta p \downarrow c_t} \mathbb{E}[ETF_t|\Delta p_t = \Delta p, \mathbf{X_t}]} = \beta_{l,h}. \tag{3}$$

As recommended by Hahn, Todd, and Van der Klaauw (2001) and Porter (2003), we can devise local linear regression estimators for the left-hand side to obtain an estimate of $\beta_{l,h}$. We use the optimal bandwidth proposed by Calonico, Cattaneo, and Titiunik (2014) and estimate $\beta_{l,h}$ using two-stage least squares and report Newey-West standard errors.

The difficulty in implementing the above approach, however, is that the cut-off is not necessarily known. While the cutoffs utilized by the BoJ were apparently known to the public in some periods, they were not known in other periods. In order to formally detect the cutoffs utilized by the BoJ, we estimate the cut-off with the presumption that the BoJ follows a cut-off rule, using the approach proposed by Porter and Yu (2015). They develop a method to estimate the discontinuity point and show that there is no loss of efficiency with the regression discontinuity estimator using the estimated cutoff. In implementing this approach, we proceed as follows. We first split the sample period to allow time variation in the policy rule. We assume the cut-off is a constant within the sample split. Then, in each sample split, we consider a set of possible cutoffs, $\mathbb{C} \equiv \{\bar{c}_1, \bar{c}_2, \ldots, \bar{c}_K\}$. For each $\bar{c} \in \mathbb{C}$, we estimate the jump of $\Pr_t(ETF_t > 0|\Delta p)$ around $\bar{c}$, which is

$$J_t(\bar{c}) \equiv \lim_{\Delta p \uparrow \bar{c}} \Pr_t(ETF_t > 0|\Delta p) - \lim_{\Delta p \downarrow \bar{c}} \Pr_t(ETF_t > 0|\Delta p).$$

We select $\bar{c}$ that maximizes the square of the jump, $J_t^2(\bar{c})$: $c_t^* \in \arg\max_{\bar{c} \in \mathbb{C}} J_t^2(\bar{c})$.

We implement the above approach with the following specifications. First, we split the sample period based on the BoJ's announcements regarding ETF purchases. The BoJ made six announcements that stated changes in the target amount of ETF purchases on March 4, 2013, October 31, 2014, December 18, 2015, July 29, 2016, July 31, 2018, and March 16, 2020. We further divide each period between two announcements, based on whether the TOPIX closing price falls relative to the opening price for the last two consecutive days. We make this choice based on widespread claims in the media,[4] and we have indeed found that it has strong explanatory power. Second, we consider the set of potential cutoffs ranging from -1% to 0% with 0.05% intervals. We estimate the jump of $\Pr_t(ETF_t > 0|\Delta p)$ around the potential cutoffs using local linear regressions with the optimal bandwidth computed from Calonico, Cattaneo, and Titiunik (2014).

Figure 2A shows the path of estimated cutoffs. The estimated cutoffs align well with the widely held consensus. During 2010-2013, it was widely believed that the BoJ followed a so-called "1% rule", in which the BoJ buys ETFs whenever the TOPIX falls by more than 1% in the morning session,[5] and our estimates confirm this view. Since April 2013, the BoJ appears to use different cutoffs depending on whether the daily change in the TOPIX has been negative for the past two consecutive days. Specifically, when the index has declined for two consecutive days, our estimates show a clear pattern in which the BoJ intervenes whenever the TOPIX falls below 0% in the morning session, confirming the claim in the media described earlier. Since March 2018, the cutoffs appear to be 0.5% when there is no consecutive fall in the past two days, which is again consistent with the so-called "0.5% rule."

Figure 2B shows the binned scatter plot of the size of the BoJ intervention against changes in the TOPIX in the morning session of the same day relative to the cutoffs. We confirm that there is a discrete jump in the size of the BoJ interventions around zero. The implied jump in the full sample is 0.0083% of the market capitalization with a standard error of 0.0005%. The Cragg-Donald F-statistic is 1821, and the Kleibergen-Paap F-statistic is 261, eliminating weak identification concerns. This discontinuity comes from the discontinuity in the likelihood of intervention, with a jump in the probability of intervention of 86% with a standard error of 0.02%. Importantly, we find strong evidence of discontinuity in any split of the sample.[6]

[4]See, for example, Bloomberg article "The BoJ's ETF Purchase Conditions Likely to Ease if Stocks Continue to Fall" (written in Japanese) (`https://www.bloomberg.co.jp/news/articles/2020-07-22/-0-3-kcwteezj`).

[5]For example, Nikkei Asia writes "the BoJ was widely thought to be following an unwritten rule, dubbed the 1% rule: it would buy ETFs when the Topix index of all issues on the first section of the Tokyo Stock Exchange fell more than 1% in the morning session." (https://asia.nikkei.com/Business/Finance/BOJ-steps-up-REIT-buying-scales-back-ETF-purchases)

[6]We report the discontinuity for each sample split in Appendix A.2 and Figure B.1.

