Quantitative Easing and Equity Prices: Evidence from the ETF Program of the Bank of Japan

Andrea Barbon and Virginia Gianinazzi *

This version : July 2019

Abstract

Since the introduction of its Quantitative and Qualitative Easing program in 2013, the Bank of Japan has been increasing its holdings of Japanese equity through large scale purchases of index-linked ETFs, to lower risk premia. We exploit the cross-sectional heterogeneity of the supply shock to identify a positive and persistent impact on stock prices, consistent with a portfolio balance channel. The evidence suggests that long-run demand curves for stocks are downward sloping with unitary price elasticity. We show that the purchases of ETFs tracking the price-weighted Nikkei 225 generate pricing distortions relative to a value-weighted benchmark.

Keywords: ETFs, quantitative easing, portfolio balance channel, unconventional monetary policy, Bank of Japan

JEL classification: E52, E58, B26

*Andrea Barbon (andrea.barbon@usi.ch, 17 Banks St 02138 Cambridge MA, +393408457736) and Virginia Gianinazzi (virginia.gianinazzi@usi.ch) are with USI Lugano and the Swiss Finance Institute. This work was supported by the Swiss Finance Institute. We would like to thank Francesco Franzoni and Alberto Plazzi for their advice and the stimulating discussions. We are also grateful to Miguel Antón, Harjoat S. Bhamra, Andrea Buraschi, John Campbell, Marco Di Maggio, Ruediger Fahlenbrach, Thierry Foucault, Laurent Frésard, Robin Greenwood, Peter Gruber, Jonathan Kearns, Marco LiCalzi, Loriano Mancini, David Miles, Jouchi Nakajima, Jean-Charles Rochet, Suresh Sundaresan, poster session participants at the AFA meeting 2018 and participants at the OFCE/Science-Po Workshop on Empirical Monetary Economics, SFI Research Days, SFI PhD Workshop, Summer School on Market Microstructure and JEN Conference for helpful comments. All errors are our own.

With policy interest rates constrained at the zero lower bound, many central banks around the world have resorted to unconventional monetary policy tools. Within the range of unconventional measures, large-scale asset purchase (LSAP) programs have attracted particular attention because of their large size and thus their impact on central banks' balance sheets. The massive expansion of both the assets and liabilities of central banks exposes them to considerable risks and raises questions about the consequences of a potential exit from QE.

There is considerable evidence that central banks' asset purchases can have an economically significant impact on yields in the targeted markets, which has likely motivated central banks to continue these purchases over the past decade (D'Amico and King, 2013, Eser and Schwaab, 2016, Gagnon et al., 2010, Hamilton and Wu, 2012, Krishnamurthy and Vissing-Jorgensen, 2011, 2013, Neely et al., 2010, Swanson, 2011). However, despite the widespread use of LSAP programs, the debate is still ongoing with regard to the mechanisms linking asset purchases to asset prices and the persistence of the impact. Unlike policy rate targeting, asset purchases are explicit decisions on quantities and are designed to have a noticeable impact on market prices. Even though the idea of *easing through quantity* relies on the view that large purchases by the central bank reduce assets' risk premia, there is still no clear theoretical foundation for how and under which conditions this is expected to work. In general, the relationship between the outstanding quantity of an asset and its price is not yet well understood.

Since 2013, the Bank of Japan (BoJ) has been engaging in what they have named *Quantitative and Qualitative Easing* (QQE) program as an attempt to fight against deflation. As part of its broader QQE agenda, the BoJ has been vigorously increasing its domestic equity holdings through purchases of index-linked ETFs. By the end of 2016, the BoJ owned approximately ¥14 trillion worth of TOPIX and Nikkei ETFs, which corresponds to more than 2.5% of the total market capitalization. This unprecedented equity operation has the declared objective of lowering the risk premia of asset prices and reducing the cost of equity capital of Japanese companies (BoJ, 2013).

The BoJ is the first central bank to engage in purchases of domestic equities as part of its QE agenda. This intervention represents a unique laboratory to shed light on the long-standing debate on the elasticity of long-term demand curves for stocks, as well as a unique opportunity to test how QE impacts equity prices and its implications for market efficiency. In this paper we study its impact on the cross-section of stock prices. We propose a novel empirical strategy to identify and quantify the price impact of the change in assets' supply through QE. This provides new evidence on the price elasticity of long-run demand curves.

The literature on the effectiveness of QE has proposed several channels through which central banks can affect prices. A natural explanation is provided by the so-called "portfolio-balance" channel, first discussed by Brunner and Meltzer (1973), Frankel (1985) and Tobin (1969). According to this channel, when the central bank buys a particular asset, it reduces the

amount held by the private sector, effectively changing the risk composition of the aggregate portfolio held by investors. For this to be an equilibrium, prices need to adjust to ensure market clearing.

In this paper, we first propose a simple structural asset pricing model that generalizes the idea of the portfolio balance channel to the case of equities.[1] The key implication of portfolio rebalancing that we derive from the model is that the change in systematic risk of each stock is determined by: (i) the entire vector of central bank purchases and (ii) the covariance matrix of stocks' cash flows.

We then bring the model to the data in a standard event-study framework, exploiting two specific events in which the BoJ announced major expansions of its ETF purchases. On October 31, 2014, the BoJ announced a three-fold increase in the purchase of ETFs and on July 29, 2016, it communicated a further doubling of the budget amount.

We document that both announcements produced a highly heterogeneous response of equity prices at the company level. Figure 1 plots the cumulative returns of two portfolios following the 2014 announcement, formed by ranking stocks on the price impact predicted by the model. The divergence in returns is statistically and economically significant. Results from cross-sectional regressions show that the variation in event returns in the cross-section is consistent with the change in the marginal contribution of each stock to the risk of the aggregate portfolio held by private investors, as predicted by the portfolio-balance channel.

Looking at longer-horizon returns we find no evidence of reversal over a one-year window after both policy announcements, which supports the main time-series prediction of the model. We estimate the long-term net effect of the portfolio-balance channel at about 22 basis points increase in market value per trillion Yen employed in the program. Given a total equity market capitalization of about ¥500 trillion, this implies an elasticity close to 1, so that each yen invested translates into an increase in total market valuation of roughly one yen. Our estimate is in line with those provded by (Shleifer, 1986a) and (Petajisto, 2011), who find an elasticity of 1 and 0.84, respectivetly, using additions to the S&P 500 index. The analysis based on Dutch auction repurchases of (Bagwell, 1992) also results into a relatively close price elasticity of 1.65 Other authors, instead, find flatter demand curves with price elasticity ranging from the value of 8.24 estimated by (Wurgler and Zhuravskaya, 2002) to that of 10.5 by (Kaul et al., 2000).

The two expansions of the policy budget provide us with an ideal natural experiment to examine the net effect of a long-lasting change in supply on prices for three reasons. First, the purchase schedule of the central bank is exogenous to firms' fundamentals in the cross-

[1] It is easy to show that the duration channel discussed in the literature is a special case of our model when all securities in the economy are exposed to a single source of risk.

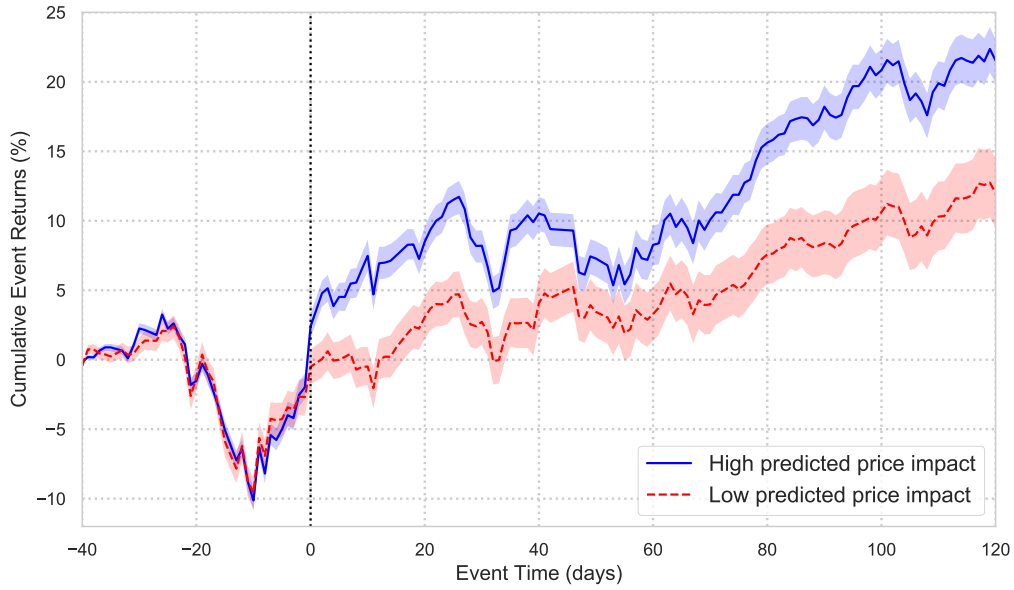


Figure 1: Heterogeneity of Event Returns. This figure shows the time series of the cumulative returns around the BoJ announcement of October 2014, of two portfolios of firms ranked by the model predicted returns. The blue line is the average for first quartile of the distribution (firms with the highest predicted price impact), while the red dashed line corresponds to the average for the last quartile (firms with the lowest predicted price impact). Bands represent bootstrapped 95% confidence intervals.

section. Second, unlike asset purchases by the Federal Reserve, the program of the BoJ affects the supply of each security according to an ex ante well-defined purchase schedule. Third, since roughly half of the capital of the central bank is allocated according to the weights of the price-weighted Nikkei 225 index, the purchases produce variation in the cross-section of supply shocks relative to market capitalization that is as good as random. In general, the identification of the impact of LSAPs on asset prices is a challenging task. The intervention of the BoJ provides us with an empirical framework that mitigates endogeneity concerns and improves the identification of the net (short-run and long-run) effect of a change in supply, which crucially relies on the exogeneity of the shock.

The non-fundamental nature of the Nikkei weighting system has already been exploited in Greenwood (2005, 2007) to establish a causal relationship between uninformed demand shocks and prices in the context of a large redefinition of the Nikkei 225 index. A major difference between the LSAP setting and the one of index redefinitions lies in the nature of the supply shocks. As the central bank buys assets, it is effectively transferring a portion of fundamental risk from the private sector to its balance sheet, and holds it for an arguably long period of time. This is at least conceptually different from an index redefinition event, in which securities merely change hands from active investors to index funds. The central bank can

be thought of as a buy-and-hold long-term investor whose portfolio holdings are not marked-to-market. Its long-term commitment to the policy induces a long-lasting change to supply, making our setting better suited than index redefinitions to identify long-run price effects due to movements along investors' long-term demand curves.

The model that we propose extends the theoretical framework in Greenwood (2005) to account for this difference in setting. As in Greenwood (2005), we consider an economy with multiple assets in finite supply and a CARA-utility representative agent that maximizes her wealth in each period. We introduce quantitative easing in the form of an exogenous shock to the supply of assets, which is first announced and then gradually carried out over a given policy horizon. The agent correctly understands that the QE program will affect the market-clearing portfolio in each future period, which determines the new vector of equilibrium risk premia. Crucially, we extend the model to an infinite horizon to relax the assumption that uncertainty is resolved at a terminal date, which mechanically drives the reversal in Greenwood (2005). In our model, prices adjust to the change in supply to reflect the new risk composition of the aggregate portfolio held by the representative agent. Unless the central bank is expected to unwind its position, this implies that we should not observe a reversal at any horizon. The fact that we observe a persistent effect in the data is consistent with this prediction of the model.

In the data, not only we find no evidence of a reversal of the initial jump in prices, but non-trivial abnormal returns are still observed one year after the announcements. Even though the policy is carried out gradually, the total size of the intervention is revealed to the market in advance. Market efficiency requires that today's prices reflect expectations about future returns, hence they should also reflect expected future changes in the supply of assets. In the model, the observed post-event drift can arise because of two reasons. First, since purchases are scattered over various dates, the model predicts that prices continue to adjust also after the announcement due to the decrease in the residual duration of the program over time. However, for realistic levels of the risk free rate, we show that this effect is expected to be quantitatively small relative to the initial price jumps. A more pronounced drift arises when we allow the representative agent to believe that the central bank will deviate from the announced purchase target. The model produces a sluggish price reaction similar to the one observed in the data when expectations about the size of the purchase program are assumed to increase over time, consistently with investors underreacting to the announcement as well as with learning about additional purchase programs in the future.[2]

We address the concern that (part of) the observed price impact and its persistence might be explained by repeated price pressure rather than a portfolio-balance channel. Large trades

[2]Beliefs are exogenous in our model and evolve deterministically over time. Extending the model to a setting where beliefs are endogenous is definitely interesting, but beyond the scope of the paper.

from the BoJ may give rise to order imbalances, thus pushing prices upwards on purchase days. Such mispricings are expected to be shortly lived in efficient markets. However, arbitrageurs may refrain from trading if the central bank is expected to buy again soon, thus failing to bring prices back to their fundamental value. This would imply a persistent price effect of the program arising from the flow of the purchases rather than the change in the supply of assets, a channel quite distinct from portfolio balance. As in D'Amico and King (2013), we will refer to this effect as the *flow effect* of the program. Even though the difference may seem subtle, disentangling between these two channels has important implications. First, the two channels lead to different conclusions about the elasticity of long-run demand curves for stocks. Second, they imply different consequences of a potential exit from QE. In particular, if QE is mainly effective through repeated price pressure, a slow-down or a suspension of the purchases would cause a sharp drop in prices. On the contrary, in our model of the portfolio-balance channel, it is not the flow into the balance sheet of the central bank that keeps prices up, but its accumulated size. Therefore, suspending the purchases should have a more limited effect on prices. We exploit both the cross-sectional and time series variation in purchase volumes to identify and quantify the flow effect of the policy in the spirit of Eser and Schwaab (2016). We then re-estimate the cross-sectional portfolio-balance channel effect using returns net of the flow-induced component. We find that price pressure effects are positive and persistent. However, this channel might explain at most a minor fraction (between 12% and 23% depending on the specification) of the estimated portfolio balance effect.

Overall, our empirical analysis confirms the concerns raised by the financial press that the intervention of the BoJ might be inducing price distortions due to the deviation of the purchase schedule from market weights. We document a significantly heterogeneous effect of the policy both at company and industry level. A modification of the QQE has the potential to address this problem. Theoretically, the only way to achieve a cross-sectionally homogeneous shift in risk premia is for the BoJ to hold each stock proportionally to the company market capitalization. At the moment, however, still roughly a quarter of the BoJ capital is allocated to the price-weighted Nikkei index.

The rest of the paper is organized as follows. Section 1 describes the ETF purchase program of the BoJ. Section 2 reviews the relevant literature. Section 3 presents the model. Section 4 describes the data, the empirical strategy and estimation procedures. Section 5 presents our main empirical findings. Section 6 considers the flow effect of direct purchases and evaluates its relative importance with respect to the portfolio balance effect. Section 7 discusses policy implications and Section 8 concludes.

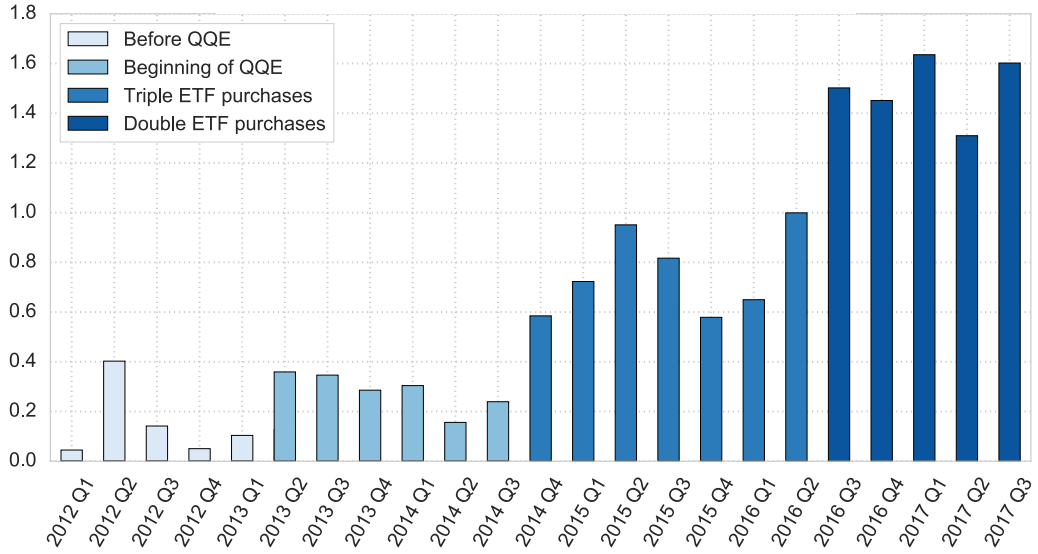


Figure 2: Quarterly ETF Purchases of the Bank of Japan in trillion yen. Changes in the bar color indicate changes in the policy target purchase amounts. In the first phase the target was set to ¥1 trillion per year, in the second phase it was tripled to ¥3 trillion and in the third phase it was additionally doubled to ¥6 trillion. Data is from the BoJ website.

1 The ETF Program of the BoJ

As part of the "Quantitative and Qualitative monetary Easing" (QQE) introduced on April 4, 2013, the BoJ embarked on a large-scale asset purchase (LSAP) program committing itself to buy large quantities of broad market equity ETFs, with the declared view of lowering risk premia (BoJ, 2013). The policy budget was initially set at ¥1 trillion per year (roughly US$ 10 billion). On two occasions, the BoJ announced a sharp expansion of the target amount: on October 31, 2014, the Bank communicated that the annual mark was tripled to ¥3 trillion, and was again doubled on July 29, 2016, to ¥6 trillion. The policy changes are clearly visible along the time series of monthly ETF purchases by the bank, as shown in Figure 2. The time series of aggregate ETF purchases is publicly available at daily frequency on the BoJ website starting from December 2010.

Its holdings accumulated rapidly, and by the end of 2016, the BoJ owned more than ¥14 trillion worth of ETFs. This corresponds to 2.5% of the total capitalization of the First Section of the Tokyo Stock Exchange (TSE), and around 3% of the Japanese GDP. The share of BoJ holdings to aggregate Assets Under Management (AUM) of targeted ETFs has grown from almost zero to more than 70% since the beginning of the program; this is even more remarkable if we consider that the ETF industry in Japan almost tripled in value between 2013 and 2016. In terms of size, the ETF program is comparable to the annual aggregate net flows

into or out of the Japanese equity fund industry and therefore economically relevant.[3]

The purchase program targets two types of ETFs: those tracking the Tokyo Stock Price Index (TOPIX) and those replicating the return of the Nikkei 225 Stock Average.[4] At inception of the program, the money allocated to each ETF was set to be proportional to its assets under management (AUM). The ratio of the aggregate AUM of ETFs tracking the TOPIX Index and those of ETFs tracking the Nikkei 225 Index is roughly 1 to 1.2. This approximately translates into half of the capital flowing into Nikkei ETFs and half into TOPIX ETFs. In turn, this then maps into a demand shock at the stock level that depends on each company's weight in the corresponding index.

The TOPIX is a value-weighted index tracking the roughly 2000 companies listed on the First Section of the TSE, while the Nikkei 225 is a *price-weighted* index of 225 TOPIX companies representative of the Japanese stock market. The constituents of the Nikkei index are typically large blue-chip companies that account for roughly two-thirds of the market capitalization of the TSE First Section on aggregate. The Nikkei 225 is the most widely traded equity benchmark in Japan.

The weighting system of the two indices implies that the BoJ allocates only half of its budget to companies proportionally to their market value. The remaining half of the budget flows instead to the Nikkei constituents proportionally to their price, not accounting for the number of shares outstanding, thus producing mis-allocation relative to market capitalization. Under market efficiency, the market value of a company should reflect all available fundamental information. The dispersion of the ratio between price weights and value weights is therefore expected to be unrelated to firms fundamentals. The relative under-weighting in the BoJ portfolio is clearly more severe for companies not included in the Nikkei index. However, there is a high degree of heterogeneity in the allocation of capital across Nikkei companies as well. This is clear from Figure A.1 in Appendix A, where we plot the distribution of the log of the ratio between the weight in the Nikkei and the weight in the TOPIX for Nikkei companies to measure the cross-sectional dispersion of the resulting allocation at the stock level.

Given the unusual weights of the BoJ purchase schedule, a sudden expansion of the policy budget produces a natural experiment where stocks are hit by an uninformed demand shock that is highly heterogeneous in the cross-section and orthogonal to firms fundamentals after

[3]Over the past 10 years, the average net flows into equity funds in Japan was roughly ¥3 trillion in absolute value per year. Data are from the Thomson Reuters Lipper Global Fund Flows database.

[4]On November 19, 2014, the BoJ started buying also ETFs tracking the JPX-Nikkei 400 Index. This approximately corresponds to 43% of the purchases flowing to ETFs tracking the TOPIX, 53% to ETFs tracking the Nikkei 225 and the remaining 4% to ETFs tracking the JPX Nikkei 400. For simplicity, in the empirical analysis we round the share of both TOPIX and Nikkei ETFs to 50%, neglecting the JPX Nikkei 400. This simplification does not affect the results of our analysis.

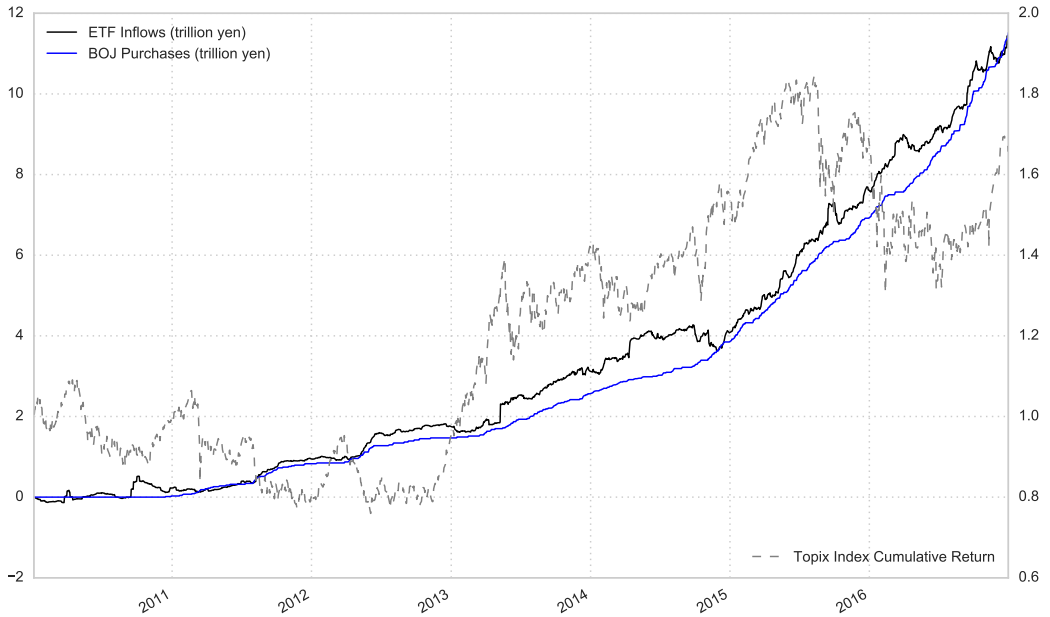


Figure 3: BoJ Purchases and ETF inflows. On the left axis we plot the daily cumulative purchases of ETFs by the BoJ (blue line) and the estimated daily cumulative inflows into ETFs tracking either the Nikkei or the TOPIX index (black line). Both are in trillion yen. On the right axis the figure shows the cumulative return from 2010 of the TOPIX index (gray dashed line).

controlling for market capitalization. In this paper, we exploit the exogenous variation in the cross-section of supply shocks to identify the causal impact of the purchase program on equity prices. We rely on a simple asset pricing model to argue that the deviation from a value-weighted allocation allows us to isolate the portfolio-balance channel of the policy impact. Section 4.2 discusses the identification strategy in detail.

Overall, the portfolio of the BoJ ends up deviating significantly from the allocation that market capitalization would dictate. To illustrate the extent of this distortion, take three companies with fairly similar market capitalization and therefore similar TOPIX weights (between 0.45% and 1% in 2014): Canon, Fast Retailing and Nintendo. Canon and Fast Retailing are both among the Nikkei constituents, though with very different weights, namely around 1.2% versus 9.5%, respectively. Nintendo, on the contrary, is not included in the Nikkei index. It follows that the BoJ allocates to Fast Retailing 4 times more capital than to Canon, and 19 times more than to Nintendo. The effects of the departure from a value-weighted allocation are reflected in the indirect ownership that the BoJ accumulated over time. According to estimates by the Financial Times, through its purchases the central bank has indirectly become the largest shareholder in a quarter of TOPIX stocks. In Table 1 we report the ten stocks with the highest estimated indirect ownership share by the BoJ.

We argue that purchases of ETFs by the BoJ translate into supply shocks at the individual

stock level. This is a consequence of the creation-redemption mechanism in the ETF primary market and the physical replication of the underlying basket. When demand exceeds supply in the ETF secondary market, new shares of ETF are issued to keep the ETF price close to its NAV. In the case of physical ETFs, creation requires the physical purchase of the basket of securities that composes the tracked index, for a value equal to the creation unit. Securities are then held by the ETF sponsor on behalf of the owner of the ETF shares, who now bears the associated risk. ETF creation thus reduces the quantity of assets available for trading in the underlying market. This mechanism is visualized in Figure A.2 in Appendix A. Given the direct correspondence, in the rest of the paper we will consider ETF purchases by the central bank equivalent to an intervention in the underlying equity market.

We can infer whether central bank purchases triggered creation of new ETF shares from data on ETFs AUM. We first obtain the list of the ETFs listed on the Tokyo Stock Exchange (TSE) that track either the Nikkei or the TOPIX index from the website of the Japan Exchange Group (JPX). We then get daily data on AUM for each ETF from Bloomberg. We estimate inflows simply as the difference between the actual increase in AUM and the increase in AUM due to the return on the index that the ETF is tracking. Figure 3 plots the time-series of ETF inflows versus the amount purchased by the BoJ. It is apparent that the flows into these ETFs are almost completely due to the asset purchase program. In turn, this implies that the purchases by the BoJ have consistently triggered creation of new ETF shares.

It must be noted that the bias towards Nikkei companies did not go unnoticed among practitioners and the BoJ was frequently accused by the financial press of distorting the market. In response to the criticism, on September 21, 2016, the BoJ amended the terms and conditions of the program and announced it will change the maximum amount of each ETF to be purchased. Since October 2016, the BoJ allocates ¥2.7 trillion a year (US$ 26.4 billion) to TOPIX ETFs, while the remaining ¥3 trillion are spread out between the TOPIX, the Nikkei 225 and the JPX-Nikkei Index 400. For the Nikkei-ETFs this means a drop from 55% to about 25% of the annual purchases by the BoJ, which brings the allocation of the flows closer to what market capitalization would justify. Yet, the accumulated balance sheet of the BoJ remains tilted away from a value-weighted allocation.

2 Related Literature

"Extraordinary times call for extraordinary measures", stated the Chairman of the Federal Reserve Ben Bernanke in 2009 (Bernanke, 2009). Since then, a number of central banks around the world have adopted unconventional monetary policy tools and most of them have been trying to support asset prices through LSAPs in order to boost economic activity in the face of severe dislocations in financial markets. With actual data on the implementation of

Company Name	BOJ Share (%)	BOJ Flow (bn JPY)	Market Cap (bn JPY)	Nikkei weight (%)
Mitsumi Electric Co Ltd	10.3	5.8	56.1	0.17
Advantest	8.9	27.5	309.1	0.63
Fast Retailing	8.7	336.4	3854.7	9.17
Taiyo Yuden	7.8	10.0	129.0	0.32
Toho Zinc	7.7	3.3	43.2	0.09
Tdk Corporation	7.4	71.0	959.0	1.41
Konami Holding	7.2	37.6	524.5	0.65
Trend Micro	7.0	36.2	514.9	0.90
Comsys Holding	6.6	18.3	275.7	0.39
Nissan Chem In	6.2	30.6	489.7	0.53
Average	6.1	3.8	255.6	0.44
Median	5.9	0.2	43.8	0.20

Table 1: BoJ indirect shareholdings. Summary statistics on indirect ownership by the BoJ for the ten companies with the highest BoJ share. BoJ Flow are the cumulative compounded BoJ purchases at company level since the beginning of QQE and Market Cap is the company's market capitalization. BoJ Share is the ratio of BoJ Flow and Market Cap. Average and median values are calculated over the universe of TOPIX firms. The values in the first three columns are as of August 31, 2016. The last column reports the average company weight in the Nikkei 225 index over the study period. Notice that the ten companies with the highest BoJ share have all positive weights in the Nikkei 225 index.

LSAPs becoming available, a large body of academic research has investigated their impact on financial markets and the real economy.

Most of the work on the impact of QE on market prices relies on evidence from purchases of government bonds by the Fed, the ECB or the BOE, and usually shows a significant impact on yields (Buraschi and Whelan, 2015, D'Amico and King, 2013, Eser and Schwaab, 2016, Gagnon et al., 2010, Hamilton and Wu, 2012, Joyce et al., 2012, Krishnamurthy and Vissing-Jorgensen, 2011, 2013, Neely et al., 2010, Swanson, 2011). There is however little empirical evidence on the large-scale purchases of the BoJ. Perhaps closest to our paper is Ueda (2013), who looks at the time series of LSAP announcements by the BoJ and finds a positive correlation with the TOPIX index and the yen-dollar exchange rate. The BoJ is the first central bank to purchase domestic equities as part of its QQE agenda, and, to the best of our knowledge, this paper is the first to study this program in depth and to analyze its impact on the cross-section of stock prices.

Although there is general agreement that LSAPs do indeed affect prices, there is less consensus regarding the channels through which these policies work. A standard explanation in the

literature is the so-called portfolio balance channel (Brunner and Meltzer, 1973, Frankel, 1985, Tobin, 1969). According to this channel, when the central bank buys a particular asset, it reduces the amount held by private investors, effectively forcing them into a different portfolio. For this to be an equilibrium, prices need to adjust to ensure market clearing. In particular, through this channel asset purchases are expected to push up the price of the target asset and of its substitutes, implying that demand curves are downward sloping. Some papers find that the observed price impact is consistent or partially consistent with portfolio balance explanations (e.g. D'Amico and King (2013), Gagnon et al. (2010), Joyce et al. (2011))[5]. However, the portfolio balance channel of monetary policy is subject of debate, in part because standard asset pricing models do not generally allow exogenous changes in the supply of a security to affect its price. For instance, Miles and Schanz (2014) argue that LSAPs by central banks since 2008 had significant effects because markets were dysfunctional and that in normal times portfolio-balance effects would be weak.

The question whether demand curves for stocks slope down has a long tradition in the asset pricing literature. The empirical evidence so far mostly comes from event studies around index redefinitions and fire sales by institutional investors (Coval and Stafford, 2007, Greenwood, 2005, Harris and Gurel, 1986, Hau et al., 2009, Mitchell et al., 2004, Petajisto, 2009, Scholes, 1972, Shleifer, 1986b, Schnitzler, 2016). The general finding is that large non-fundamental trades have a significant but temporary price impact, even though there is considerably heterogeneous evidence on the speed and the extent of reversal. The standard interpretation is that limits to arbitrage can justify temporary deviations from fundamental value: under market efficiency, uninformed shocks cannot have a long-lasting impact on prices.[6]

Quantitative easing provides an ideal laboratory in which to test asset pricing theories such as the long-held belief of flat demand curves for stocks. However, from the success of QE in pushing up prices alone, one cannot conclude much about the elasticity of demand curves. A large number of papers show that LSAPs by central banks have effects beyond those due to portfolio balance, and provide evidence of alternative transmission channels that are consistent with flat demand curves. For the case of purchases of long-term bonds, Krishnamurthy and Vissing-Jorgensen (2011) provide compelling empirical evidence that the so-called signalling channel explains a significant fraction of the drop in bond yields observed after the Federal Reserve's QE announcements. The idea behind this channel is discussed in Eggertsson and

[5]Vayanos and Vila (2009) try to reconcile the predictions of the portfolio balance channel with the observed lack of spillovers across maturities, building on market segmentation and preferred-habitat theories as proposed by Culbertson (1957) and Modigliani and Sutch (1966).

[6]The traditional view in finance is that, in a frictionless world, a simple expansion of the balance sheet of the central bank should have no effect. This neutrality result is formalized in Eggertsson and Woodford (2004) and crucially relies on the assumption of a rational infinitely lived agent with no credit restrictions, who sees no difference between its own assets and those held by the central bank.

Woodford (2004), who claim that financial markets may interpret LSAPs as signals about the central bank's intention to keep interest rates low, thus influencing long-term yields through investors' expectations about the future path of interest rates. Other papers attribute the beneficial effect of the Fed's MBS purchases on risk premia during the financial crises to a capital constraints channel motivated by the distress in the financial intermediary sector (Curdia and Woodford, 2011, He and Krishnamurthy, 2013).

In general, the identification of the impact of market interventions through a specific channel is a challenging task. Our paper contributes to this literature proposing a new identification strategy for the transmission channel of monetary policy and providing new insights on the elasticity of demand curves. Moreover, the results of the empirical literature suggest that the specific workings of LSAPs depend on the asset purchased and the economic conditions under which these purchases take place. We complement the existing evidence by documenting the effects of the ETF program by the BoJ, a unique case in which a central bank is targeting the equity market.

3 The Model

In this section we develop a theoretical framework to describe the portfolio balance channel as the transmission mechanism from LSAP to asset prices. The idea is that asset purchases shift part of the fundamental risk from the market to the balance sheet of the central bank. Because the premium demanded for a given security is proportional to its marginal risk contribution to the aggregate portfolio held by the representative agent, the price effect of the monetary intervention is proportional to the implied change in this quantity. Therefore, the net effect on asset prices through this channel is not simply proportional to the purchased amounts, but it crucially depends on the correlation structure of firms fundamentals.

Our model features the central bank only in reduced form, in the sense that the policy rule is exogenous. We also assume that asset purchases are deterministic. This assumption holds also when we allow investors to believe that the central bank will deviate from the announced purchase target. With no policy uncertainty, asset purchases do not represent a source of risk that has to be priced in equilibrium. Moreover, we assume firms fundamentals to be neutral with respect to monetary policy, excluding the possibility that asset purchases affect market prices through the change in future investment opportunities. We make these choices to keep the model simple and to focus on the direct effect of supply on prices. These assumptions also allow us to restrict our attention to the covariance-stationary equilibrium of the model, which immediately follows once we assume covariance-stationary dividends. The limitation is that the model abstracts from potential additional channels related to uncertainty about future supply and endogenous responses of firms.

3.1 Model Setup

Consider an economy with n risky assets in fixed supply $Q = (Q^1, \ldots, Q^n)$, paying dividends in every time period. The dividend $D_{i,t}$ paid at time t is

$$D_{i,t} = D_{i,0} + \sum_{s=1}^{t} \varepsilon_{i,s}, \qquad \forall i \in 1, \ldots, n \tag{1}$$

where each $\varepsilon_{i,t}$ is revealed at time t. The fundamental innovations $\varepsilon_{i,t}$ are modelled as zero-mean jointly normal random variables, iid over time.

The representative agent optimally chooses her time-t demand N_t to maximize her next period expected utility, subject to a standard budget constraint

$$\max_{N} \ \mathrm{E}_t \left(-\exp(-\gamma W_{t+1}) \right) \tag{2}$$

$$\text{s.t.} \qquad W_{t+1} = W_t(1+r) + N_t'(p_{t+1} + D_{t+1} - p_t(1+r)) \tag{3}$$

where W_t is the total wealth, N_t' denotes the transpose of the vector N_t and γ the aggregate risk-aversion. At date $t = 1$ the central bank announces share purchases described by the vector $q = (q^1, \ldots, q^n)$, distributed over M periods after the announcement. We refer to M as the policy horizon. Let q_t denote the vector of cumulative purchases by the central bank up to date t. One can think of q_t as the active side of the balance sheet of the central bank at any time t.

We assume, first, that $q_t = tq$ for $t = 1, \ldots, M$ and, second, that $q_t = Mq$ for $t > M$. The first assumption implies that in our model the central bank's balance sheet evolves deterministically and grows linearly over time. Assuming non-stochastic asset purchases allows us to abstract from policy uncertainty as a priced risk factor and to focus on how QE affects prices through the change in supply [7]. The second assumption implies that the central bank never unwinds its position nor engages in further purchases beyond horizon M. This assumption might be restrictive once we go to the data since the BoJ never announced such a stringent commitment. Still, given that the BoJ position have not been unwound (and neither announced to be so) over the window of our empirical analysis, we believe it to be a reasonable benchmark.

The realized demand shocks negatively affect the net supply of assets in each period. Setting $Q_0 = Q$ yields

$$Q_t = Q - q_t \tag{4}$$

[7] The assumption that the central bank spreads its purchases equally over the policy horizon is instead innocuous. Relaxing this assumption does not improve the economic intuition and only adds technical complexity to the model.

Asset purchases by the central bank affect the quantity at which the equity market clears given the equilibrium condition $N_t = Q_t$. Notice that equation (4) also implies that the quantity of assets available to the market can only change through purchases of the central bank. This excludes the possibility for companies to respond endogenously to changes in prices by issuing new stocks or buying back those outstanding.

The central bank buys the vector q of securities in exchange for cash. We assume that the representative agent invests the proceeds in the risk-free asset and, since risk-free returns are uncorrelated with those of Japanese equities, omitting the risk-free asset from the model does not change the predicted policy impact on stock prices. This assumption may be interpreted as a form of market segmentation, in that the representative agent cannot re-invest the proceeds in assets outside the Japanese equity universe. We discuss the implications of this assumption in Section 5.4. The fact that we do not model other asset classes that equity investors might hold in their portfolios does not affect the model predictions even in case of non-zero correlation with equities. It is easy to show that including securities that are not targeted by the asset purchase program has no effect on the predicted price impact on stock prices. Stock purchases will spillover to correlated asset classes, but in this paper we are not interested in these effects.

Appendix B shows that the pricing equation of the covariance-stationary equilibrium, in matrix notation, is given by

$$p_t = \frac{1}{r}\left(D_t - \gamma V \Omega_t\right) \tag{5}$$

where $V \equiv \mathrm{Var}_t(p_{t+1} + D_{t+1})$ is the stationary covariance matrix of asset returns and Ω_t is the vector of time-t expected future asset supply, properly discounted by time, defined as

$$\Omega_t \equiv \frac{r}{1+r} \sum_{i=0}^{\infty} \frac{\mathrm{E}_t[Q_{t+i}]}{(1+r)^i} \tag{6}$$

Notice that the covariance matrix V is not time-varying because the supplies of each asset are fixed and the schedule of purchases by the central bank is deterministic, so there is no uncertainty on future shocks to the asset supply. Equation (6) shows that at any point in time prices reflect the path of future asset supply. Given the time-discounting, today's prices are less sensitive to quantities further into the future.