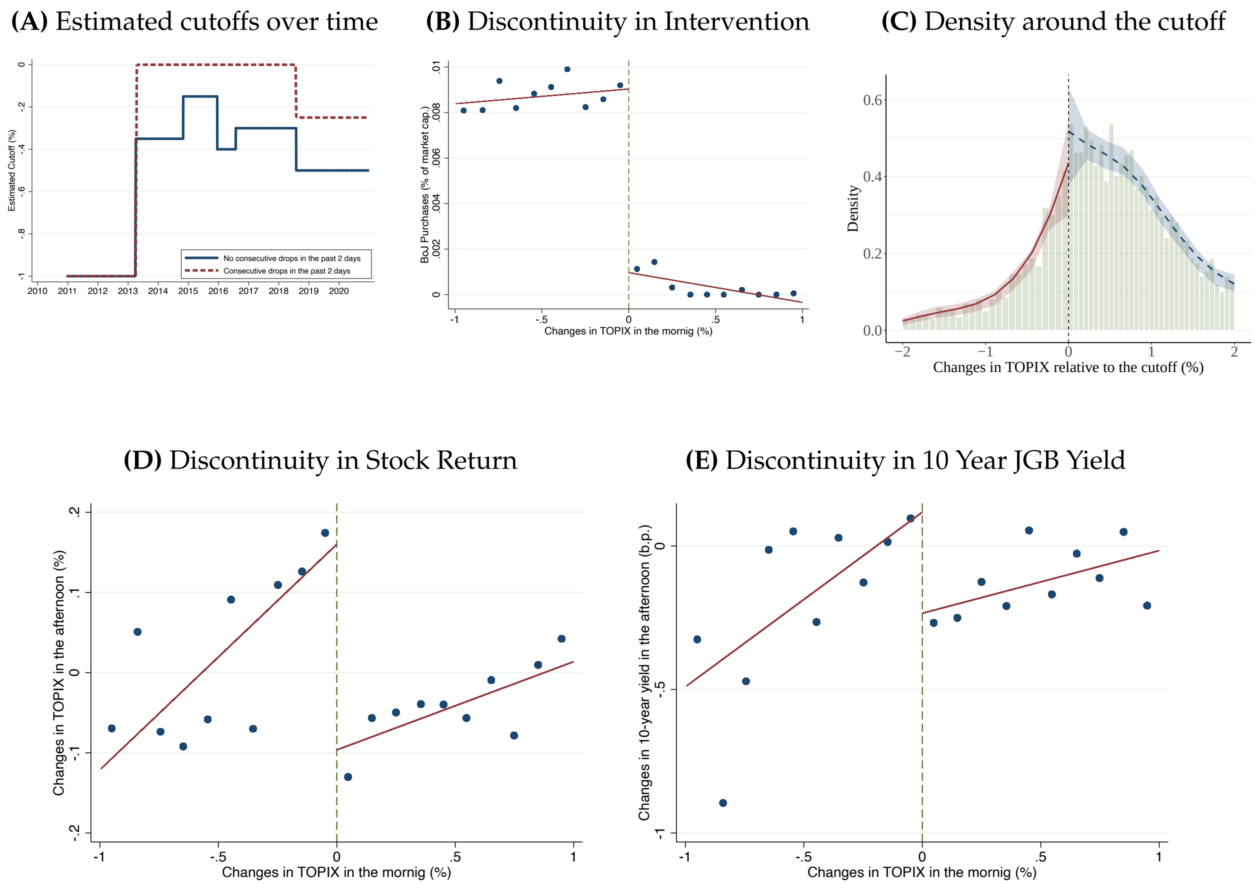

Figure 2: Discontinuities in the Policy Rule and Financial Market Outcomes

Notes: Panel A plots the path of estimated cutoffs over our sample period. Panel B shows the discontinuity in the amount of BoJ stock purchases in the range of -1% to 1% around the estimated cutoff. Each dot represents the binned scatter plot with a 0.1% bin width and the red line represents the linear fit on each side of the cutoff. Panel C shows the histogram and the density of changes in TOPIX relative to the cutoff. The shaded area is the 95% confidence interval. We use the local polynomial density estimator of Cattaneo et al. (2020) with order 2. Panel D shows the binned scatter plot of the log changes in TOPIX in the afternoon (from 11AM to 3PM) against changes in TOPIX in the morning session relative to the cutoff. The bin width is 0.1%. The line represents the best fit from the linear regression. Panel E is analogous to Panel D, with the vertical axis being the changes in the 10-year JGB yield in basis points (b.p.) in the afternoon (from 11AM to 3PM).

A natural concern for a discontinuity-based research design is that whether the index falls below the cutoff might not be random. For example, to the extent that the policy rule is widely known among investors, the optimal responses of investors may induce bunching of stock price changes in the morning session just below the cutoff. Under such manipulation of the stock price changes, the regression discontinuity estimator lacks causal interpretation, as explained in McCrary (2008).

Figure 2C plots the density of stock price changes in the morning session relative to the estimated cutoff. Under manipulation of stock price changes, we would expect bunching in the density on either side of the cutoff. However, we do not observe such bunching. Furthermore, Appendix A.3 formally tests the presence of manipulation using the methodology proposed by Cattaneo et al. (2020). We do not find evidence of manipulation. This supports our assumption that the likelihood of falling below the cutoff is random. This likely reflects the presence of uninformed ("noise") traders in the financial market, making it difficult for investors to manipulate stock prices.

4 Empirical Results

Armed with the cutoff estimates, we implement the regression discontinuity design to assess the impact of the BoJ's ETF purchases on the financial market.

4.1 Homogeneous Effect for the Entire Sample Period

Figure 2D shows the results for stock price changes. Specifically, it reports the binned scatter plot of changes in the TOPIX in the afternoon (from 11AM to 3PM) against changes in the TOPIX in the morning relative to the estimated cutoff. The figure shows that the stock returns were around 0.2 p.p. higher when the TOPIX fell slightly below the cutoff in the morning than when it fell slightly above the cutoff. Since the BoJ's intervention was likely to occur on days when the TOPIX fell below the threshold, and since the BoJ submitted the order to purchase ETFs during the lunch break, this suggests that the BoJ's intervention had a large impact on stock prices within the day. The magnitude is large considering that the BoJ purchased on average around 0.01% of market capitalization in each intervention.

Figure 2E focuses on the 10-year Japanese government bond yield as an outcome variable. Perhaps surprisingly, we also see a discontinuity in the long-term interest rate. The long-term interest rate is 0.4 basis points higher on the left side of the cutoff than on the right side. These within-day discontinuities motivate the dynamic specification in (1), which scales the price responses by intervention size and traces the dynamic responses.

Figures 3A and 3B plot the impulse response functions of stock prices and the 10-year JGB yield. Formally, we plot estimates of $\beta_{l,h}$ in equation (1) for each l and h, where l represents the number of days since the intervention and h represents the time in hours. In Figure 3A, we see an immediate and large stock price response in the afternoon of the intervention. It implies that a stock purchase of 0.01% of market capitalization, a typical size of the intervention, increases the stock price by 0.4%. This coefficient is statistically significant. Over the next five days, the coefficient is roughly halved and the standard error is larger, but it does not revert to zero. Reassuringly, we do not find any evidence of pre-trends, which is consistent with our identifying assumptions.

Figure 3B shows that the 10-year JGB yield also rises sharply after the intervention. Moreover, the effect appears to be quite persistent, and it remains statistically significant even several days after the intervention. The magnitude is again substantial. In response to a typical purchase size (0.01% of market capitalization), the 10-year JGB yield rises by around 0.4-0.5 basis points.

4.2 Heterogeneous Effects and Yield Curve Control

Figure 3B showed that the stock market purchases were accompanied by a rise in the long-term interest rate. In standard theoretical models, the rise in interest rates puts downward pressure on stock prices. For this reason, we expect that the ability of the interest rate to respond will be a critical determinant of the stock price responses to central bank stock purchases.

Another unconventional policy of the BoJ, the so-called "yield curve control," provides an ideal laboratory to explore this hypothesis. On September 21, 2016, the BoJ introduced an explicit target for the 10-year Japanese government bond yield at 0%. Figure 1C indeed shows that the long-term rate has stabilized at around 0% since the introduction of yield curve control. The daily standard deviation of the long-term rate is 0.37% before the introduction of yield curve control, but it falls to 0.08% after the introduction. If the BoJ does its best to stabilize the long-term interest rate at 0%, we would expect to see a much smaller response of the long-term rate in response to the stock purchases. To test this, we split our sample into periods before and after the introduction of yield curve control and rerun our analysis.

Figures 3C–3F show the main results of this paper. Figures 3C and 3D show the impulse response of stock prices before and after the introduction of yield curve control, respectively. Before the introduction of yield curve control, we find no evidence that the