Staring at the vector of risk premia $\gamma V \Omega_t$ in the pricing equation (5) one can see that, for each stock, priced risk is an increasing function of the stock's covariance with the market portfolio and the risk aversion parameter γ. The vector $V\Omega_t$ admits an interpretation very similar to the CAPM beta and should be thought of as a measure of systematic risk.[8] This is easier to

[8] While the model is written in price changes, market betas are usually defined in terms of returns. In Appendix C we derive an expression of systematic risk that determines expected returns in the model. While the notation becomes messier, the intuition carries through.

see in the absence of monetary policy shocks, in which case $V\Omega_t$ reduces to VQ.

By plugging equation (4) into equation (5) it is easy to see how the portfolio balance mechanism works in the model. Asset purchases change the amount of each security in the market clearing portfolio. This affects systematic risk and in turn prices. Notice that this change in systematic risk is fully consistent with our assumption of a constant covariance matrix V, since what determines systematic risk is the product $V\Omega_t$, and central bank purchases affect only the latter term in the model.

Let's now turn to how the representative agent builds expectations about future asset supply, before and after the purchase program is announced. These expectations enter the pricing equation (5) and so determine the impact of the policy. At time $t > 0$, the representative agent's expectation about the quantity in period $h = t + 1, \ldots, M$ is given by

$$E_t[Q_h] = E_t[Q_0 - q_h] = Q - E_t[q_h] \tag{7}$$

We assume that the central bank intervention is fully unexpected at $t = 0$ before the announcement, i.e. $E_t[q_h] = 0$ for every $t \leq 0$ and $h \geq t$. At each period $t \geq 1$ after the announcement date, we allow the investor to believe that the central bank will deviate from its purchase target. We restrict to a family of investors beliefs parametrized by a time-varying scalar λ_t. Formally, let $\lambda_t \geq 0$ be real numbers such that

$$\mathrm{E}_t\left[Q_h\right] = Q - \lambda_t q_h, \qquad t \geq 1 \tag{8}$$

The parameter λ_t is assumed to change over time in a deterministic fashion. The path of λ_t determines how the representative agent updates her beliefs regarding the size of the purchase program. We first solve the model for a general mapping $t \mapsto \lambda_t$ and then we present, in the next section, results for the special case $\lambda_t \equiv 1$, in which the pricing equation takes a simpler form that better conveys the intuition for the portfolio balance channel. Imposing $\lambda_t \equiv 1$ is equivalent to assuming that the representative agent expects the central bank to commit to the announced target exactly. This implies that she also expects the central bank to never unwind its positions and to never engage in additional purchase programs in the future. Changing the parameter λ_t allows us to study how deviations from this benchmark case impact the effect of the policy.

As shown in Appendix B, the model's pricing equation predicts price changes given by

$$p_t - p_{t-1} = \frac{1}{r}\left(\varepsilon_t + \gamma\xi(t)Vq\right) \tag{9}$$

where the function $\xi(t)$ is defined piece-wise as follows

$$\xi(t) = \begin{cases} 0 & \text{if} \quad t \leq 0 \ \text{ or } \ t > M \\ \lambda_1(M - \varphi(1)) & \text{if} \quad t = 1 \\ \Delta\lambda_t M - (\lambda_t \varphi(t) - \lambda_{t-1}\varphi(t-1)) & \text{if} \quad 1 < t \leq M \end{cases} \tag{10}$$

and $\varphi(t) < M$, defined in Appendix B, is a deterministic function of time representing the residual duration of the purchase program.

In the first part of equation (10), $\xi(t) = 0$ implies that both before the announcement ($t \leq 0$) and after the purchase program has been completely carried out ($t > M$), price changes only reflect shocks to dividends and are therefore unpredictable. The functional form of $\xi(t)$ in the second and third pieces of the domain determine event ($t = 1$) and post-event price changes ($1 < t \leq M$), respectively. Even when future supply changes are fully predictable and the average path of future prices can be perfectly anticipated ($\lambda_t \equiv 1$), the shock to supply is impounded into stock prices immediately after the policy announcement only up to the term $\varphi(1)$. Prices will then continue to adjust in the following days. The reason why prices do not fully adjust on the event day is that until purchases are actually realized at future dates, the representative agent bears dividend risk and requires a compensation for it. Consistent with this intuition, $\varphi(t)$ is decreasing in t and increasing in M. So, even though $\xi(t)$ drives predictable post-event price changes, these are fully consistent with market efficiency and do not represent an arbitrage opportunity.

The relative magnitude of the initial price reaction and the subsequent adjustments depend on the level and the dynamics of λ_t. More specifically, the price jump at $t = 1$ is increasing in the initial expectation of future supply λ_1 since prices are effectively responding to a purchase program of size $\lambda_1 Mq$. Post-event price changes are then linked to the time series evolution of λ_t. An increasing λ_t over time means that the agent is revising upward her expectations about the size of the program. One can think of different reasons for why this might happen. For example, the agent may not immediately believe that the central bank will commit to the full size of the program and thus update her expectations only once she observes the purchases actually being carried out. Or, she may start believing over time that the central bank will engage in additional purchases beyond the announced policy horizon M. Similarly, a decreasing λ_t means that the agent revises downward the expected size of the program, either because she starts to believe that the central bank will not complete the announced program or that it will unwind the portfolio soon after. In the Internet Appendix we show simulations of the price dynamics implied by different functional forms for λ_t.

Even though we mainly think of λ_t as controlling the agent's beliefs on the central bank actions conditional on time-t information, this reduced-form suits a number of non-mutually exclusive interpretations. For instance, as in Barberis and Thaler (2003), the slow reaction

may be due to the bounded rationality of agents who fail to correctly process the consequences of the BoJ announced program.

3.1.1 Benchmark Case

In this section we focus on the special case where $\lambda_t \equiv 1$, which implies that expected and realized purchases are the same at any point in time

$$E_t[q_h] = q_h = h\, q \tag{11}$$

It follows directly from equation (9) that the price adjustment at $t = 1$ is given by

$$p_1 - p_0 = \frac{1}{r}\left(\varepsilon_1 + \gamma V(Mq - \varphi(1)q)\right) \tag{12}$$

Ignoring fundamental innovations, equation (12) predicts a positive price jump of magnitude $\gamma V(Mq - \varphi(1)q)$. This swing in prices is due to the fact that the policy is unexpected at $t = 0$, but it is impounded into prices as soon as it is revealed.

In the following periods ($t \geq 1$), price changes are instead given by

$$p_{t+1} - p_t = \frac{1}{r}\left(\varepsilon_{t+1} - \gamma V(\varphi(t+1) - \varphi(t))q\right), \qquad t = 1, \dots M \tag{13}$$

Equation (13) shows that price changes in the post-announcement period include a non-stochastic component $\gamma V(\varphi(t+1) - \varphi(t))q$ which accounts for the time delay between the announcement of the supply shocks and their realizations. Since the cross-sectional distribution of these predictable price adjustments is always parallel to and of the same sign as the initial price impact, they add up to create a propagation (drift) of the initial cross-sectional effect.

3.2 Testable Predictions

In this section we derive testable predictions from the model. To make these predictions more suitable to be tested in the data, we state them in terms of returns. In order to go from the expressions in price changes derived in Section 3.1 to predictions about returns, we first need to introduce some new notation. We define u as the vector of yen amount purchased by the BoJ of each security, so that

$$u_i \equiv p_{i,t}\, q_i, \qquad \forall i \in 1 \dots, n \tag{14}$$

where q_i is the number of shares purchased of stock i and $p_{i,t}$ the stock price at time t. Then, we define Σ to be the stationary covariance matrix of stock returns, i.e.

$$\Sigma_{i,j} \equiv \text{Cov}(R_i, R_j), \qquad \forall i, j \in 1 \ldots, n \tag{15}$$

where R_i is the daily percentage return of stock i.

Dividing equation (5) by p_{t-1} leads to the following two propositions about event returns and post-event returns. Proofs are in Appendix B.

Proposition 1 (Event returns)**.** *The vector of returns $R_1 = (p_1 - p_0)/p_0$ on the announcement day is positively related to the vector $\pi \equiv \Sigma u$ in the cross-section.*

Proposition 2 (Post-event returns)**.** *Assume $\Delta\lambda_{t+1} = \lambda_{t+1} - \lambda_t \geq 0$. Then the vector of post-event returns R_{t+1} is positively related to $\pi = \Sigma u$ in the cross-section for every $t = 1, \ldots, M$. Moreover, the vector of expected cumulative returns is given by*

$$\sum_{s=1}^{t} \text{E}[R_s] = \theta_t \pi \tag{16}$$

where $\theta_t = \frac{\gamma}{r} \sum_{s=1}^{t} \xi(s)$ is a positive and increasing function of t, which follows from the definition of $\xi(t)$ in equation (10).

Proposition 1 states that through the portfolio balance channel, the policy announcement leads to abnormal event returns proportional to the change in systematic risk captured by the vector $\pi = \Sigma u$. Notice that if the central bank were to buy stocks proportionally to their market weight, abnormal event returns would be proportional to the product of Σ and the vector of market capitalizations, i.e. the vector of each stock's covariance with the market portfolio. Proposition 1 therefore implies that an exogenous shock to supply parallel to the market portfolio would cause price adjustments proportional to market betas. At the same time, it also implies that shocks to supply that are orthogonal to market capitalization produce abnormal returns orthogonal to market betas. This prediction is key to identify the effect of the policy shock from the cross-section of realized event returns in the empirical part of the paper.

As summarized in Proposition 2, the model predicts post-event returns in the same direction of event returns, i.e. proportional to π, until the purchase target is met at $t = M$. This generates a post-event drift, whose magnitude depends both on the value of the risk free rate and the beliefs dynamics parametrized by λ_t. In the Internet Appendix, we show analytically and from model simulations that for realistic value of the risk free rate and λ_t constant, this drift is small. The model produces a more pronounced drift under the assumption that the representative agent revises her expectations on the size of the program over time ($\Delta\lambda_{t+1} > 0$).

From Proposition 2 it follows that a permanent change in the supply of assets generates a permanent change of risk premia, and hence of prices. Unless the central bank unwinds its positions, prices will not revert to the pre-event level.[9] By stating that changes in supply can have long-lasting impacts on prices, the proposition implies downward sloping demand curves for stocks through the portfolio balance mechanism.

4 Data and Empirical Methodology

4.1 Data Sources

From Compustat Global we collect stock-level data on daily returns, volumes and shares outstanding for the roughly 2000 stocks of the TOPIX universe for the period 1990-2016. Daily returns and volume data for the TOPIX index as well as the monthly time-series of TOPIX and Nikkei 225 index weights for every stock in our sample are obtained from Thomson Reuters Datastream. The USD/JPY exchange rate is from Japan Macro Advisors Inc. The time-series of ETF purchases by the BoJ is publicly available at daily frequency on its website.

4.2 Identification Strategy

To test the model predictions from Proposition 1 and Proposition 2 we estimate the following cross-sectional regression at different horizons H around the two policy announcements made by the BoJ

$$R^H_{i,e} = \alpha_e + \beta^H_e \pi_{i,e} + \delta'_e W_{i,e} + \eta_{i,e} \tag{17}$$

where R^H_i is the cumulative return of stock i computed over H days from the event day and W is a matrix of stock-level observed covariates. The estimation of the vector π is described in the next section. All variables are event specific and therefore indexed by the subscript $e \in (2014, 2016)$. Regression coefficients are also indexed by the event because we estimate the model separately for the two announcements.

The coefficient of interest β^H measures the portfolio balance effect of the policy and is identified from the cross-sectional heterogeneity of the model-implied change in systematic risk π. Notice that β^H has a similar interpretation as the coefficient on the interaction term in a diff-in-diff estimation where π measures the intensity of the treatment.

[9] Notice that a reversal would be observed as soon as investors update the expected path of future supply to include a sale of the portfolio of the central bank ($\Delta\lambda_t < 0$). Also, we would observe a reversal if the central bank was to surprise the market by ceasing the purchases before reaching the expected target. Since we do not provide any empirical evidence of an exit from LSAP, we do not formalize this scenario into a proposition.

Following Proposition 1, if stock returns respond to the exogenous shock to supply through the mechanism described in the model, we expect $\hat{\beta}^H$ to be positive and significant at short horizons. As our baseline specification of the short-run effect, we choose $H = 10$ days.

Proposition 2 implies a positive and significant coefficient at any horizon H. We therefore look at $\hat{\beta}^H$ estimated from a regression of cumulative returns over longer horizons (one month, three months, six months and one year) on π. Estimating a positive $\hat{\beta}^H$ at short horizons followed by a lower $\hat{\beta}^H$ at longer horizons would indicate that the initial event return is, at least partially, reversed after some time. Such evidence would be inconsistent with the portfolio balance channel described by our model and would rather suggest a temporary price pressure story, where arbitrageurs with limited capital need some time to absorb the demand shock coming from the central bank.

Proposition 2 also implies that $\hat{\beta}^H$ should be found to be weakly increasing in H. An increasing $\hat{\beta}^H$ indicates that the divergence in the cross-section of returns in the direction of the vector π not only does not vanish, but it becomes larger with time.

In the recent literature on the effect of monetary policy, the common approach to address endogeneity concerns is to employ high-frequency data and to focus on very short windows around policy announcements (e.g. Andrade and Ferroni, 2018, Cochrane and Piazzesi, 2002, Hanson and Stein, 2015, Nakamura and Steinsson, 2018). Estimating long-run effects is challenging in this type of studies that rely on time-series variation for identification due to increasing exposure to confounding factors when moving away from the time of the announcement. Since our identification relies on the cross-section of returns, we can expand the post-event window to estimate the long-run effect of the policy and say something about the elasticity of long-run demand curves for stocks.

Our empirical framework relies on a set of assumptions to conclude that the observed returns are due to the supply shock induced by the policy.

First, to allow for a causal interpretation of the coefficient on π we need the OLS exogeneity assumption to be satisfied. The cross-section of asset purchases would be endogenous if, for instance, the central bank were to overweight underpriced stocks and underweight overpriced ones. Clearly, in this case one could not claim a causal link from the supply shock to the observed subsequent returns. Endogeneity is a concern in papers that try to pin down portfolio balance effects in the case of interventions by the Federal Reserve in Treasury markets where the purchases were likely targeting particularly illiquid or underpriced securities. Since the BoJ is buying ETFs tracking market indices, concerns about the endogeneity of the purchases at stock level are mitigated.

Second, in order for the expected future change in supply to be correctly and immediately impounded into price, we need to assume that the purchase schedule is known by the market.

If, instead, there was uncertainty about what securities the central bank would purchase, event returns would react to these expectations rather than to the actual change in supply. Unlike other asset purchase programs by the Federal Reserve or the ECB, the BoJ announced it would buy according to a well defined and predetermined rule. Since index weights are public information, this ensures that the vector q of supply shocks at stock level is known at the announcement.

Proposition 1 from the model shows that a supply shock parallel to market capitalization would produce returns proportional to stock market betas in the cross-section. Since in the model monetary policy affects prices *only* through the portfolio balance channel, those returns are by definition caused by the change in assets' supply. In reality, asset purchases can affect stock prices through channels other than portfolio balance, possibly in proportion to their market beta. In such case, using the market model as our benchmark, we would not be able to identify the portfolio balance effect of the policy in the cross-section of returns. Our empirical strategy relies on the model prediction that a purchase schedule orthogonal to market valuations would leave a characteristic footprint in the cross-section of abnormal returns computed against a market model. We are able to perform this test empirically since the BoJ is tilting its purchases away from market capitalization. Our identification strategy therefore relies on a third assumption, namely that there is no transmission mechanism of monetary policy other than portfolio balance that would affect prices proportionally to π.

Finally, our specification assumes that the theory-implied measure π of predicted price impact reflects the change in systematic risk correctly. Recall that π is derived in the model under the assumption of market segmentation, i.e. that agents keep the proceeds from the sale of the securities in cash. We discuss violations of this assumption in the Internet Appendix.

4.3 Variable Construction and Summary Statistics

This section presents the data and defines the empirical proxies for the vector u of expected purchases and of the covariance matrix Σ of asset returns, defined in Section 3.2.

To be conservative, in Section 5 we test the model predictions separately for the two policy announcements of the BoJ. All variables are therefore calculated or estimated twice, in order to have two sets of variables, one for each event. For all estimated variables in our analysis we use an estimation window of one year, ending two trading weeks before each BoJ announcement.

4.3.1 Expected Purchases

In the guidelines to the LSAP program, the BoJ states that it would spread its purchases among index-tracking ETFs proportionally to the aggregate AUM of each ETF. In practice, this roughly corresponds to a 50-50 allocation of capital between TOPIX and Nikkei ETFs.

When we go to the data, we assume that this allocation rule not only holds on aggregate over the policy horizon, but also each time the central bank makes a purchase. Under this assumption, the vector u of purchases by the BoJ (in yen) can then be expressed as

$$u_i = T w_{i,T} + N w_{i,N} \tag{18}$$

where T and N indicate the amount of BoJ capital allocated to TOPIX and Nikkei ETFs, respectively, and $w_{i,T}$, $w_{i,N}$ are the weight of stock i in the TOPIX and Nikkei indices.

Since $T \cong N$ in the current purchase program, for the empirical analysis we compute the vector u simply as $w_{i,T} + w_{i,N}$. The vector u proxied in this way still encodes the cross-sectional variation in purchases at the heart of our identification strategy. Given that index weights are time varying, for each event we take $w_{i,T}$ and $w_{i,N}$ to be the index weights as of the end of the month preceding the announcement.

Figure A.4 in Appendix A shows in the top row of each panel, the cross-sectional distribution of stock weights in the TOPIX and the Nikkei 225 index in the month before the event. The top-right panel plots the cross-sectional distribution of the resulting stock-level weights in the BoJ purchase vector $(w_{i,T} + w_{i,N})$. The percentile plots in logarithmic scale clearly show that variation in weights across stocks is substantial.

4.3.2 Covariance Matrix

We estimate the variance-covariance matrix Σ of stock returns using daily returns data from Compustat Global. Because the cross-sectional dimension of our data is larger than the sample size, the sample covariance matrix of returns is a poor estimator of Σ. We therefore use the shrinkage method proposed by Ledoit and Wolf (2004) to obtain a well-conditioned and more accurate estimator, which also ensures that the resulting matrix is always positive definite.

In the model described in Section 3 returns are driven only by fundamentals innovations and changes in supply. However, when we go to the data, this assumption may not hold. We are especially concerned about the impact on the returns moments of other monetary policy announcements during the estimation window. To address this concern, we look at stock

returns net of market returns and we estimate Σ as the cross-sectional covariance of the fitted residuals $\hat{e}_{i,t}$ from a simple market model specified as

$$R_{i,t} = \alpha_i + \beta_i^{mkt} R_{mkt,t} + e_{i,t} \tag{19}$$

where $R_{i,t}$ are daily returns of stock i and $R_{mkt,t}$ is the return on the TOPIX Index used as proxy for the market portfolio. As reported in Table A.2 of the Appendix, our results are robust to estimating Σ on raw returns rather than abnormal returns.

4.3.3 Control Variables

We estimate stocks' sensitivities to changes in the exchange rate by running the following regression separately for each stock i

$$R_{i,t} = \alpha_i + \beta_i^{mkt} R_{mkt,t} + \beta_i^{F} F_t + e_{i,t} \tag{20}$$

Here F_t is the daily percentage change in the exchange rate from US Dollar to Japanese Yen. Estimation results for market and Forex betas are reported in Table A.1 and Figure A.4 in Appendix A. In the bottom rows we plot the cross-sectional distributions of pre-event market betas and of Forex betas together with companies' market capitalization. Table A.1 in Appendix A presents a break down of the summary statistics by Nikkei and non-Nikkei companies.

5 Empirical Results

In this section we test the empirical predictions of the model described in Section 3. Our results show that the ETF program of the BoJ had a significant impact on stock prices and that both the cross-sectional and time-series patterns of the price effect are consistent with a portfolio balance channel. We first perform event studies around the two BoJ announcements and show that the observed price impact is positively related at the stock-level with our ex-ante measure of systematic risk change. We then look at the effect over different horizons and conclude that the impact of the policy is persistent and increasing in time. We propose a simple back-of-the-envelope calculation to quantify the net aggregate portfolio balance effect of the implemented program. We estimate a 22 basis points increase in aggregate market valuation per trillion Yen invested, which corresponds to a unitary price elasticity.

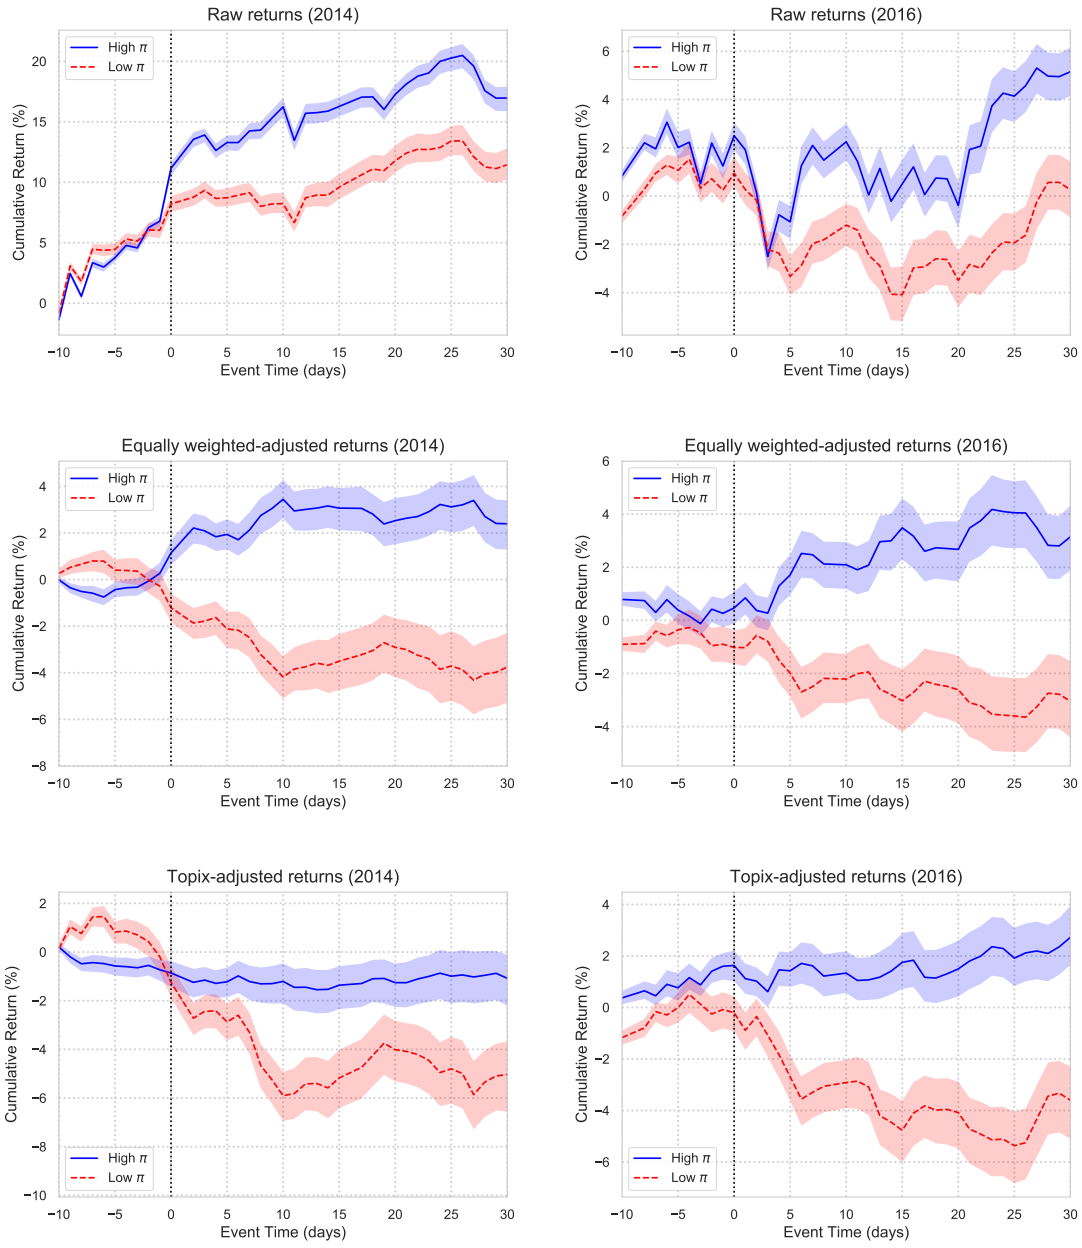


Figure 4: Cumulative returns of high versus low π stocks (in percentage). This figure shows the time series of the mean cumulative returns around the BoJ announcements of stocks with high predicted price impact π against that of low π stocks. The plots on the left refer to the announcement on October 31st, 2014, while those on the right show the reaction to the announcement on July 29th, 2016. The two top panels plot the unadjusted returns. In the four remaining panels returns are adjusted using a market model estimated in a window of one year, as described in Section 4.3.3. An equally-weighted portfolio of stocks in the TOPIX universe is used a proxy for the market portfolio in the middle panels, while the return of the TOPIX index is used in the bottom panels. The blue line is the average for the first quartile of the distribution (firms with the highest predicted price impact), while the red dashed line corresponds to the average for the last quartile (firms with the lowest predicted price impact). Bands represent bootstrapped 95% confidence intervals.

5.1 Event Study

Proposition 1 states that we should observe a positive relationship between each security abnormal event return and the change in its marginal contribution to the risk of the aggregate portfolio.

As a preliminary test of this relationship we rank stocks in the TOPIX universe by the predicted abnormal event return $\pi_i = (\Sigma u)_i$ into four equally-weighted portfolios. Figure 4 presents cumulative returns of the low and high π portfolios. Plots on the left show the event returns around the first policy change in 2014 (when the target purchase amount of ETFs was tripled), while those on the right present the effect of the second change in 2016 (when the target was doubled further). We consider raw returns and abnormal returns based on two versions of the market model with different proxies for the market portfolio, the TOPIX index and an equally-weighted index, respectively. The reported bands represent bootstrapped 95% confidence intervals.

Each plot shows a sizeable and highly significant spread between the returns of high and low π firms opening after the two announcements. While for the 2014 event the reaction seems to be slightly anticipated, in 2016 the effect is delayed by a couple of days. Overall, the pattern of abnormal returns is similar for the two events, with the performance of the high π portfolio being significantly higher than that of the low π portfolio. There is no sign of reversal over 30 days after the announcement, and rather the gap between the two groups appears to increase over time. This preliminary evidence is consistent with both predictions of the model.

5.2 Cross-Sectional Regressions

One might be concerned that, by sorting on π, we are implicitly ranking stocks based on firms' characteristics such as size, export share or market beta, which might explain the heterogeneous response to the announcements and thus the divergence in returns. We therefore run security-level cross-sectional regressions of event returns on the predicted price impact π_i and a set of control variables

$$R_i^H = a_0 + a_1\,\pi_i + a_2\,u_i + a_3\,\log(cap_i) + a_4\,\beta_i^{mkt} + a_5\,\beta_i^F + a_6\,\text{Amihud}_i + \eta_i \qquad (21)$$

For the purpose of these regressions, event returns are defined as the cumulative returns computed over the 10 trading days following the announcement ($H = 10$). We control for each security's weight in the purchase schedule of the BoJ (u), the natural logarithm of its market capitalization, its market beta, its Forex beta and its Amihud ratio as a proxy for illiquidity.

Panel A: October 31st, 2014

	Raw Returns				Abnormal Returns			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
π	57.86***	59.15***	31.92***	23.27***	36.06***	37.75***	39.95***	30.60***
	(8.59)	(8.00)	(4.40)	(3.49)	(4.94)	(4.81)	(5.54)	(4.62)
u		-0.00	-0.02***	-0.02***		-0.00	-0.02***	-0.02***
		(-0.95)	(-3.83)	(-3.43)		(-1.34)	(-3.62)	(-3.20)
Market Beta			0.040	0.025			-0.05	-0.07*
			(0.75)	(0.51)			(-1.48)	(-1.97)
Forex Beta			0.040*	0.043**			0.040*	0.041**
			(1.87)	(2.32)			(1.96)	(2.32)
log(Market Cap)			0.007	0.007			0.005	0.006
			(1.17)	(1.36)			(0.97)	(1.16)
Amihud			0.000	0.000			0.000	4.618
			(0.53)	(0.49)			(0.22)	(0.03)
Observations	1,851	1,851	1,807	1,701	1,851	1,851	1,807	1,701
R-squared	0.106	0.106	0.162	0.203	0.046	0.047	0.108	0.160
Industry FE	NO	NO	NO	YES	NO	NO	NO	YES

Panel B: July 29th, 2016

	Raw Returns				Abnormal Returns			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
π	14.07*	14.30*	12.17*	11.88	15.33*	16.33*	17.49**	16.69**
	(2.09)	(1.93)	(1.68)	(1.78)	(2.10)	(2.08)	(2.43)	(2.52)
u		-0.00	0.001	0.002		-0.00	0.004	0.004
		(-0.29)	(0.19)	(0.38)		(-1.33)	(0.56)	(0.67)
Market Beta			0.006	-0.00			0.002	-0.01
			(0.12)	(-0.06)			(0.06)	(-0.26)
Forex Beta			0.016	0.012			0.019	0.014
			(0.78)	(0.66)			(0.94)	(0.80)
log(Market Cap)			-0.00	-0.00			-0.00	-0.00
			(-0.10)	(-0.05)			(-0.58)	(-0.49)
Amihud			0.000	0.000			0.000	0.000
			(0.43)	(0.18)			(0.41)	(0.11)
Observations	1,905	1,905	1,839	1,734	1,905	1,905	1,839	1,734
R-squared	0.017	0.017	0.021	0.043	0.019	0.019	0.028	0.050
Industry FE	NO	NO	NO	YES	NO	NO	NO	YES

Table 2: Cross-sectional regressions. The tables report the regression coefficients of the cross-sectional regression of returns (in percentage points) on the predicted price impact π and a set of control variables (standardized). Regressions are run separately for the two events. The dependent variable in columns 1-3 is the cumulative raw return, while in columns 4-6 is the cumulative abnormal return with respect to the market model estimated in the pre-event window. Cumulative returns are computed over a 10 days horizon after the announcement date. t-statistics from placebo regressions are in parenthesis; asterisks denote conventional significance levels (***=1%, **=5%, *=10%) based on empirical p-values.

The Internet Appendix A reports summary statistics of the control variables by quartile of π. Consistent with the fact that the policy is heavily skewed towards Nikkei companies, which are on average larger than non-Nikkei ones, we find a positive correlation between π and market capitalization. Market capitalization is therefore an omitted variable in a regression of stock returns on π and we need to control for it. Moreover, the policy announcement could affect equity prices through its impact on the foreign exchange market. Since π is weakly negatively correlated with the Forex exposure β^F, if the yen depreciated as a consequence of the announcement, we would spuriously observe returns proportional to π. We therefore control for the exposure to the exchange-rate by adding β^F to the regressions. We also control for market betas. Notice though that there is no obvious relationship between π and market betas. We include the weights of the BoJ purchase schedule u_i to control for alternative explanations based on the direct effect of purchases in which LSAPs affect asset prices proportional to the amount purchased. Finally, we include industry fixed effects to some specifications, to make sure that our results hold within industries.

We run these regressions on the entire universe of TOPIX firms. Panel A of Table 2 investigates the cross-sectional effect of the BoJ announcement on October 31, 2014 (when the target purchase amount of ETFs was tripled), while Panel B analyzes event returns following the announcement on July 29, 2016 (when the target was further doubled). In either cases, no change was made to the weighting scheme of the purchases. In the first four columns the dependent variable is the cumulative raw return of the stock, while in the last columns the left-hand side variable is the cumulative abnormal return from a market model calculated using pre-event market betas.

On a given day, stock returns are expected to be correlated in the cross-section and therefore the OLS assumption of iid residuals is likely to be violated. We therefore run placebo regressions on the period from January 2009 to March 2013 to get the empirical distribution of the coefficients in the absence of policy shocks, which we use to compute robust standard errors. The placebo event days are chosen randomly on non-overlapping periods to ensure that the empirical distribution is constructed from independent draws. For regressions involving short-horizon returns (up to 3 months) we impose that placebo event periods do not include BoJ meetings on which important monetary policy announcements were made. Namely, we exclude the meetings of February 1st 2013, March 25th 2013, June 18th 2012 and the announcement of the post-tsunami intervention in March 14th 2011. On all regression tables of this paper we report the empirical p-values computed using this methodology.

Consistent with Proposition 1, the coefficient on the predicted price impact π is positive and significant across specifications and events. For the 2014 announcement, the baseline specification with raw returns reported in columns (1) shows a remarkable R^2 above 10%, suggesting that our expected price impact π is crucial to explain the heterogeneity of event

returns. As it was already visible from the plots in the previous section, the results are weaker for the 2016 event. In particular, the portion of explained variance for the 2016 policy announcement is significantly lower, consistent with the smaller change in the target purchase amount. Still, the coefficient on π is positive and significant at the 10% confidence level in most specifications. The coefficient turns however insignificant when we include industry fixed effects in the specification with raw returns.

Results show that the effect of π is robust to the inclusion of the vector u of purchased amounts. This horse race provides additional support for the portfolio-balance channel against a local channel where spillovers are negligible. The model predicts that the effect of u should be insignificant once we control for π. This is indeed what we find in the second specification of each panel. The coefficient on u turns however negative and significant in columns (3) and (4) of panel A. We find that this is due to the inclusion of the control for market capitalization since u and market cap are highly correlated, as it is natural to expect. Still, this does not affect the size and the significance of the coefficient on π.

Results also show that controlling for the exposure to the exchange rate does not impair the significance of the coefficient on π. The coefficient on β_F is positive and significant in 2014, when the BoJ announcement was followed by a rise in the Forex. In 2016, on the other hand, the coefficient on β_F is not significant, consistent with the fact that the Forex did not move significantly (see Figure A.5 in the Appendix).

In column (2) of the regression using raw returns as dependent variable, the coefficient on π drops significantly. This is, as expected, due to the fact that control variables play an important role in explaining cross-sectional returns variation, as documented by a significantly larger coefficient of determination. In particular, market beta, Forex beta and market capitalization incrementally increase the regression's R^2 and dampen the coefficient on π. In columns (3) the number of observations drops slightly because of missing data on trading volume needed to estimate the Amihud ratio. In columns (4) the sample is further reduced because of missing information on industry classification.

5.3 Time-Series Pattern

In this section, we test the long-run predictions of the model summarized in Proposition 2. To this end, we estimate the cross-sectional model specified in equation (21) at different horizons H over which cumulative returns are calculated. Results are reported in Table 3.

	Abnormal Returns 2014						Abnormal Returns 2016					
	5	10	21	63	126	252	5	10	21	63	126	
π	17.78**	39.95***	28.31***	72.21***	171.9***	305.5***	17.43**	17.49**	33.70***	40.86**	93.52***	1
	(3.09)	(5.54)	(2.89)	(4.52)	(8.97)	(11.91)	(3.03)	(2.43)	(3.44)	(2.56)	(4.88)	
u	-0.00	-0.02***	-0.01	-0.03***	-0.03**	-0.07***	0.011**	0.004	0.026**	0.016	0.052***	0
	(-0.30)	(-3.62)	(-1.21)	(-3.20)	(-2.46)	(-3.43)	(2.22)	(0.56)	(2.76)	(1.48)	(3.31)	
Market Beta	-0.02	-0.05	-0.07	-0.16*	-0.29**	-0.40***	0.026	0.002	0.025	0.025	0.041	
	(-0.84)	(-1.48)	(-1.36)	(-2.02)	(-2.04)	(-2.51)	(0.86)	(0.06)	(0.47)	(0.32)	(0.28)	
Forex Beta	0.036**	0.040*	0.101***	0.097**	0.043	-0.13	-0.00	0.019	0.061*	0.070*	0.220***	
	(2.17)	(1.96)	(3.83)	(2.29)	(0.61)	(-0.02)	(-0.23)	(0.94)	(2.30)	(1.67)	(3.10)	
log(Market Cap)	0.001	0.005	0.001	0.001	0.005	-0.00	-0.00*	-0.00	-0.01*	-0.02	-0.07***	-
	(0.32)	(0.97)	(0.17)	(0.06)	(0.25)	(-0.27)	(-1.60)	(-0.58)	(-1.81)	(-1.80)	(-3.52)	
Amihud	0.001	0.000	-1.02	0.001	0.016***	0.018***	0.000	0.000	-0.00	-0.00	-0.00**	
	(0.78)	(0.22)	(-0.01)	(0.70)	(6.27)	(7.17)	(0.47)	(0.41)	(-0.47)	(-0.92)	(-3.12)	
Observations	1,807	1,807	1,807	1,807	1,807	1,807	1,839	1,839	1,839	1,839	1,839	
R-squared	0.055	0.108	0.073	0.098	0.153	0.119	0.051	0.028	0.079	0.077	0.178	

Table 3: Cross-sectional regressions over different horizons. The table report the coefficients of cross-sectional regressions of cumulative returns (in percentage points) computed at different horizons on the predicted price impact π and a set of control variables (standardized). Regressions are run separately for the two events. The dependent variable is the cumulative abnormal return with respect to the market model estimated in the pre-event window. t-statistics from placebo regressions are in parenthesis; asterisks denote conventional significance levels (***=1%, **=5%, *=10%) based on empirical p-values.

At portfolio level, Figure 4 suggests that the cross-sectional effect of the BoJ announcements is long-lasting and weakly increasing over time. The regression analysis confirms that the evidence holds at stock-level and after controlling for security specific characteristics. The vector π is positively and significantly related to cross-sectional stock returns at every horizon H after the announcement. In other words, the model implied changes in systematic risk estimated ex-ante are a significant predictor of post-event abnormal returns across stocks.

We find no evidence of reversal of the initial price impact even one year after the event. The positive coefficients on π at longer horizons provide evidence of a persistent effect of the purchase program on prices. This result is consistent with the portfolio-balance channel, whereby a permanent reduction in the stock of assets held by private investors leads to a permanent decrease in risk premia.

Claims about the persistence of the effect of QE are hard to make in event studies where identification relies on time-series evidence over a short window around the announcement. The main concern in those type of studies is that long-lasting effects might be due to changes in expectations about the future path of policy rates (signalling channel), and more generally that results might be confounded by the release of macro news in the subsequent days. Since the shock to supply induced by the BoJ has a unique cross-sectional shape, it is unlikely that the observed effect is due to shocks other than the purchase program, suggesting a causal effect of the policy.