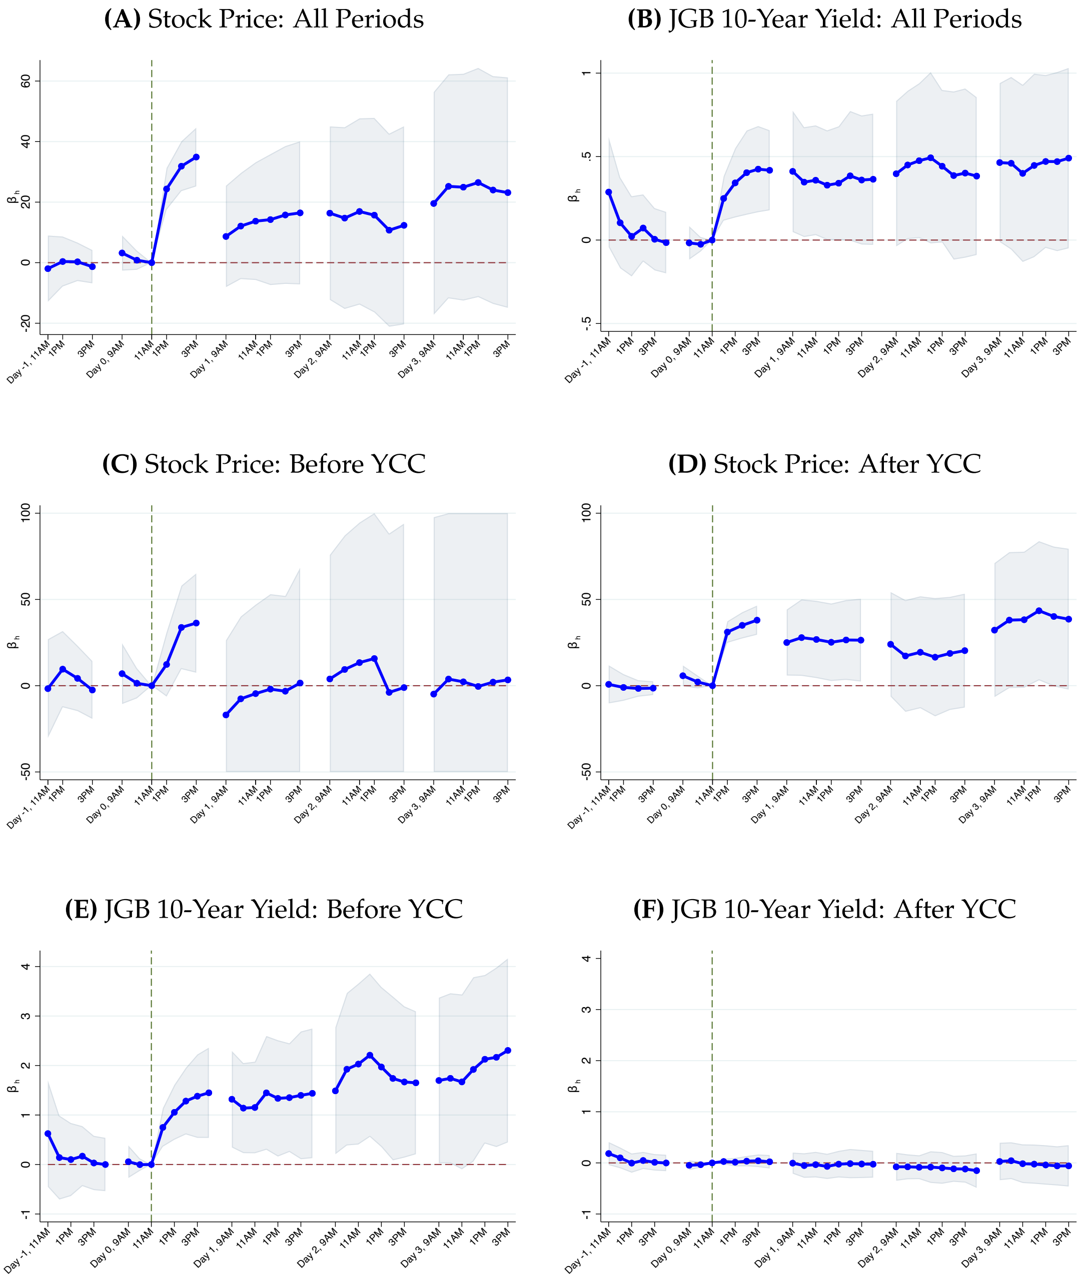

Figure 3: The Impact on Stock Prices and Long-Term Interest Rates

Notes: Panels A and B show the impulse response functions of stock prices and the 10-year JGB yield by plotting the coefficient $\beta_{l,h}$ in equation (1). The stock price coefficient measures the log changes in stock prices in response to stock purchases of 1% of market capitalization, and the yield coefficient measures the percentage point changes in the yield in response to stock purchases of 1% of market capitalization. Panels C and D show the impulse response of stock prices separately estimated before and after yield curve control, which is analogous to Panel A. Panels E and F show the impulse response of the 10-year JGB yield separately estimated before and after yield curve control, which is analogous to Panel B. In all figures, the shaded areas represent 90% confidence intervals, which account for heteroskedasticity and autocorrelation.

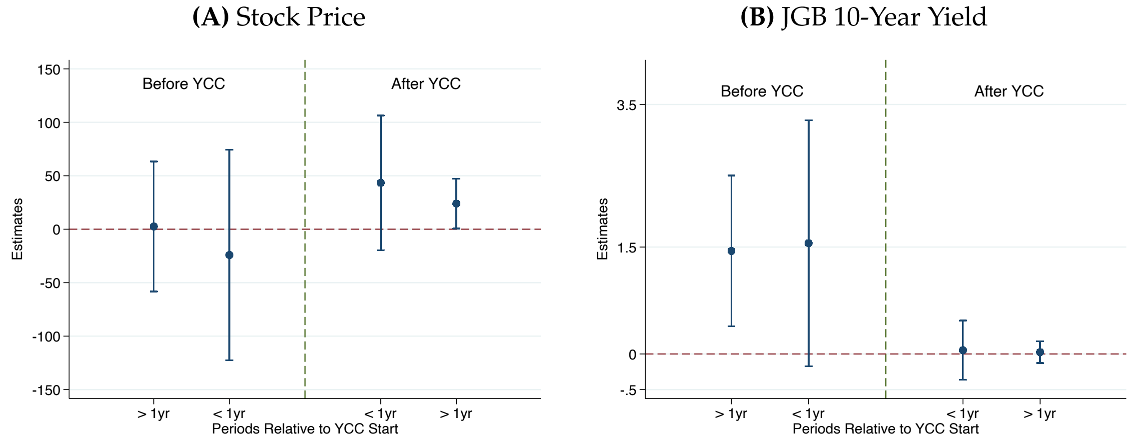

Figure 4: Heterogeneity Around the Introduction of Yield Curve Control

Notes: Panels A and B show, respectively, the responses of stock prices and 10-year JGB yields, divided into two periods relative to the introduction of the YCC on September 21, 2016: (i) more than one year before or after YCC (labeled "> 1yr"), and (ii) within one year of the YCC's introduction (labeled "< 1yr"). The estimates reflect the response to changes from 11 AM on the intervention day to 9 AM the following day. All confidence intervals account for heteroskedasticity and autocorrelation.

stock market responded positively beyond one day to the BoJ intervention, although the standard error is large. In stark contrast, stock prices rise persistently and in a statistically significant manner under yield curve control. Quantitatively, a purchase of 0.01% of stock market capitalization, a typical size of the intervention, causes around a 0.2-0.3% increase in stock prices several days after the intervention.

Figures 3E and 3F explain why. Figure 3E shows that the long-term interest rate responds positively before yield curve control. Quantitatively, a purchase of 0.01% of stock market capitalization by the BoJ causes around a one-basis-point increase in long-term rates, and the effect is statistically significant. However, under yield curve control, long-term rates stopped responding and the effect is precisely estimated as zero. These results imply that yield curve control was successful in stabilizing the long-term interest rate.

One potential concern here is that the observed heterogeneity may not reflect the causal impact of yield curve control. For example, it is possible that the effects of the BoJ's policy intervention gradually increased or decreased over time, irrespective of the implementation of yield curve control. To address this concern, we estimate the effects of the BoJ's stock purchases by focusing on the one-year periods immediately before and after the introduction of yield curve control.

Figures 4A and 4B plot the estimated next-day responses of the stock price and long-term interest rates, dividing the periods into four subsample periods: more than one year before, less than one year before, less than one year after, and more than one year after the start of yield curve control. We continue to find the stark heterogeneity even if we focus on the periods shortly before and after the implementation of yield curve control. While standard errors are large, these results support the notion that the heterogeneity reflects the causal effect of yield curve control.

4.3 Robustness

Table 1 conducts a battery of robustness checks and shows that our results are robust to various alternatives to the baseline specifications. Rows 1-2 vary the bandwidth, row 3 uses a quadratic local polynomial, row 4 controls for BoJ purchases over the previous two days, rows 5-6 control for lagged stock returns and long-rate changes, and row 7 drops observations within one week of cutoff changes. In all cases, the estimates are similar to the baseline estimates. Appendix Figure B.4 explores bandwidths more systematically with similar conclusions, and Appendix A.4 shows that falling below the cutoff today does not predict purchases on other days.

4.4 Discussion

We close our empirical analysis by discussing several potential concerns. First, Appendix C reports placebo tests at cutoffs that should not trigger BoJ purchases. The placebo estimates are generally indistinguishable from zero and, when statistically significant, have the opposite sign from the baseline estimates. This supports the interpretation that our results are driven by the BoJ policy rule rather than by other factors that may cause discontinuities at arbitrary cutoffs.

Second, if the BoJ is selling long-term government bonds at the same time as the stock purchases, then it is not surprising that the long-term interest rate rises in response to the stock purchases. However, during our sample period, the BoJ sold government bonds only twice, on March 24, 2017, and March 23, 2020, both of which were after yield curve control. Therefore, we can rule out this concern.

Third, while our results so far concern nominal stock prices and interest rates, one may be concerned that our results are driven entirely by inflation expectations. We show in Appendix B.5 that the response of the inflation swap rate is indistinguishable from zero.

Fourth, the interest rate response is not specific to the 10-year JGB yield but is widespread across the yield curve. In Appendix D, we show that before yield curve control, yields

Panel A. Stock Price Response

	All		Before YCC		After YCC	
	Same Day	Next Day	Same Day	Next Day	Same Day	Next Day
0. Baseline	34.90	8.65	36.31	-16.95	35.24	22.23
	(5.89)	(10.18)	(17.50)	(26.27)	(4.68)	(11.13)
1. Narrower	41.65	16.22	30.14	-21.25	47.26	33.08
Bandwidth	(7.16)	(16.09)	(22.18)	(39.57)	(7.15)	(16.25)
2. Wider	29.49	6.07	30.63	-2.44	29.41	15.66
Bandwidth	(4.84)	(9.36)	(13.46)	(20.74)	(4.69)	(9.14)
3. Polynominal	43.40	16.98	41.13	-22.70	47.24	30.63
Order 2	(7.47)	(15.64)	(21.64)	(42.42)	(7.14)	(14.39)
4. Control Past	39.50	12.56	54.57	-16.44	37.32	29.65
Interventions	(8.09)	(15.10)	(25.67)	(34.04)	(6.61)	(14.48)
5. Control Past	34.91	8.62	36.15	-17.33	35.18	20.51
Stock Returns	(5.89)	(10.20)	(17.70)	(26.21)	(4.77)	(11.76)
6. Control Past	35.96	6.63	40.91	-23.41	35.06	21.90
10-Year Yield	(5.84)	(9.94)	(18.33)	(25.42)	(4.67)	(11.13)
7. Drop Around	34.96	8.00	31.21	-20.53	35.07	24.56
the Cutoff Changes	(6.05)	(10.92)	(16.52)	(28.22)	(5.52)	(12.45)