The absence of reversal is a key prediction of the portfolio-balance channel. Still, we cannot completely rule out other explanations. While price pressure and limits to arbitrage would generally predict a temporary effect, the fact that the purchase program is ongoing over the entire sample period might prevent us to observe a reversal. In Section 6, we address the possibility that (part of) the effect might be due to a continuous price pressure that prevents prices from reverting to the pre-announcement level.

A second finding from Table 3 is that the estimated coefficients on π are generally increasing in H.[10] This suggests that the effect of the BoJ policy is not immediately reflected on prices, but it is increasingly impounded over time. Post-event returns in the same direction of the announcement effects are predicted by the model through the decrease in residual duration of the program. However, as we discuss in Section 3, this effect is expected to be small for realistic values of the interest rate. The significant post-event abnormal returns observed in the data are consistent with investors expectations about the size of the program increasing over time. In terms of the model, the results are consistent with a λ_t increasing in t. This suggests that investors might be extrapolating current purchases above and beyond the policy horizon or that they might not believe to a full commitment of the central bank to the announced

[10]Notice that the cumulative returns on the left-hand side of the regression are computed as cumulative sums rather than cumulative products in order to avoid a mechanical effect when increasing the horizon.

purchase target at first, but slowly update their beliefs. In the current setting we cannot disentangle between these explanations, nor convincingly claim that post-event returns are in fact driven by updating in beliefs. The question is therefore open for future research.

In Appendix A we present some robustness evidence. We show that the observed price impact cannot be explained by industry effects in Table A.3, which presents regression results including industry fixed effects. To show that the difference in returns is not simply driven by an over-performance of Nikkei stocks versus non-Nikkei stocks, in Table A.4 we re-run the analysis including a dummy variable for Nikkei stocks. Finally, Table A.5 includes both industry and Nikkei fixed effects. The results remain largely unchanged across specifications.

5.4 Quantification of Portfolio-Balance Effects

In this section we propose a simple back-of-the-envelope calculation to try to quantify the net aggregate portfolio balance effect of the BoJ intervention from the coefficient estimated in the cross-section. From this quantity we then derive an estimate of the aggregate elasticity of equity demand curves.

For this calculation we want to use the average effect of the policy across the two events, so we first run again our main regression in equation (21) over the pooled sample, including event fixed-effects FE_e to allow for a different intercept across the two announcements. Precisely, we estimate the following regression model for each daily horizon $h \in 1, \ldots, 252$

$$R^h_{i,e} = \beta^h \pi_{i,e} + \gamma^h X_{i,e} + \delta^h \, \text{FE}_e + \varepsilon_{i,e} \tag{22}$$

where X is a vector of control variables that depends on the regression specification and $e \in (2014, 2016)$ is an index numbering the events. Since we are considering both events together, we need to rescale the π vectors to take into account the different magnitude of the announcements. Therefore we multiply π_{2014} by 3 and π_{2016} by 6 to reflect the magnitude of the target amount announced by the BoJ in the two events, respectively. We include market capitalization in each specification to control for the size factor, which is expected to become more relevant as the horizon increases. In the second specification we also control for market and Forex betas. In the third specification we additionally include each stock's Amihud ratio to control for liquidity.

Given $\hat{\beta}^h$ from the estimation, the predicted net return through the portfolio balance channel for security i is $\hat{R}^h_{i,e} = \hat{\beta}^h \pi_{i,e}$ [11]. To aggregate the effect at market level, we calculate for each

[11] Notice that the estimated $\hat{\beta}^h$ allows us in principle to compare the impact of alternative purchase schedules u' that the central bank could have implemented, conditional on the same covariance matrix Σ. In this section we are interested in the estimated portfolio balance effect of the actual purchase portfolio.

	1 week	1 month	3 months	6 months	1 year
(1) Baseline	3.54	10.10	10.28	25.17	22.32
(2) Control for market and Forex	2.23	7.72	9.53	22.08	22.45
(3) Control for market, Forex and liquidity	1.56	7.17	9.56	19.80	22.05

Table 4: Portfolio Balance Effects. The table presents the estimated net portfolio balance effect on the market, expressed in basis points per Trillion Yen invested by the central bank into the ETF purchase program. We report point estimates for the net effect impounded into prices over increasing horizons, from three models employing different sets of control variables defined in the text.

event $e \in (2014, 2016)$ the predicted market return as the value-weighted sum of security level predicted returns at every horizon

$$\hat{R}_e^h = \hat{\beta}^h \sum_i w_{i,e} \pi_{i,e} \tag{23}$$

We then divide by the capital commitment by the central bank to obtain the induced market return per trillion yen. Considering the two-year policy horizon, this amounts to 6 trillion Yen for 2014 and 12 trillion Yen for 2016, with the underlying assumptions that each announcement was completely unexpected and that investors are reacting to the announced program size. Thus, the per yen estimated average market return induced by the policy through the portfolio-balance channel is calculated as

$$\hat{R}^h = \frac{1}{2} \left(\hat{R}^h_{2014}/6 + \hat{R}^h_{2016}/12 \right) \tag{24}$$

Results of this exercise are reported in Table 4 for the three specifications. The last column shows an estimated long-term impact of about 22 basis points increase in market value per trillion yen employed. With about ¥500 trillion of total market capitalization, this implies an elasticity close to one since each yen invested translates into an increase of the market valuation by roughly one yen.

Figure 5 plots the time-series evolution of the point estimate for the third specification, showing that the portfolio balance effects are slowly impounded into prices. Consistent with the qualitative prediction of our model, a momentum-like pattern is visible over the first 100 trading days following the announcement.

Notice that the quantity we are estimating in this section is the aggregate portfolio-balance effect of the policy on the returns of stocks that are included in the TOPIX index. While the policy is expected to have additional effects through different channels, our empirical methodology allows us to identify and quantify the portfolio balance channel. In turn, this allows us to derive an estimate of the price elasticity of the demand curve for stocks, which has to be intended as local to the First Section of the TSE. We acknowledge that the policy

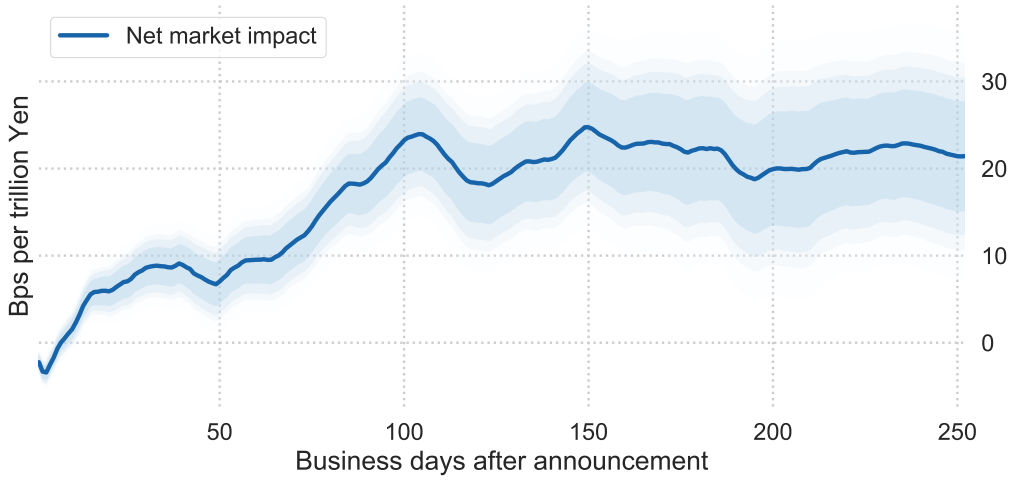


Figure 5: Portfolio Balance Effects. This figure plots the time-series evolution of the estimated portfolio balance effect induced by the BoJ purchase program, expressed in basis points per trillion Yen invested. The estimates are based on specification (3), which includes controls for stocks liquidity, market beta and exposure to the US-JPN Forex exchange rate. Thus the estimated market impact can be interpreted as the counter-factual policy effect, net of of alternative channels and confounding factors. Shaded areas denote 10%, 5% and 1% confidence intervals.

might have produced spillover to other asset classes (this is indeed a prediction of the model), but we are not considering them in this paper.[12]

The derivation of the aggregate portfolio balance effect described in this section relies on two main assumptions. First, it depends on our assumption that the representative agent in our model re-invests the proceeds from the sale of stocks to the central bank at the constant risk-free rate.[13] The key point here, is that the risk-free asset is uncorrelated with the Japanese stocks. While this is not a concern for the identification of the portfolio-balance effect, it is more problematic when we try to quantify the effect, since our approach might be providing a biased estimate of the aggregate effect if the assumption is not valid.

To understand the direction and magnitude of the potential bias, in the Internet Appendix we extend the model to allow the representative agent to re-invest the proceed in a security correlated with the targeted assets. We derive an expression of the resulting bias if we incorrectly assume the above mentioned assumption, which can be seen as a form of market

[12]Spillovers to unaffected stocks is already a key point in Greenwood (2005) and is to be expected in this setting as well. We believe that the cross-sectional heterogeneity among the stocks in the TOPIX is sufficient to support our arguments. Moreover, the TOPIX index covers all First Section companies in the Tokyo Stock Exchange (TSE), which are the majority of public companies in Japan and is by far largest section of the TSE in terms of market capitalization and trading volume.

[13]We thank the Anonymous Referee for pointing out this issue.

segmentation. The bias is a function of the market weighted average of the covariance of the omitted variable and π, where the omitted variable is the vector of covariances between the re-investment security and the stocks. The sign of the bias is ambiguous and depends on Σ, u, the market weights and the re-investment security. We therefore run simulations of the model using the parameters estimated in the data and assuming different re-investment securities, namely S&P500, 10-year US Treasury bonds and 10-year JGBs. The estimated bias is positive using the S&P500 and negative using long-term government bonds. The magnitude of the bias is relatively small, around (positive or negative) 10%. Depending on which direction the bias is going, the estimated elasticity of 1 might be slightly over- or under-estimating the true elasticity of Japanese equities.

A second reason why our approach might be delivering a biased estimate of the aggregate effect is that in our calculation we are assuming that the market is reacting to the announced size of the program. If the market is in fact reacting to expectations of a smaller program, either because investors think the BoJ will not reach the announced target or because it will soon unwind its portfolio, then the estimated elasticity of 1 represents an upper-bound for the true price elasticity. Vice versa, the estimated elasticity of 1 would be a lower-bound if investors believe that the BoJ will continue the purchase program beyond M.

6 Portfolio Rebalancing or Price Pressure?

The previous section shows that the reaction of stock prices to the upward revisions of the purchase target is consistent with a portfolio balance channel. Even at long horizons, the expected change in systematic risk is key to explain the cross-sectional variation in returns after the policy announcement. We interpret the persistence of the effect as evidence of downward sloping long-run demand curves.

An alternative explanation for the observed persistence relies on the continued pressure exercised by the BoJ through repeated purchases. If short-run demand curves are downward sloping, abnormal volumes induced by the BoJ during intervention days might push prices above fundamentals. Such effects are usually motivated by limits-to-arbitrage and are expected to revert quickly. However, as the central bank is expected to buy repeatedly, arbitrageurs may refrain from betting against mispricings and fail to bring prices back to fundamentals. If this was the case, the absence of reversal could not be interpreted as evidence for long-run demand curves sloping down. In the spirit of D'Amico and King (2013), we will refer to this kind of dynamics as *flow effect* of the policy. This naming highlights that under this explanation the price impact is caused directly by the trading volume (or *flow*) of the central bank, rather than by reduction in systematic risk which underlies the portfolio balance channel.

The repeated price pressure story implies that we should observe higher positive abnormal returns on intervention days and that these should be proportional in the cross-section to the abnormal trading volume generated by the intervention. In this section we first introduce a reduced form model that exploits the time-series and cross-sectional variation in daily purchases by the BoJ to estimate the flow effect of the policy. We then use the predicted returns from that model to remove the flow effect component from stock returns. Finally, we re-run the analysis of Section 5.2 on these *net* returns. By comparing the coefficient estimated in this way to the one in the previous section, we can assess how much of the observed price impact and its persistence is due to repeated price pressure rather than the portfolio balance mechanism.

Evaluating the relative magnitude of these two channels is essential to draw conclusions on the elasticity of long-run demand curves for stocks. The distinction between the two explanations has also practical consequences for policy makers regarding the exit strategy from the purchase program. A repeated price pressure story predicts prices to revert as soon as the buying pressure from the central bank stops, making the accumulated size of the balance sheet *de facto* irrelevant beyond that point. On the contrary, in the model of Section 3 the aggregate impact of the policy is unaffected by the timing of the purchases. In the extreme case where the central bank buys everything on the announcement day, the model predicts an immediate, complete and permanent price adjustment.

6.1 Purchase Frequency and Volumes

The QQE was announced on April 4, 2013, and the asset purchases were then gradually carried out. In its official statements, the BoJ does not commits itself to any particular purchase frequency and does not reveal in advance the days in which it will buy. Ex-post, we can see from Panel A of Figure 6 that the bank has been buying fairly consistently once to twice a week over the sample period. The blue crosses in Panel B of Figure 6 indicate intervention days. On the y-axis we report the ratio between the amount purchased and the aggregate trading volume in the underlying stock market on that day.

Since the purchase frequency remained stable over the policy horizon, the upward revisions of the annual target in October 2014 and in July 2016 translated into an increase of the daily purchased amount. However, even in the last period, the quantity purchased by the BoJ represented less than 5% of the daily market volume, a threshold which is often used by practitioners as guideline for when a trade is expected to have a significant price impact. At the stock level, the purchases by the BoJ account for more than 5% of the daily trading volume on average only for 5.2% of the targeted stocks.

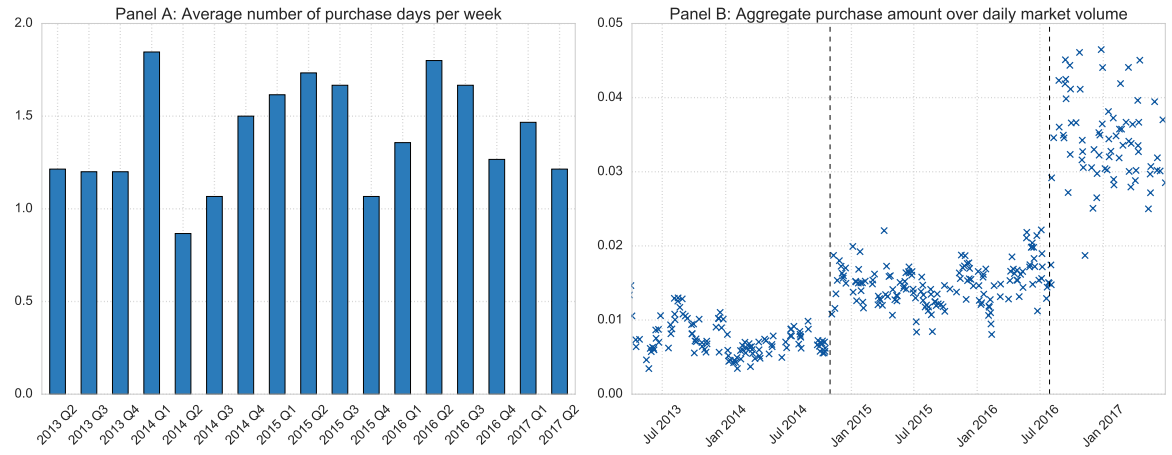


Figure 6: Purchase Frequency and Volume. Panel A plots the average number of purchase days per week at quarterly frequency. Panel B plots the ratio between the yen amount purchased by the BoJ on a given day and the aggregate trading volume in yen on that day. The aggregate trading volume is computed as the sum of the trading volume of the securities targeted by the policy.

6.2 Empirical Setup and Results

To quantify the direct price impact of purchases we estimate a dynamic model in the spirit of Eser and Schwaab (2016) that relates daily stock returns to daily flows from the central bank. The model is specified as

$$AR_{i,t} = \alpha + \beta_0 AV_{i,t} + \beta_1 AV_{i,t-1} + \beta_2 \left(\sum_{k=2}^{K} \rho^{k-2} AV_{i,t-k} \right) + \varepsilon_{i,t} \tag{25}$$

where the left-hand side variable $AR_{i,t}$ is the daily abnormal return of stock i relative to the market model, estimated following the methodology outlined in Section 4.3. On the right-hand side, the BoJ-induced abnormal volume $AV_{i,t}$ is defined by

$$AV_{i,t} := \frac{\text{BoJ Flow}_{i,t}}{\text{E}\left[\text{Volume}_{i,t}\right]} \tag{26}$$

and measures the size of the purchased amount of stock i on day t relative to the average market volume of that stock. The purchased amount $\text{BoJ Flow}_{i,t}$ is computed as $\frac{1}{2}(w_{i,T}+w_{i,N})A_t$, where $w_{i,T}$ is the weight of stock i in the TOPIX index, $w_{i,N}$ is the weight of stock i in the Nikkei 225 index and A_t is the value of ETFs purchased by the BoJ on day t. Here we assume that each trade of the BoJ in the ETF market translates into proportional shocks to the underlying basket on the same day.[14] The average daily volume $\text{E}\left[\text{Volume}_{i,t}\right]$ is estimated over

[14]Underlying securities inherit shocks that occur in the ETF market both through primary market arbitrage as well as through the arbitrage that takes place continuously in the secondary market and that is carried

a backward-looking window of six months excluding days in which the BoJ is intervening. On non-purchase days the abnormal volume is therefore zero for every stock in our sample, while it is strictly positive on purchase days.

The model includes lagged values of AV to capture the permanent component of the price pressure, net of transitory and delayed effects of purchases. The long-run effect of the flow-induced price impact can be computed from the estimated coefficients as

$$F = \beta_0 + \beta_1 + \beta_2 \left(\sum_{k=2}^{K} \rho^{k-2} \right) \tag{27}$$

The parameter $\rho \in (0, 1)$ determines how long it takes for prices to adjust following an intervention. If ρ is close to zero the dynamic of the flow effect is exhausted after two days. $F \approx 0$ implies that temporary price impacts, if any, are fully reverted. This in turn would mean that the price pressure story does not contribute to explain the persistence of the policy impact documented in Section 4. On the contrary, $F > 0$ implies that (part of) the persistence attributed to the portfolio balance mechanism might be due to the direct impact of the flow of BoJ purchases.

The identification of the direct impact of the purchases (flow effect) in this panel regression framework relies both on the exogeneity of the cross-sectional variation of the purchases and on the predetermination of the purchase amounts with respect to prices. The exogeneity in the cross-section is discussed extensively in Section 1 and mainly relies on the fact that the weighting system of the Nikkei 225 introduces significant variation in the cross-section of purchases that is unrelated to firms' fundamentals. Predetermination of the purchases is not straightforward in the current context. The criteria used by the BoJ to decide whether and how strongly to intervene on a particular day are not public information, however there are reasons to believe that the BoJ tends to intervene on days when the market is falling. In fact, the median stock return is significantly lower on intervention days (-0.6%) relative to non-intervention days (0.3%). To tackle this potential endogeneity of BoJ flows we specify the regression model in terms of abnormal returns. Given that the BoJ might be using the return on the market as a signal for whether to intervene, removing the contemporaneous return on the market should mitigate the issue. The fact that mean and median abnormal returns are

out by hedge funds and high-frequency traders (Ben-David et al., 2018). Secondary market arbitrageurs make profits by opening their positions when the price of the ETF deviates from NAV and holding them until prices converge. Our identification of the flow effect of the Policy relies on arbitrageurs trading on the same day as the BoJ. For secondary market arbitrage, this is a reasonable assumption. Competition among arbitrageurs implies that hedge funds and high-frequency traders will open their positions as soon as they observe the ETF trading at a premium over the NAV. That such arbitrage opportunities exist on the days when the BoJ buys is consistent with the evidence in Figure 3, since growth in AUM is consistent with upward pressure on ETF prices. The results reported in Table 5 provide further support for the validity of this assumption since they show that most of the price impact seem to take place on the event day and the day after.