Panel B. JGB 10-Year Yield Response

	All		Before YCC		After YCC	
	Same Day	Next Day	Same Day	Next Day	Same Day	Next Day
0. Baseline	0.47	0.37	1.52	1.41	0.04	-0.00
	(0.18)	(0.23)	(0.59)	(0.68)	(0.07)	(0.13)
1. Narrower	0.54	0.43	1.96	1.96	-0.00	-0.13
Bandwidth	(0.22)	(0.30)	(0.84)	(0.96)	(0.09)	(0.20)
2. Wider	0.40	0.35	1.28	1.05	0.01	-0.04
Bandwidth	(0.16)	(0.19)	(0.47)	(0.50)	(0.06)	(0.10)
3. Polynominal	0.52	0.43	1.81	1.75	0.07	0.01
Order 2	(0.22)	(0.27)	(0.77)	(0.91)	(0.09)	(0.18)
4. Control Past	0.47	0.37	1.73	1.79	-0.06	-0.17
Interventions	(0.23)	(0.31)	(0.76)	(0.90)	(0.11)	(0.21)
5. Control Past	0.47	0.37	1.53	1.43	0.03	-0.03
Stock Returns	(0.18)	(0.23)	(0.59)	(0.68)	(0.08)	(0.13)
6. Control Past	0.46	0.37	1.51	1.42	0.04	-0.01
10-Year Yield	(0.17)	(0.23)	(0.60)	(0.69)	(0.07)	(0.13)
7. Drop Around	0.37	0.28	1.12	0.91	0.02	-0.05
the Cutoff Changes	(0.16)	(0.21)	(0.44)	(0.48)	(0.06)	(0.12)

Table 1: Robustness

Notes: Table 1 shows robustness checks against various modifications of our benchmark specifications. Panels A and B show the responses of stock prices and the JGB 10-year yield, respectively. In each panel, row 0 reports the baseline estimates. Row 1 considers a bandwidth that is 50% less than the original one. Row 2 considers a bandwidth that is 50% larger than the original one. Row 3 considers a local polynomial regression of order 2 instead of 1. Row 4 controls for the BoJ's stock purchases over the past two days. Row 5 controls for stock market returns over the past two days. Row 6 controls for changes in the 10-year yield over the past two days. Row 7 drops observations within one week before and after the cutoff changes. The same-day response indicates the changes in the outcome variable from 11AM to 3PM on the same day of the intervention. The next-day response indicates the changes in the outcome variable from 11AM on the day of the intervention to 9AM on the next day. Standard errors, which account for heteroskedasticity and autocorrelation, are reported in parentheses.

rose at all maturities with larger effects at longer maturities, while after yield curve control all maturities stopped responding.

The final issue concerns the interpretation of our results. One may be concerned that our results are driven by the signaling channel, whereby the BoJ's stock purchases send signals about the BoJ's future policy stance. However, the BoJ announces the target amount of stock purchases each year in advance. Therefore, whether or not the BoJ purchases stocks today should not reveal the BoJ's future policy stance.

5 An Explanation

In this section, we provide a simple explanation for our empirical findings. We consider a model with two assets: stocks and bonds. We do not make a distinction between money and bonds.[7] We show that a model with an inelastic stock market and an even more inelastic bond market can account for our findings. We focus on the key equilibrium condition here, and Appendix E provides the details of the setup and the derivation, as well as the calibration and quantification of the model mechanisms.

5.1 Model

Time is discrete, and the horizon is infinite. Throughout, we focus on the steady-state equilibrium. There is no aggregate uncertainty. The economy is populated by a representative household that invests only in bonds, a continuum of financial intermediaries (funds) that invest in both stocks and bonds on behalf of households, and a consolidated central bank and government that can purchase stocks, financed by bond issuance or lump-sum taxes. Stocks are claims to the exogenously given output of the economy, and we normalize the supply of stocks to one.

Most importantly, we assume that households derive utility from the liquidity services that bonds provide, as in Krishnamurthy and Vissing-Jorgensen (2012), and the funds' portfolio choice is imperfectly elastic to the return differences between stocks and bonds, as in Gabaix and Koijen (2021).

The bond market is characterized by the following modified Euler equation after imposing the market-clearing conditions:

$$u'(Y) = \beta R^b u'(Y) + v'(B^g), \tag{4}$$

[7] This is a valid assumption if the elasticity of substitution between near-money assets like bonds and money is high, which is empirically the case (Nagel, 2016; Krishnamurthy and Li, 2023).

where $u'(\cdot)$ is the marginal utility of consumption, β is the discount factor, Y is the aggregate output, B^g is the government bond issuance, R^b is the bond interest rate, and $v'(\cdot)$ is the marginal utility from the liquidity services that bonds provide. We assume that $v'(\cdot)$ is decreasing.