		Panel A					Panel B		
Model	K	β_0	β_1	β_2	ρ	F	$\tilde{a}_1$	Flow Effect	Port Balance
(1)	0	0.011				0.011	138.644	5.63%	94.37%
		(9.356)					(11.175)		
(2)	1	0.004	0.015			0.019	132.288	9.96%	90.04%
		(3.314)	(11.358)				(10.662)		
(3)	2	0.004	0.015	-0.001		0.018	132.608	9.74%	90.26%
		(3.358)	(10.979)	(-0.679)			(10.688)		
(4)	5	0.004	0.015	-0.001	0.001	0.018	132.567	9.77%	90.23%
		(3.396)	(10.997)	(-0.694)	(0.035)		(10.685)		
(5)	10	0.004	0.015	-0.001	0.001	0.018	132.567	9.77%	90.23%
		(3.396)	(10.997)	(-0.694)	(0.039)		(10.685)		

Table 5: Flow Effect The table reports results from the estimation of the dynamic model described in (25), where a different value for the number of lags K is used in each specification. The models are estimated with maximum likelihood assuming normally distributed error terms and constraining the persistence parameter ρ in the unit interval. Panel A presents the estimated model parameters and the implied long-run effect F. Panel B shows OLS estimates of the coefficient $\tilde{a}_1$ resulting from a cross-sectional regression of cumulative abnormal returns, purified from the estimated flow effects, on the predicted price impact π resulting from the portfolio balance model of Section 3. The decomposition into *flow* and *portfolio balance* components is obtained by comparing $\tilde{a}_1$ with the coefficient a_1 from Section 5.2 based on standard CARs.

not significantly different from zero in both intervention and non-intervention days supports our claim.

We do not include an announcement dummy in the specification because on those days no ETF purchases were made by the BoJ. Looking at the time series of BoJ purchases, we see that the bank intervened two weeks before and one week after the first upward revision of the purchase target on October 31, 2014. Similarly, no purchases were made on July 29, 2016. Purchases are registered on the previous day and four days after.

We estimate five different specifications of model (25). We start considering only contemporaneous volumes ($K = 0$), then we augment the specification to K equal to $1, 2, 5$ or 10. The estimated parameters are reported in Panel A of Table 5 together with the implied long-run impact. The positive coefficients on β_0 and β_1 suggest that abnormal returns are significantly higher during purchase days for stocks experiencing a higher degree of buying pressure. The negative but not significant value of β_2 and a persistence parameter ρ close to zero suggest that such a price impact is not reverted in the next trading weeks and give rise to a positive long-run component F in every specification.

The results indicate a positive and persistent flow effect of the policy, which might lead to an overestimation of the portfolio balance channel in the previous section. To quantify the

consequences of not taking flows into account, we construct the flow induced returns as the fitted values of the estimated model

$$\widehat{AR}_{i,t}^{Flow} = \hat{\beta}_0 AV_{i,t} + \hat{\beta}_1 AV_{i,t-1} + \hat{\beta}_2 \left(\sum_{k=2}^{K} \hat{\rho}^{k-2} AV_{i,t-k} \right) \tag{28}$$

which we subtract from stock returns to remove the direct impact of the BoJ purchases

$$\widetilde{AR}_{i,t} = AR_{i,t} - \widehat{AR}_{i,t}^{Flow} \tag{29}$$

We then estimate our main regression (21) using $\widetilde{AR}$ instead of AR, computing the cumulative abnormal returns over a one-year horizon following the two event dates and we regress them on the predicted price impact vector π. We pool the 2014 and 2016 events together to obtain a unique estimate $\tilde{a}_1$ of the price impact of the policy through the portfolio balance channel. The ratio between $\tilde{a}_1$ and its counterpart $\hat{a}_1$ obtained estimating the model with the cumulative returns computed from AR, gives us the fraction of the estimated portfolio balance impact that might be explained by the price pressure channel.

Panel B of Table 5 summarizes the results of this second step, showing that the fraction of the observed cross-sectional pattern explained by the price pressure channel ranges between 5% and 10% depending on the specification. It must be noted that these figures represent upper bounds for the persistent flow effect of the policy, since this might be amplified by expectation updates consistent with the portfolio balance model, if investors learn about the commitment of the central bank through the realization of its purchases.

Taken together, the results of this section suggest that the price pressure generated by the central bank at the stock level plays a limited role in explaining the impact of the policy. We conclude that the observed cross-sectional pattern of stock returns is mostly generated by the portfolio balance channel rather than continued price-pressure arising from the central bank flows.

7 Policy Implications

In this section we discuss the policy implications of our results. Based on our theoretical framework, we show formally that the heterogeneity uncovered by our empirical analysis could be avoided if the central bank would buy the value-weighted market portfolio, since this would lead to a homogeneous reduction of firms' cost of capital in the cross-section.

Recall that in our model the cost of capital of each firm is proportional to its marginal risk contribution to the market portfolio (systematic risk). Formally, the vector of risk premia

prior to the BoJ intervention is proportional to VQ, where V is the variance-covariance matrix of fundamentals and $Q \in \mathbb{R}^n$ is the vector of shares outstanding.

As soon as the central bank purchases a quantity $q \in \mathbb{R}^n$, the cost of capital is affected and converges to $V(Q - Mq)$. In particular, firm i experiences a percentage shift in its perceived cost of capital equal to

$$\Delta k_i = \frac{(V(Q - Mq))_i}{(VQ)_i} - 1 \tag{30}$$

Notice that Δk_i is not necessarily negative, thus some firms may experience an increase in their financing costs ($\Delta k_i > 0$), even if the central bank buys some of their shares ($q_i > 0$).

It follows that a homogeneous impact on risk premia can be achieved with a vector of purchases proportional to Q. If the purchase schedule is $q^* = aQ$ for $a \in \mathbb{R}$, the effect on firm i is

$$\Delta k_i^* = \frac{(V(Q - Mq^*))_i}{(VQ)_i} - 1 = \frac{((1 - Ma)VQ)_i}{(VQ)_i} - 1 = \frac{((1 - Ma)VQ)_i}{(VQ)_i} - 1 = -Ma \tag{31}$$

which does not depend on i and is thus homogeneous across companies.

In the case of Japan, a purchase schedule q parallel to Q corresponds to the BoJ limiting its purchases of ETFs to those tracking the value-weighted TOPIX Index[15]. Buying ETFs tracking the price-weighted Nikkei 225, on the other hand, introduces a component in q which is orthogonal to Q. This, in turn, leads to heterogeneous consequences for firms financing costs, which can be interpreted as a distortion of the market allocation mechanisms. Figure A.6 shows that the distortion is evident also at the industry-level.

Under the assumption that a homogeneous effect is the preferred outcome of the policy, we infer from the model that the central bank should stop buying Nikkei-indexed ETFs. More precisely, the central bank should schedule future purchases with the objective of re-shaping its equity portfolio in a value-weighted fashion.

A change of policy in this direction was solicited by a number of critics of the purchasing program, and on September 2016 the BoJ changed the guidelines for its asset purchases, reducing the share of capital flowing to ETFs tracking the Nikkei 225 Index and increasing its holdings of ETFs tracking the TOPIX. This brought the cross-sectional allocation of capital closer to what market capitalization would justify.

To date, the BoJ has not completely abandoned the price-weighted Nikkei Index, nor it is bringing its already accumulated holdings towards value-weighted proportions. According to our model, the BoJ should make sure to bring its holdings proportional to companies market capitalizations if it wants to amend the allocational side-effects of the policy.

[15] A purchase of $u' = aW_{\text{Topix}}$ in Yen corresponds to $u = aQ$ in shares, since the TOPIX is value-weighted.

8 Conclusion

In this paper we study asset pricing implications of the ETF purchase program undertaken by the BoJ since April 2013. The analysis is supported by a dynamic asset pricing model, featuring multiple assets with time-varying supply due to open market operations of the central bank.

To identify the net portfolio balance effect of the policy our empirical analysis exploits the exogeneity and the cross-sectional dimension of the BoJ's purchase schedule, which mitigates endogeneity problems characteristic of other studies.

We show that the intervention has a positive and persistent effect on domestic equity prices, thus reducing the cost of equity capital of domestic companies. We provide empirical evidence that the effect is consistent with a portfolio balance channel both in the cross-section and in the time-series.

This evidence suggests that demand curves for stocks are downward sloping in the long-run. We estimate an economically significant increase of 22 basis points in aggregate market valuation per trillion Yen invested into the program, which corresponds to a price elasticity of 1. The mechanism behind downward sloping demand curves in the model is through the change in the structure of systematic risk held by the private sector induced by the central bank's intervention. This change in the composition of risk leads to a new discount factor and consequently to price adjustments.

We also show that the outright purchases of the BoJ generate positive and persistent pressure on prices. Our estimates of the portfolio-balance channel remain significant after accounting for these flow effects of the policy.

Our results shed light on the side-effects of the LSAP, uncovering a highly heterogeneous impact in the cross-section of firms' cost of equity capital, both at the firm and at the industry level. Using our theoretical framework to evaluate the impact of arbitrary purchase schedules, we find that the observed heterogeneity in the price effects mainly arises from the weight given to the Nikkei 225 price-weighted index. Capital injections shaped according to market weights would instead induce a cross-sectionally homogeneous change in the cost of capital.

References

Andrade, P. and Ferroni, F. (2018). Delphic and odyssean monetary policy shocks: Evidence from the euro area.

Bagwell, L. S. (1992). Dutch auction repurchases: An analysis of shareholder heterogeneity. *The Journal of Finance*, 47(1):71–105.

Barberis, N. and Thaler, R. (2003). A survey of behavioral finance. *Handbook of the Economics of Finance*, 1:1053–1128.

Ben-David, I., Franzoni, F., and Moussawi, R. (2018). Do ETFs increase volatility? *The Journal of Finance*, 73(6):2471–2535.

Bernanke, B. S. (2009). *Federal Reserve policies to ease credit and their implications for the Fed's balance sheet: a speech at the National Press Club Luncheon, National Press Club, Washington, DC, February 18, 2009*. Number 443.

BoJ (2013). Introduction of the Quantitative and Qualitative Monetary Easing, April 4th.

Brunner, K. and Meltzer, A. H. (1973). Mr. Hicks and the "Monetarists". *Economica*, 40(157):44–59.

Buraschi, A. and Whelan, P. (2015). Bond markets and unconventional monetary policy. *Handbook of Fixed-Income Securities*, First Edition.

Cochrane, J. H. and Piazzesi, M. (2002). The Fed and interest rates - A high-frequency identification. *American Economic Review*, 92(2):90–95.

Coval, J. and Stafford, E. (2007). Asset fire sales (and purchases) in equity markets. *Journal of Financial Economics*, 86(2):479–512.

Culbertson, J. M. (1957). The term structure of interest rates. *The Quarterly Journal of Economics*, 71(4):485–517.

Curdia, V. and Woodford, M. (2011). The central-bank balance sheet as an instrument of monetary-policy. *Journal of Monetary Economics*, 58(1):54–79.

D'Amico, S. and 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(2):425–448.

Eggertsson, G. B. and Woodford, M. (2004). *Optimal monetary policy in a liquidity trap*. National Bureau of Economic Research (No. w10840).

Eser, F. and Schwaab, B. (2016). Evaluating the impact of unconventional monetary policy measures: Empirical evidence from the ECBs securities markets programme. *Journal of Financial Economics*, 119(1):147–167.

Frankel, J. A. (1985). Portfolio crowding-out, empirically estimated. *The Quarterly Journal of Economics*, 100(Supplement):1041–1065.

Gagnon, J., Raskin, M., Remache, J., and Sack, B. P. (2010). *Large-scale asset purchases by the Federal Reserve: did they work?* FRB of New York Staff Report 441.

Greenwood, R. (2005). Short-and long-term demand curves for stocks: theory and evidence on the dynamics of arbitrage. *Journal of Financial Economics*, 75(3):607–649.

Greenwood, R. (2007). Excess comovement of stock returns: Evidence from cross-sectional variation in nikkei 225 weights. *The Review of Financial Studies*, 21(3):1153–1186.

Hamilton, J. D. and Wu, J. C. (2012). The effectiveness of alternative monetary policy tools in a zero lower bound environment. *Journal of Money, Credit and Banking*, 44(s1):3–46.

Hanson, S. G. and Stein, J. C. (2015). Monetary policy and long-term real rates. *Journal of Financial Economics*, 115(3):429–448.

Harris, L. and Gurel, E. (1986). Price and volume effects associated with changes in the S&P 500 list: New evidence for the existence of price pressures. *the Journal of Finance*, 41(4):815–829.

Hau, H., Massa, M., and Peress, J. (2009). Do demand curves for currencies slope down? Evidence from the MSCI global index change. *The Review of Financial Studies*, 23(4):1681–1717.

He, Z. and Krishnamurthy, A. (2013). Intermediary asset pricing. *The American Economic Review*, 103(2):732–770.

Joyce, M., Lasaosa, A., Stevens, I., Tong, M., et al. (2011). The financial market impact of quantitative easing in the united kingdom. *International Journal of Central Banking*, 7(3):113–161.

Joyce, M., Miles, D., Scott, A., and Vayanos, D. (2012). Quantitative easing and unconventional monetary policy – an introduction. *The Economic Journal*, 122(564).

Kaul, A., Mehrotra, V., and Morck, R. (2000). Demand curves for stocks do slope down: New evidence from an index weights adjustment. *The Journal of Finance*, 55(2):893–912.

Krishnamurthy, A. and Vissing-Jorgensen, A. (2011). *The effects of quantitative easing on interest rates: channels and implications for policy*. National Bureau of Economic Research.

Krishnamurthy, A. and Vissing-Jorgensen, A. (2013). The ins and outs of LSAPs. In *Federal Reserve Bank of Kansas City's Jackson Hole Symposium on the Global Dimensions of Unconventional Monetary Policy, Jackson Hole, Wyoming, August*.

Ledoit, O. and Wolf, M. (2004). A well-conditioned estimator for large-dimensional covariance matrices. *Journal of multivariate analysis*, 88(2):365–411.

Miles, D. and Schanz, J. (2014). The relevance or otherwise of the central bank's balance sheet. *Journal of International Economics*, 92:S103–S116.

Mitchell, M., Pulvino, T., and Stafford, E. (2004). Price pressure around mergers. *The Journal of Finance*, 59(1):31–63.

Modigliani, F. and Sutch, R. (1966). Innovations in interest rate policy. *The American Economic Review*, 56(1/2):178–197.

Nakamura, E. and Steinsson, J. (2018). High-frequency identification of monetary non-neutrality: The information effect. *The Quarterly Journal of Economics*, 133(3):1283–1330.

Neely, C. J. et al. (2010). *The large scale asset purchases had large international effects*. Federal Reserve Bank of St. Louis, Research Division.

Petajisto, A. (2009). Why do demand curves for stocks slope down? *Journal of Financial and Quantitative Analysis*, 44(5):1013–1044.

Petajisto, A. (2011). The index premium and its hidden cost for index funds. *Journal of Empirical Finance*, 18(2):271–288.

Schnitzler, J. (2016). S&P 500 inclusions and stock supply. *Working Paper*.

Scholes, M. S. (1972). The market for securities: Substitution versus price pressure and the effects of information on share prices. *The Journal of Business*, 45(2):179–211.

Shleifer, A. (1986a). Do demand curves for stocks slope down? *The Journal of Finance*, 41(3):579–590.

Shleifer, A. (1986b). Do demand curves for stocks slope down? *The Journal of Finance*, 41(3):579–590.

Swanson, E. T. (2011). Let's twist again: a high-frequency event-study analysis of operation twist and its implications for QE2. *Brookings Papers on Economic Activity*, 2011(1):151–188.

Tobin, J. (1969). A general equilibrium approach to monetary theory. *Journal of money, credit and banking*, 1(1):15–29.

Vayanos, D. and Vila, J.-L. (2009). *A preferred-habitat model of the term structure of interest rates.* National Bureau of Economic Research (No. w15487).

Wurgler, J. and Zhuravskaya, E. (2002). Does arbitrage flatten demand curves for stocks? *The Journal of Business*, 75(4):583–608.

Appendix

A Additional Material: Figures and Tables

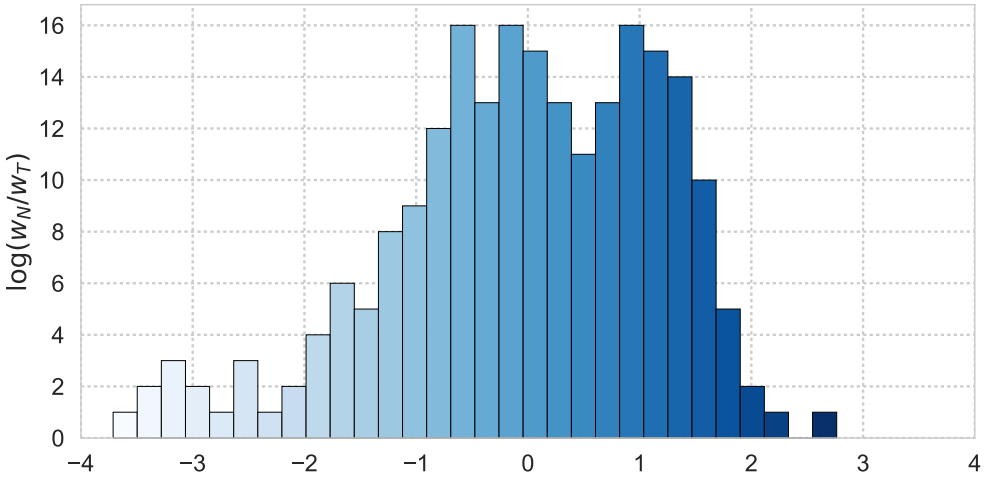


Figure A.1: Distortion. The figure plots the distribution of the log ratio between the Nikkei weight w_N and the TOPIX weight w_T for Nikkei firms only. The histogram shows a significant dispersion, confirming that Nikkei weights induce significant cross-sectional variation of purchased quantities relative to market capitalization.

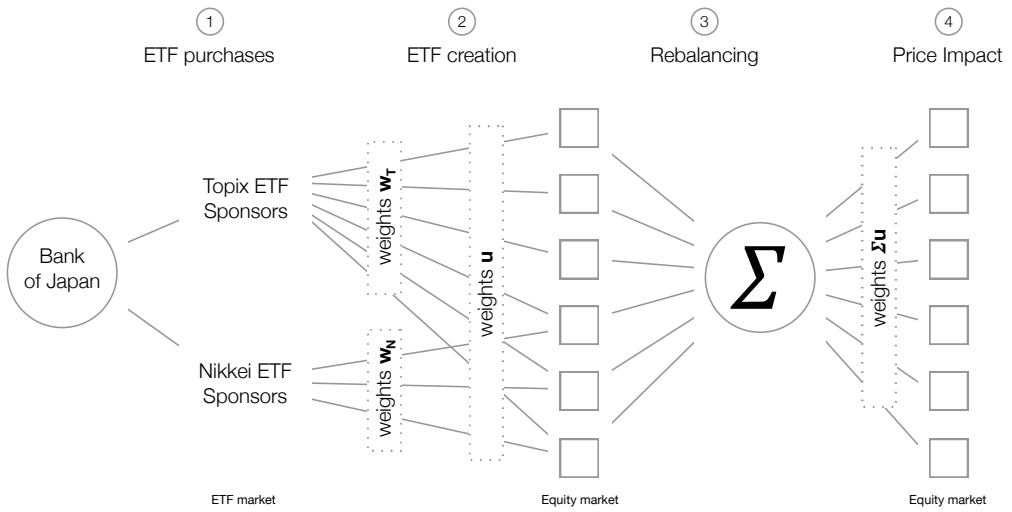


Figure A.2: From ETFs to equity This figure describes the channel though which ETF purchases of the central bank may have an impact on equity prices. As the BoJ buys TOPIX- and Nikkei-linked ETFs, these are created by ETF sponsors and/or authorized participants. The securities needed to form the ETF basket are collected by these intermediaries in the equity market, thus effectively reducing the supply of equity shares available to private investors.