The stock market is characterized by the following market clearing condition:

$$Q = QS^g + \mathcal{S}(Q, R^b), \tag{5}$$

where Q is the stock price, S^g is the government's stock holdings, and $\mathcal{S}$ is the stock holdings of the funds. The left-hand side is the supply of stocks (in value terms), and the right-hand side is the demand for stocks, which comes from the government and the funds. Naturally, we impose that the demand function for stocks is decreasing in the stock price and the bond interest rate. As argued in Gabaix and Koijen (2021), many standard models feature an extremely high elasticity of stock demand $\mathcal{S}$ with respect to prices (Q, R^b). Our formulation here allows arbitrary elasticity, which we microfound through heterogeneous beliefs in Appendix E.

Given the government policies $\{S^g, B^g\}$, the equilibrium stock price and the bond interest rate, $\{Q, R^b\}$, solve (4) and (5).

5.2 Central Bank Stock Purchases in the Model

We model the central bank stock purchases as a small permanent increase in S^g around an equilibrium with $S^g = B^g = 0$. We consider two scenarios that differ in how the stock purchases are financed. In the first experiment, the central bank finances the stock purchases with an equal amount of bond (money) issuance, $dB^g_t = QdS^g$. This experiment aims to replicate the Bank of Japan's stock purchases before the implementation of yield curve control. In the second experiment, the government finances the stock purchases entirely through a lump-sum tax on households so that the bond interest rate is kept constant. This experiment aims to replicate the Bank of Japan's stock purchases after the implementation of yield curve control.[8]

We have the following expression for the stock price impact of central bank stock

[8] Empirically, throughout our sample period, the stock purchases were financed by central bank reserves (money), and the commitment to peg the long-term rate was sufficient to maintain the peg without actual bond-market intervention. Our model abstracts from these features. See Itskhoki and Mukhin (2025) for a related theoretical mechanism in the context of exchange rate pegs.

purchases:

$$\frac{d\ln Q}{dS^g} = \underbrace{\frac{1}{1+\epsilon^s_Q}}_{\geq 0} \underbrace{-\frac{\epsilon^s_{R^b}}{1+\epsilon^s_Q}\frac{d\ln R^b}{dS^g}}_{\leq 0}, \tag{6}$$

where $\epsilon^s_Q \equiv -\frac{\partial \ln \mathcal{S}}{\partial \ln Q} > 0$ and $\epsilon^s_{R^b} \equiv -\frac{\partial \ln \mathcal{S}}{\partial \ln R^b} > 0$ are the demand elasticities of stocks with respect to the stock price and the bond interest rate, respectively. The first term captures the direct effect of the stock purchases on the stock price, holding the bond interest rate constant. This term is positive, as an increase in stock demand naturally puts upward pressure on the stock price. The second term captures the indirect effect through the endogenous response of the bond interest rate to the stock purchases.

The endogenous response of the bond interest rate in the absence of yield curve control is given by

$$\frac{d\ln R^b}{dS^g} = \frac{-v''(B^g)Q}{\beta R^b u'(Y)} > 0, \tag{7}$$

which comes from the fact that an increase in the supply of liquidity puts upward pressure on the bond interest rate. The increase in the bond interest rate, in turn, reduces the demand for stocks, which puts downward pressure on the stock price. This is why the indirect effect is negative. The existing literature on the inelastic stock market, such as Gabaix and Koijen (2021), typically features a perfectly elastic bond market ($v'(\cdot) = 0$), which implies that $\frac{d\ln R^b}{dS^g} = 0$, contrary to what we see in the data.

The above results provide a unifying explanation for our empirical findings. Before yield curve control, the bond interest rate rises through equation (7), and we would not expect a robust response of the stock price because of the countervailing force from the increase in the interest rate. After yield curve control, because the central bank imposes $\frac{d\ln R^b}{dS^g} = 0$, we start to see a robust response of the stock price given by (6). In Appendix E, we show that our model can quantitatively replicate our empirical findings provided that the bond market is sufficiently inelastic (i.e., $v'(\cdot)$ is sufficiently strongly decreasing).

Taken together, our empirical findings, when interpreted through the lens of the model, underscore the importance of taking into account inelasticity in both the stock and the bond market. In the absence of the explicit constraint on interest rate adjustment, the stock price response from an inflow into the stock market might be substantially dampened by the endogenous response of the bond interest rate. Researchers who do not take into account the inelasticity in the bond market may conclude that the stock market is more elastic than it actually is. Put differently, the inelasticity in the stock market can be masked by the presence of an even more inelastic bond market.

6 Conclusion

How does a flow into the stock market impact financial markets? To answer this question, we exploit discontinuities in the Bank of Japan's policy rule. We empirically show that central bank stock purchases have a far-reaching impact on financial markets. Through the lens of the model, we argue that, while these results support the notion that the stock market is inelastic, taking into account an inelastic bond market is crucial to account for our empirical findings.

References

ADACHI, K., K. HIRAKI, AND T. KITAMURA (2021): “Supplementary Paper Series for the "Assessment" (1): The Effects of the Bank of Japan’s ETF Purchases on Risk Premia in the Stock Markets,” Tech. Rep. 21-E-3, Bank of Japan.

ANGRIST, J. D. AND G. W. IMBENS (1995): “Two-Stage Least Squares Estimation of Average Causal Effects in Models with Variable Treatment Intensity,” *Journal of the American Statistical Association*, 90, 431–442.

BARBON, A. AND V. GIANINAZZI (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.

CABALLERO, R. J., T. E. CARAVELLO, AND A. SIMSEK (2024): “Financial Conditions Targeting,” .

CALONICO, S., M. D. CATTANEO, AND R. TITIUNIK (2014): “Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs,” *Econometrica : journal of the Econometric Society*, 82, 2295–2326.

CATTANEO, M. D., M. JANSSON, AND X. MA (2020): “Simple Local Polynomial Density Estimators,” *Journal of the American Statistical Association*, 115, 1449–1455.

CHAROENWONG, B., R. MORCK, AND Y. WIWATTANAKANTANG (2021): “Bank of Japan Equity Purchases: The (Non-) Effects of Extreme Quantitative Easing,” *Review of Finance*, 25, 713–743.

CHODOROW-REICH, G. (2014): “Effects of Unconventional Monetary Policy on Financial Institutions,” .

CHUNG, J.-C. (2020): “The Bank of Japan’s Exchange-Traded Fund Purchase Program and Its Impact on Daily Stock Market Returns,” 53.

DA, Z., B. LARRAIN, C. SIALM, AND J. TESSADA (2018): “Destabilizing Financial Advice: Evidence from Pension Fund Reallocations,” *The Review of Financial Studies*, 31, 3720–3755.

DROSTE, M., Y. GORODNICHENKO, AND W. RAY (2021): “Unbundling Quantitative Easing: Taking a Cue from Treasury Auctions,” *Manuscript, UC Berkeley*.

FUCHS, W., S. FUKUDA, AND D. NEUHANN (2023): "Demand-System Asset Pricing: Theoretical Foundations," *Available at SSRN 4672473*.

FUKUDA, S.-I. AND M. TANAKA (2022): "The Effects of Large-scale Equity Purchases during the Coronavirus Pandemic," 30.

GABAIX, X. AND R. S. J. KOIJEN (2021): "In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis," SSRN Scholarly Paper ID 3686935, Social Science Research Network, Rochester, NY.

HAHN, J., P. TODD, AND W. VAN DER KLAAUW (2001): "Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design," *Econometrica : journal of the Econometric Society*, 69, 201–209.

HARADA, K. AND T. OKIMOTO (2021): "The BOJ's ETF Purchases and Its Effects on Nikkei 225 Stocks," *International Review of Financial Analysis*, 77, 101826.

HARTZMARK, S. M. AND D. H. SOLOMON (2023): "Market Wide Predictable Price Pressure," SSRN Scholarly Paper 3853096, Social Science Research Network, Rochester, NY.

HATTORI, T. AND J. YOSHIDA (2023): "The Impact of Bank of Japan's Exchange-Traded Fund Purchases," *Journal of Financial Stability*, 65, 101102.

ICHIUE, H. (2024): "The Bank of Japan's Stock Holdings and Long-term Returns," *Available at SSRN 4802165*.

ITSKHOKI, O. AND D. MUKHIN (2025): "Mussa puzzle redux," *Econometrica*, 93, 1–39.

KATAGIRI, M., K. TAKAHASHI, AND J. SHINO (2022): "Bank of Japan's ETF purchase program and equity risk premium: a CAPM interpretation," .

KRISHNAMURTHY, A. AND W. LI (2023): "The demand for money, near-money, and treasury bonds," *The Review of Financial Studies*, 36, 2091–2130.

KRISHNAMURTHY, A. AND A. VISSING-JORGENSEN (2011): "The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy," .

——— (2012): "The Aggregate Demand for Treasury Debt," *Journal of Political Economy*, 120, 233–267.

Kuersteiner, G. M., D. C. Phillips, and M. Villamizar-Villegas (2018): "Effective Sterilized Foreign Exchange Intervention? Evidence from a Rule-Based Policy," *Journal of International Economics*, 113, 118–138.

Li, J. J., N. D. Pearson, and Q. Zhang (2021): "Impact of Demand Shocks on the Stock Market: Evidence from Chinese IPOs," SSRN Scholarly Paper 3731219, Social Science Research Network, Rochester, NY.

McCrary, J. (2008): "Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test," *Journal of Econometrics*, 142, 698–714.

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

Parker, J. A., A. Schoar, and Y. Sun (2023): "Retail financial innovation and stock market dynamics: The case of target date funds," *The Journal of Finance*, 78, 2673–2723.

Porter, J. (2003): "Estimation in the Regression Discontinuity Model," *Unpublished Manuscript, Department of Economics, University of Wisconsin at Madison*, 2003, 5–19.

Porter, J. and P. Yu (2015): "Regression Discontinuity Designs with Unknown Discontinuity Points: Testing and Estimation," *Journal of Econometrics*, 189, 132–147.

Selgrad, J. (2023): "Testing the portfolio rebalancing channel of quantitative easing," Tech. rep., Working Paper.

Shirota, T. (2018): "Evaluating the Unconventional Monetary Policy in Stock Markets : A Semi-parametric Approach," 24.