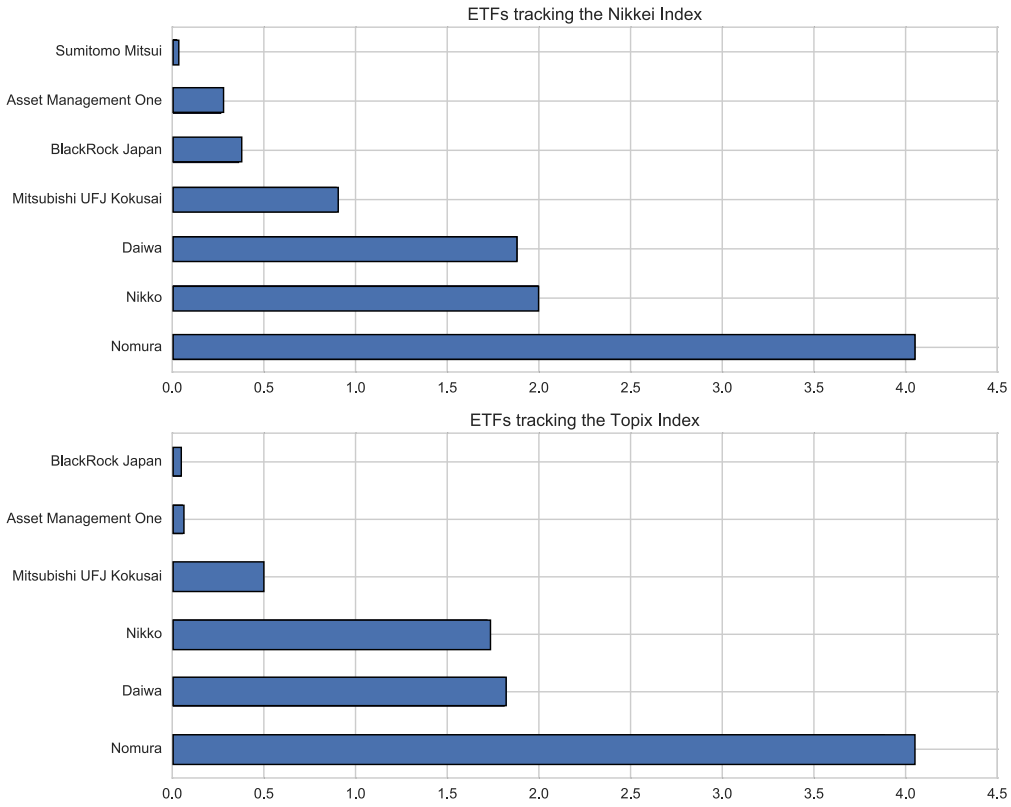
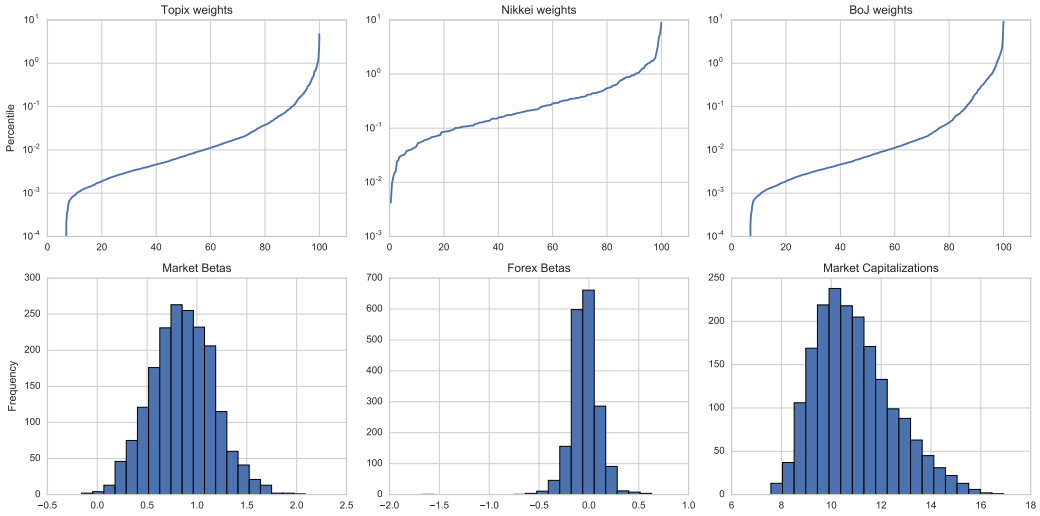


Figure A.3: Assets Under Management (AUM) by Provider (in trillion yen). This figure shows the Assets Under Management of ETFs aggregated at Provider level. The values are computed as of December 30th, 2016.

Panel A – 2014 Event



Panel B – 2016 Event

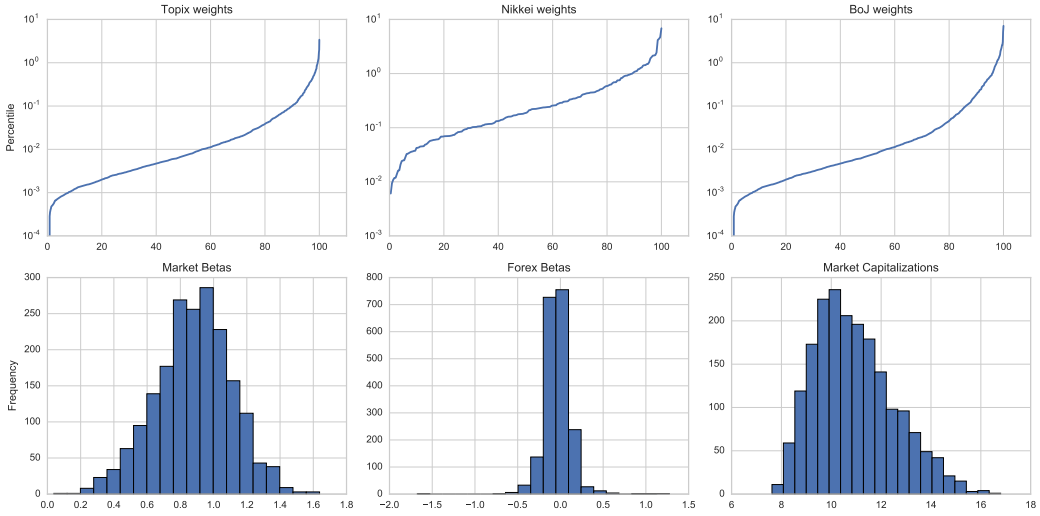


Figure A.4: Weights, betas and market capitalizations. The plots display cross-sectional heterogeneity of the variables of interest at the time of the BoJ announcements. Panel A refers to the announcement in 2014, Panel B to the announcement in 2016. The first row of each panel plots the percentile functions in logarithmic scale of the TOPIX weights (ω_T), the Nikkei weights (ω_N) and the BoJ weights(q). BoJ weights are computed as $\omega_T + \omega_N$ and correspond to the elements of the vector q in the model. The second row of each panel shows the distribution of stock-level market betas, Forex betas and market values. Market betas and Forex betas are estimated following the procedure explained in Section 4.3.3. Companies market capitalizations are in logs.

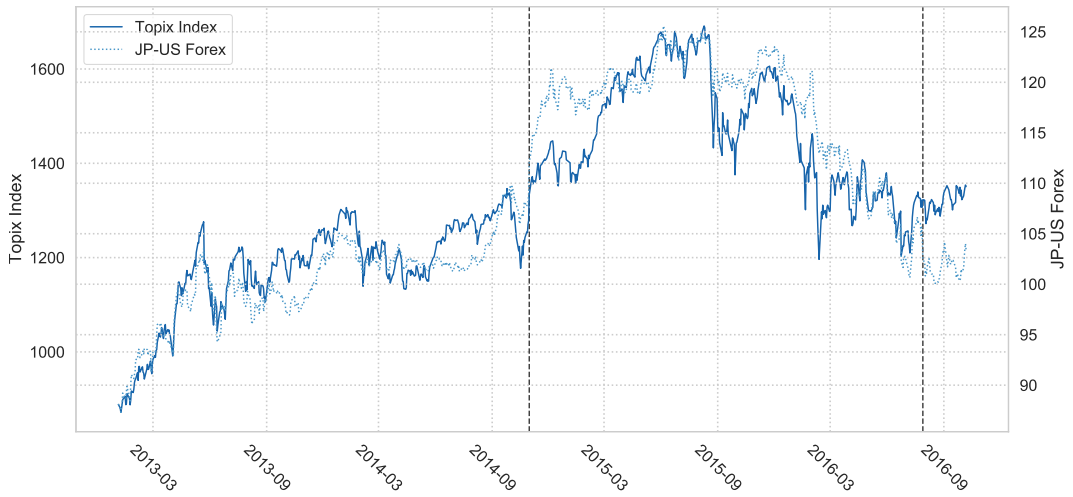


Figure A.5: TOPIX Index and JP-US Exchange Rate. This figure shows the time-series of the TOPIX Index over our sample period (green solid line, left axis) and of the exchange rate from US Dollar to Japanese Yen (purple dotted line, right axis).

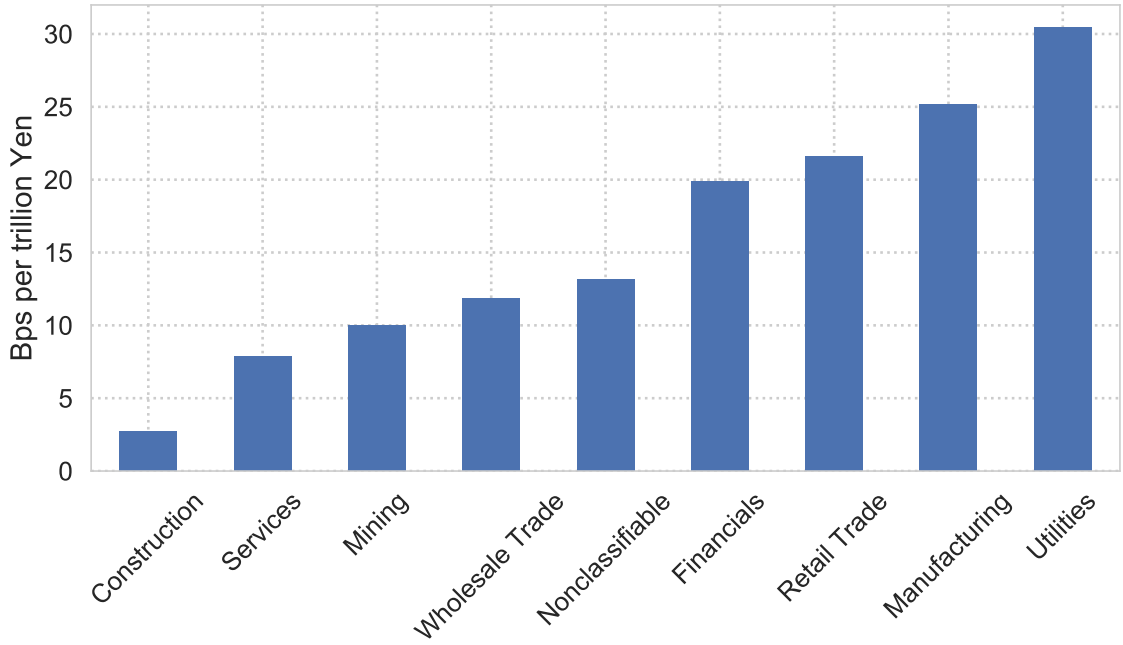


Figure A.6: Portfolio Balance Effect across Industries This figure shows the estimated portfolio balance impact of the policy, expressed in basis points per trillion Yen, computed separately for each sector.

	Mean	Std Deviation	Min	25%	50%	75%	Max	Obs
Market Cap (Billions Yen)								
TOPIX	252	853	2	17	45	148	22210	3824
Nikkei 225	1352	2110	28	294	683	1522	22210	442
Not Nikkei 225	108	251	2	15	35	94	4434	3382
Forex Beta								
TOPIX	-0.04	0.15	-1.68	-0.12	-0.04	0.04	1.27	3824
Nikkei 225	0.02	0.13	-0.39	-0.07	0.02	0.10	0.46	442
Not Nikkei 225	-0.05	0.15	-1.68	-0.12	-0.05	0.03	1.27	3382
Market Beta								
TOPIX	0.87	0.27	-0.16	0.69	0.88	1.05	2.08	3824
Nikkei 225	1.05	0.19	0.46	0.90	1.04	1.18	1.71	442
Not Nikkei 225	0.85	0.27	-0.16	0.67	0.85	1.02	2.08	3382
BoJ Weight								
TOPIX	0.05	0.22	0.00	0.00	0.00	0.01	4.65	3824
Nikkei 225	0.37	0.53	0.01	0.10	0.20	0.43	4.65	442
Not Nikkei 225	0.01	0.03	0.00	0.00	0.00	0.01	0.42	3382
Nikkei 225 Weight								
TOPIX	0.05	0.31	0.00	0.00	0.00	0.00	8.91	3824
Nikkei 225	0.45	0.82	0.00	0.09	0.20	0.45	8.91	442
Not Nikkei 225	0.00	0.00	0.00	0.00	0.00	0.00	0.00	3382
TOPIX Weight								
TOPIX	0.05	0.18	0.00	0.00	0.01	0.03	4.70	3824
Nikkei 225	0.30	0.44	0.01	0.06	0.15	0.36	4.70	442
Not Nikkei 225	0.02	0.05	0.00	0.00	0.01	0.02	0.85	3382

Table A.1: Summary Statistics. This table provides summary statistics for various stock characteristics by index membership. TOPIX stocks represent our entire sample of stocks. Nikkei stocks are those included in the Nikkei 225 index, while Not Nikkei stocks are those that only appear in the TOPIX index. All Nikkei companies also belong to the TOPIX index. All statistics are computed using pre-event information and pooling both events together.

Panel A: October 31st, 2014

	Raw Returns				Abnormal Returns			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
π	95.77***	91.16***	39.82***	38.19***	18.13***	13.14**	47.69***	43.85***
	(16.84)	(15.62)	(4.49)	(4.32)	(3.17)	(2.24)	(5.50)	(5.09)
u		0.01***	-0.00	-0.00		0.01***	-0.00	-0.00
		(3.32)	(-0.86)	(-0.73)		(3.58)	(-0.99)	(-0.77)
Market Beta			0.02***	0.01			-0.08***	-0.09***
			(2.70)	(1.38)			(-11.04)	(-12.29)
Forex Beta			0.07***	0.07***			0.07***	0.07***
			(7.58)	(7.45)			(8.30)	(7.88)
log(Market Cap)			0.01***	0.01***			0.01***	0.01***
			(6.87)	(6.33)			(7.13)	(6.60)
Amihud			0.00	0.00			0.00	0.00
			(0.96)	(0.51)			(0.94)	(0.50)
Observations	1,851	1,851	1,807	1,701	1,851	1,851	1,807	1,701
R-squared	0.13	0.14	0.19	0.24	0.01	0.01	0.11	0.17
Industry FE	NO	NO	NO	YES	NO	NO	NO	YES

Panel B: July 29th, 2016

	Raw Returns				Abnormal Returns			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
π	18.00***	16.27***	15.80***	18.59***	18.00***	16.64***	18.59***	20.24***
	(6.29)	(5.57)	(3.59)	(4.10)	(6.11)	(5.52)	(4.10)	(4.33)
u		0.01***	0.01*	0.01**		0.01**	0.01*	0.01**
		(2.75)	(1.82)	(2.00)		(2.10)	(1.80)	(1.97)
Market Beta			-0.00	-0.02*			-0.01	-0.02**
			(-0.57)	(-1.80)			(-0.99)	(-2.22)
Forex Beta			0.02**	0.01			0.02***	0.02*
			(2.06)	(1.42)			(2.63)	(1.94)
log(Market Cap)			0.00	0.00			-0.00	0.00
			(0.99)	(1.11)			(-0.02)	(0.17)
Amihud			0.00	0.00			0.00*	0.00*
			(1.48)	(1.53)			(1.75)	(1.78)
Observations	1,905	1,905	1,839	1,734	1,905	1,905	1,839	1,734
R-squared	0.02	0.02	0.03	0.05	0.02	0.02	0.03	0.05
Industry FE	NO	NO	NO	YES	NO	NO	NO	YES

Table A.2: Robustness: Alternative Covariance Matrix Estimation The tables report results for specifications similar to those in Table 2, but where the main explanatory variable $\pi = \Sigma u$ is constructed using the covariance matrix Σ estimated on raw returns. Regressions of event returns on the predicted price impact π are run separately for the two events. The dependent variable in columns 1-3 is the cumulative raw return, while in columns 4-6 is the cumulative abnormal return with respect to the market model estimated in the pre-event window. Cumulative returns are computed over a 10 days horizon after the announcement date. t-statistics are in parenthesis; asterisks denote conventional significance levels (***=1%, **=5%, *=10%).

	Abnormal Returns 2014						Abnormal Returns 2016					
Horizon (days)	5	10	21	63	126	252	5	10	21	63	126	252
π	11.06**	30.60***	22.51**	63.85***	152.1***	275.0***	15.52**	16.69**	30.48***	34.18**	82.20***	110.4***
	(2.11)	(4.62)	(2.52)	(4.34)	(8.50)	(10.32)	(2.96)	(2.52)	(3.42)	(2.32)	(4.60)	(4.14)
u	-0.00	-0.02***	-0.01	-0.02***	-0.02	-0.04*	0.010*	0.004	0.024**	0.014	0.045**	0.102***
	(-0.39)	(-3.20)	(-1.30)	(-2.52)	(-1.65)	(-1.99)	(2.06)	(0.67)	(2.55)	(1.35)	(2.91)	(4.70)
Market Beta	-0.04*	-0.07*	-0.08*	-0.17**	-0.32**	-0.45***	0.016	-0.01	0.009	0.012	0.007	0.041
	(-1.47)	(-1.97)	(-1.70)	(-2.32)	(-2.22)	(-2.68)	(0.59)	(-0.26)	(0.19)	(0.16)	(0.05)	(0.24)
Forex Beta	0.039**	0.041**	0.095***	0.087**	0.019	-0.12	-0.00	0.014	0.053**	0.049	0.190***	0.166
	(2.81)	(2.32)	(4.11)	(2.26)	(0.30)	(-0.03)	(-0.48)	(0.80)	(2.32)	(1.28)	(2.92)	(0.04)
log(Market Cap)	0.002	0.006	0.001	-0.00	0.005	-0.00	-0.00	-0.00	-0.01	-0.02	-0.06***	-0.10***
	(0.63)	(1.16)	(0.22)	(-0.07)	(0.26)	(-0.72)	(-1.40)	(-0.49)	(-1.60)	(-1.70)	(-3.43)	(-8.54)
Amihud	0.000	4.618	0.001	0.000	0.019***	0.005	0.000	0.000	-0.00	-0.00	-0.00**	-0.01***
	(0.66)	(0.03)	(0.75)	(0.36)	(7.19)	(1.67)	(0.18)	(0.11)	(-0.59)	(-1.31)	(-3.23)	(-3.11)
Observations	1,701	1,701	1,701	1,701	1,701	1,701	1,734	1,734	1,734	1,734	1,734	1,734
R-squared	0.114	0.160	0.102	0.120	0.180	0.141	0.071	0.050	0.101	0.111	0.203	0.191
Industry FE	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES

Table A.3: Cross-sectional regressions with industry fixed effects. The table report the coefficients of cross-sectional regressions of cumulative returns (in percentage points) computed at different horizons on the predicted price impact π and a set of control variables (standardized). In this specification we include industry fixed effects, based on the first 3 digits of the Standard Industry Classification Code (SIC-3). Regressions are run separately for the two events. The dependent variable is the cumulative abnormal return with respect to the market model estimated in the pre-event window. t-statistics from placebo regressions are in parenthesis; asterisks denote conventional significance levels (***=1%, **=5%, *=10%) based on empirical p-values.

	Abnormal Returns 2014						Abnormal Returns 2016					
	5	10	21	63	126	252	5	10	21	63	126	252
π	19.04***	42.17***	28.77***	72.38***	172.5***	300.7***	17.61**	17.34*	31.88***	38.51**	87.69***	112.9***
	(3.22)	(5.69)	(2.86)	(4.57)	(9.25)	(11.60)	(2.98)	(2.34)	(3.16)	(2.43)	(4.70)	(4.36)
u	0.008*	-0.00	-0.00	-0.03***	-0.03**	-0.11***	0.014***	0.001	-0.00	-0.02*	-0.04**	0.001
	(1.90)	(-1.32)	(-0.97)	(-3.44)	(-2.06)	(-4.64)	(3.28)	(0.27)	(-0.24)	(-2.15)	(-2.47)	(0.07)
Market Beta	-0.02	-0.05	-0.07	-0.16*	-0.29**	-0.40***	0.026	0.002	0.027	0.028	0.048	0.076
	(-0.84)	(-1.48)	(-1.37)	(-2.06)	(-2.05)	(-2.14)	(0.86)	(0.07)	(0.51)	(0.36)	(0.33)	(0.41)
Forex Beta	0.037**	0.042*	0.102***	0.097**	0.043	-0.13	-0.00	0.018	0.053*	0.060	0.196***	0.186***
	(2.20)	(2.02)	(3.81)	(2.28)	(0.61)	(-1.72)	(-0.18)	(0.90)	(2.00)	(1.43)	(2.73)	(2.33)
log(Market Cap)	0.001	0.006	0.001	0.001	0.005	-0.00	-0.00	-0.00	-0.01*	-0.03	-0.07***	-0.11***
	(0.33)	(0.94)	(0.16)	(0.06)	(0.24)	(-0.33)	(-1.48)	(-0.55)	(-1.80)	(-1.79)	(-3.58)	(-7.59)
Amihud	0.001	0.000	4.552	0.001	0.016***	0.018***	0.000	0.000	-0.00	-0.00	-0.00**	-0.00
	(0.79)	(0.25)	(0.00)	(0.67)	(6.13)	(7.14)	(0.48)	(0.39)	(-0.53)	(-0.96)	(-3.27)	(-2.45)
Nikkei	-0.01	-0.02*	-0.00	-0.00	-0.00	0.048*	-0.00	0.003	0.038*	0.049*	0.122***	0.153***
	(-1.26)	(-1.72)	(-0.27)	(-0.08)	(-0.18)	(1.54)	(-0.38)	(0.24)	(2.23)	(2.32)	(4.03)	(4.88)
Observations	1,807	1,807	1,807	1,807	1,807	1,807	1,839	1,839	1,839	1,839	1,839	1,839
R-squared	0.06	0.11	0.07	0.10	0.15	0.12	0.05	0.03	0.08	0.08	0.19	0.15
Industry FE	NO	NO	NO	NO	NO	NO	NO	NO	NO	NO	NO	NO

Table A.4: Cross-sectional regressions controlling for Nikkei. The table report the coefficients of cross-sectional regressions of cumulative returns (in percentage points) computed at different horizons on the predicted price impact π and a set of control variables (standardized). In this specification we add a dummy variable *Nikkei* that indicates stocks belonging to the Nikkei 225 Index. Regressions are run separately for the two events. The dependent variable is the cumulative abnormal return with respect to the market model estimated in the pre-event window. t-statistics from placebo regressions are in parenthesis; asterisks denote conventional significance levels (***=1%, **=5%, *=10%) based on empirical p-values.

	Abnormal Returns 2014						Abnormal Returns 2016					
Horizon (days)	5	10	21	63	126	252	5	10	21	63	126	252
π	12.21**	32.64***	22.95**	63.97***	150.7***	266.7***	15.62**	16.37**	28.39**	31.48**	75.75***	102.7***
	(2.28)	(4.80)	(2.51)	(4.38)	(8.72)	(10.01)	(2.91)	(2.41)	(3.11)	(2.16)	(4.38)	(3.86)
u	0.006	-0.00	-0.00	-0.02***	-0.03**	-0.10***	0.012**	0.000	-0.00	-0.02*	-0.04**	-0.00
	(1.35)	(-1.41)	(-1.16)	(-2.89)	(-2.32)	(-4.28)	(2.72)	(0.05)	(-0.56)	(-2.42)	(-2.95)	(-0.21)
Market Beta	-0.04*	-0.07*	-0.08*	-0.17**	-0.33**	-0.46***	0.016	-0.00	0.011	0.015	0.014	0.050
	(-1.48)	(-1.98)	(-1.71)	(-2.38)	(-2.23)	(-2.28)	(0.59)	(-0.25)	(0.24)	(0.21)	(0.10)	(0.25)
Forex Beta	0.040**	0.043**	0.095***	0.087**	0.018	-0.12***	-0.00	0.012	0.044*	0.038	0.163**	0.134***
	(2.83)	(2.38)	(4.11)	(2.25)	(0.28)	(-1.93)	(-0.45)	(0.71)	(1.93)	(0.98)	(2.49)	(2.01)
log(Market Cap)	0.002	0.006	0.001	-0.00	0.004	-0.00	-0.00	-0.00	-0.01	-0.02	-0.06***	-0.10***
	(0.61)	(1.09)	(0.21)	(-0.07)	(0.24)	(-0.81)	(-1.29)	(-0.48)	(-1.61)	(-1.69)	(-3.43)	(-9.10)
Amihud	0.001	0.000	0.001	0.000	0.019***	0.005	0.000	0.000	-0.00	-0.00	-0.00**	-0.01***
	(0.67)	(0.06)	(0.70)	(0.34)	(6.97)	(1.52)	(0.18)	(0.09)	(-0.64)	(-1.32)	(-3.33)	(-3.15)
Nikkei	-0.01	-0.01	-0.00	-0.00	0.011	0.074***	-0.00	0.006	0.039*	0.050*	0.120***	0.143***
	(-1.04)	(-1.39)	(-0.22)	(-0.05)	(0.36)	(2.15)	(-0.17)	(0.46)	(2.24)	(2.29)	(3.62)	(4.17)
Observations	1,701	1,701	1,701	1,701	1,701	1,701	1,734	1,734	1,734	1,734	1,734	1,734
R-squared	0.12	0.16	0.10	0.12	0.18	0.14	0.07	0.05	0.11	0.12	0.21	0.20
Industry FE	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES

Table A.5: Cross-sectional regressions controlling for Nikkei and industry. The table report the coefficients of cross-sectional regressions of cumulative returns (in percentage points) computed at different horizons on the predicted price impact π and a set of control variables (standardized). In this specification we add a dummy variable *Nikkei* that indicates stocks belonging to the Nikkei 225 Index and industry fixed effects based on the first 3 digits of the Standard Industry Classification Code (SIC-3). Regressions are run separately for the two events. The dependent variable is the cumulative abnormal return with respect to the market model estimated in the pre-event window. t-statistics from placebo regressions are in parenthesis; asterisks denote conventional significance levels (***=1%, **=5%, *=10%) based on empirical p-values.

B Model Derivation

The model features a representative investor who chooses time-t demand N_t of shares to maximize its next period exponential utility subject to a standard budget constraint

$$\max_{N} \ \mathrm{E}_t\left(-\exp(-\gamma W_{t+1})\right) \tag{32}$$

$$\text{s.t.} \qquad W_{t+1} = W_t(1+r) + N_t'(p_{t+1} + D_{t+1} - p_t(1+r)) \tag{33}$$

From the first order condition it follows that

$$N_t = \frac{1}{\gamma}[\mathrm{Var}_t(p_{t+1} + D_{t+1})]^{-1}(\mathrm{E}_t[p_{t+1} + D_{t+1} - p_t(1+r)]) \tag{34}$$

We restrict our attention to the covariance stationary equilibrium. Imposing market clearing and substituting $V = \mathrm{Var}_t(p_{t+1} + D_{t+1})$ yields

$$(1+r)p_t = \mathrm{E}_t[p_{t+1} + D_{t+1}] - \gamma V Q_t \tag{35}$$

Iterating forward up to time T and applying the law of iterated expectations we get

$$(1+r)p_t = \mathrm{E}_t\left[\frac{p_T}{(1+r)^{T-t-1}}\right] + \sum_{i=0}^{T-t-1}\frac{D_t}{(1+r)^i} - \gamma V \sum_{i=0}^{T-t-1}\frac{\mathrm{E}_t[Q_{t+i}]}{(1+r)^i} \tag{36}$$

Taking the limit $T \to \infty$ and imposing the no-bubble condition yields

$$p_t = \frac{D_t}{r} - \frac{\gamma V}{(1+r)}\left(\sum_{i=0}^{\infty}\frac{\mathrm{E}_t[Q_{t+i}]}{(1+r)^i}\right) = \frac{1}{r}\left(D_t - \gamma V \Omega_t\right) \tag{37}$$

where we introduced the notation

$$\Omega_t = \frac{r}{1+r}\sum_{i=0}^{\infty}\frac{\mathrm{E}_t[Q_{t+i}]}{(1+r)^i} \tag{38}$$

which can be interpreted as the discounted time-t expected future supply of the assets. This term is crucial for our analysis, representing the channel through which the central bank is able to affect risk premia. Under no expectation of monetary policy intervention we have $\mathrm{E}_t(Q_{t+i}) = Q$, so that the resulting pricing equation collapses to

$$p_t = \frac{1}{r}\left(D_t - \gamma V Q\right) \tag{39}$$

where the vector γVQ can be interpreted as the cross-sectional vector of risk premia required by investors in equilibrium. In our context, this is the pricing equation that applies before the policy announcement at $t = 1$.

In the following sections we look at what happens to prices if the central bank unexpectedly commits itself to a large-scale purchase of assets over a defined period, thus affecting the expected path of future supply Ω_t.

We now solve the model in its most general form, allowing for the possibility that agents expectations on future supply change over time. We assume that for each $t \geq 1$ there exist a scalar $\lambda_t \geq 0$ such that the time-t expectation is

$$\begin{cases} \mathrm{E}_t(Q_{t+i}) = Q & \text{for } i \geq 0 \text{ and } t < 1 \\ \mathrm{E}_t(Q_{t+i}) = Q - \lambda_t(t+i)q & \text{for } i \geq 0 \text{ and } t = 1, \ldots, M \\ \mathrm{E}_t(Q_{t+i}) = Q - \lambda_t M q & \text{for } i \geq M - t \text{ and } t \geq 1 \end{cases} \tag{40}$$

The parameter λ_t can be interpreted as the degree of confidence of investors in the BoJ commitment or, in other words, as the conditional probability they attach to the continuation of the program.

Assuming that investors increase their confidence as time passes – and they observe more actual purchases by the BoJ – amounts to assume that λ_t is increasing in time.

After the BoJ announcement, for $t \geq 1$, the expected supply can be written as

$$\Omega_t = \frac{r}{1+r} \sum_{i=0}^{\infty} \frac{\mathrm{E}_t[Q_{t+i}]}{(1+r)^i} \tag{41}$$

$$= \frac{r}{1+r} \left(\sum_{i=0}^{M-t-1} \frac{Q - \lambda_t(t+i)q}{(1+r)^i} + \sum_{i=M-t}^{\infty} \frac{Q - \lambda_t M q}{(1+r)^i} \right) \tag{42}$$

$$= Q - \frac{\lambda_t r}{1+r} \left(\sum_{i=0}^{M-t-1} \frac{(t+i)q}{(1+r)^i} + \sum_{i=M-t}^{\infty} \frac{Mq}{(1+r)^i} \right) \tag{43}$$

$$= Q - \frac{\lambda_t r}{1+r} \left(\sum_{i=0}^{M-t-1} \frac{(t+i-M)q}{(1+r)^i} + \sum_{i=0}^{\infty} \frac{Mq}{(1+r)^i} \right) \tag{44}$$

$$= Q - \lambda_t M q + \frac{\lambda_t r}{1+r} \sum_{i=0}^{M-t-1} \frac{(M-t-i)}{(1+r)^i} q \tag{45}$$

$$= Q - \lambda_t M q + \lambda_t \varphi(t) q \tag{46}$$

where we introduced the real-valued function

$$\varphi(t) = \frac{r}{1+r} \sum_{i=0}^{M-t-1} \frac{(M-t-i)}{(1+r)^i}, \qquad t \geq 1 \tag{47}$$

This quantity represents the residual duration of the program at time t. In the Internet Appendix we show that, for realistic values of the risk-free rate r, the function $\varphi(t)$ enjoys the following properties:

(i) $\varphi(t+1) - \varphi(t) < 0$

(ii) $\varphi(t) < M$ for $t \geq 1$

(iii) $\varphi(t) = 0$ for $t \geq M$

Given the assumption that by the end of the policy horizon $t = M$ the central bank will have purchased exactly Mq as announced and afterwards it will not engage in further market operations, it follows that the pricing equation (39) takes the form

$$\begin{cases} p_t = \frac{1}{r}\left(D_t - \gamma V Q\right) & \text{for } t < 1 \\ p_t = \frac{1}{r}\left(D_t - \gamma V (Q - \lambda_t M q + \lambda_t \varphi(t) q)\right) & \text{for } t = 1, \dots M \\ p_t = \frac{1}{r}\left(D_t - \gamma V (Q - Mq)\right) & \text{for } t \geq M \end{cases} \tag{48}$$

and the price change at the announcement day $t = 1$ can be written as

$$p_1 - p_0 = \frac{1}{r}\left(\varepsilon_1 + \lambda_1 \gamma V (Mq - \varphi(1) q)\right) \tag{49}$$

Dividing by p_0 coordinate-wise proves Proposition 1. The equation also shows that the size of the price jump is increasing in the initial belief parameter λ_1.

On the days following the announcement, price changes depend on the time-series evolution of λ_t. Denoting the updates in beliefs by $\Delta\lambda_{t+1} = \lambda_{t+1} - \lambda_t$ we have

$$\begin{aligned} p_{t+1} - p_t &= \frac{1}{r}\left(\varepsilon_{t+1} - \gamma V ((\lambda_{t+1}\varphi(t+1) - \lambda_t \varphi(t)) - \Delta\lambda_{t+1} M) q\right) && (50) \\ &= \frac{1}{r}\left(\varepsilon_{t+1} + \gamma \xi(t+1) V q\right), \qquad t = 1, \dots M && (51) \end{aligned}$$

Given $\Delta\lambda_{t+1} > 0$, the following inequalities show that $\xi(t) > 0$

$$\lambda_{t+1}\varphi(t+1) - \lambda_t \varphi(t) < \lambda_{t+1}\varphi(t) - \lambda_t \varphi(t)) = \Delta\lambda_{t+1}\varphi(t) < \Delta\lambda_{t+1} M \tag{52}$$

Therefore we conclude that if $\Delta\lambda_{t+1} > 0$ for every $t = 1, \dots, M$, then we should observe a positive relationship between Vq and the cross-section of price changes. To complete the proof of the first part of Proposition 2 we need to show that this conclusion also applies to the relationship between returns $R_{i,t+1} = (p_{i,t+1} - p_{i,t})/p_{i,t}$ and $\pi = \Sigma u$. This follows from the definitions of $\Sigma_{i,j}$, u_i and π_i

$$R_{i,t+1} = \frac{1}{r}\left(\varepsilon_{i,t+1}/p_{i,t} + \gamma \xi(t+1)(Vq)_i / p_{i,t}\right) = \frac{1}{r}\left(\varepsilon_{i,t+1}/p_{i,t} + \gamma \xi(t+1)\pi_i\right) \tag{53}$$

Finally taking the expectation of the cumulative returns we get

$$\sum_{s=1}^{t} \mathrm{E}[R_s] = \sum_{s=1}^{t} \frac{1}{r}\left(\gamma \xi(s) \pi\right) = \theta_t \pi \tag{54}$$

where $\theta_t = \sum_{s=1}^{t} \frac{\gamma}{r}\xi(s)$ is a positive and increasing function of t, which follows from $\xi(s) > 0$ for $s = 1, \dots, M$ as shown above. This concludes the proof of Proposition 2.

C Systematic Risk in the Model

In our model the systematic risk of security i is measured as $(VQ)_i$, where V is the covariance matrix of price innovations and Q is the vector of shares outstanding. This quantity represents the covariance

of the security's price changes ε with the change in the wealth of the representative agent (i.e. the value of the market portfolio) and it therefore admits an interpretation similar to the market beta. Denoting the value of the market portfolio by MP and the covariance of the price of stock i with MP by $\beta(P)_i$ we have

$$\beta(P)_i = \mathrm{Cov}(\Delta MP, \Delta p_i) \propto \mathrm{Cov}(\varepsilon_t' Q, \varepsilon_{i,t}) = (VQ)_i \tag{55}$$

Market betas are usually defined in terms of returns, not of price changes. Thus an empirically more relevant definition of the systematic risk of security i is given by $(\Sigma W)_i$, where Σ is the covariance matrix of returns and W is the vector of percentage weights of the market portfolio. This quantity is proportional to the market beta of stock i, denoted by $\beta(R)_i$

$$\beta(R)_i = \frac{\mathrm{Cov}(R^{mkt}, R_i)}{\mathrm{Var}(R^{mkt})} \propto \mathrm{Cov}(R'W, R_i) = \sum_{i,j} \mathrm{Cov}(R_j, R_i) W_j = (\Sigma W)_i \tag{56}$$

Let $Mq = Q^{\mathrm{post}} - Q$ denote the announced change in the supply of assets and β^{post} the implied vector of market beta after the announcement. It follows immediately from the above definitions that the change in (price-level) systematic risk is proportional to the product between V and q:

$$\beta(P)_i^{\mathrm{post}} - \beta(P)_i \propto (VQ^{\mathrm{post}} - VQ)_i \propto -(Vq)_i \tag{57}$$

Similarly, from the definition of $\pi = \Sigma u$, where $u_i = p_i q_i$ is the announced change in the supply of stock i expressed in yen, it follows that

$$\beta(R)_i^{\mathrm{post}} - \beta(R)_i \propto (\Sigma W^{\mathrm{post}} - \Sigma W)_i \propto -(\Sigma u)_i = -\pi_i \tag{58}$$

For each stock i, π_i can thus be interpreted as the change in the stock beta (i.e. the systematic risk) induced by the supply shock. Notice that this change is induced by the policy through a modification of the portfolio held by the representative agent, while the fundamental covariance structure of returns is unchanged.