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

Andrea Barbon
USI Lugano and Swiss Finance Institute

Virginia Gianinazzi
USI Lugano and Swiss Finance Institute

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 premiums. 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. (*JEL* E52, E58, B26)

Received XXXX XX, XXXX; editorial decision XXXX XX, XXXX by Editor XXXXXXXXXXXX.

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 assets and

This work was supported by the Swiss Finance Institute. We thank Francesco Franzoni and Alberto Plazzi for their advice and 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, and 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. Send correspondence to Virginia Gianinazzi, USI Lugano and Swiss Finance Institute, Via Buffi 13, Lugano, Switzerland 6904; telephone: +41797267614. E-mail: virginia.gianinazzi@usi.ch. All errors are our own.

Published by Oxford University Press on behalf of The Society for Financial Studies 2019.
doi:10.1093/rapstu/Sample Advance Access publication July 3, 2019

liabilities of central banks exposes them to considerable risks and raises questions about the consequences of a potential exit from QE.

Substantial evidence indicates that central banks' asset purchases can have an economically significant impact on yields in targeted markets (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 LSAPs, the debate about the mechanisms linking asset purchases to asset prices and the persistence of the impact is still ongoing. Even though the idea of *easing through quantity* relies on the view that large purchases by the central bank reduce assets' risk premiums, 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 exchange-traded funds (ETFs). By the end of 2016, the BoJ owned approximately ¥14 trillion worth of TOPIX and Nikkei ETFs, or more than 2.5% of the total market capitalization. This unprecedented equity operation has the declared objective of lowering risk premiums and reducing the cost of equity capital of Japanese companies (BoJ, 2013).

The BoJ is the first central bank to engage in large purchases of equity securities. Given the large cross-section of stocks, this intervention represents a unique laboratory to test how QE affects asset prices and an opportunity to shed some light on the long-standing debate on the elasticity of long-term demand curves for stocks.

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 asset pricing model that generalizes the idea of the portfolio balance channel to the case of equities.[1] The key implication that we derive from the model is that

[1] We can easily 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.

the change in systematic risk of each stock is determined by (1) the entire vector of central bank purchases and (2) 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 in 2014 and 2016. We document that both announcements produced a highly heterogeneous response of equity prices at stock level. Figure 1 plots the cumulative returns, around the 2014 announcement, of two portfolios formed by ranking stocks based 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 1-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 purchased. Given a total equity market capitalization of roughly ¥500 trillion, this implies an elasticity close to 1, that is, each yen invested translates into an increase in total market valuation of roughly 1 yen. Our estimate is in line with those provided by Shleifer (1986a) and Petajisto (2011), who find an elasticity of 1 and 0.84, respectively, 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. Wurgler and Zhuravskaya (2002) and Kaul et al. (2000) find instead flatter demand curves, with estimated elasticities of 8.24 and 10.5, respectively.

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-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, because 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.

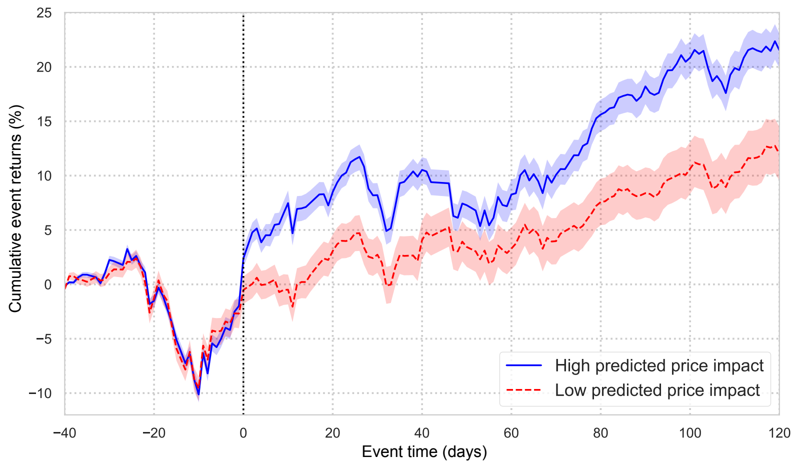

Figure 1
Heterogeneity of event returns. This figure shows the time series of the cumulative returns around the BoJ announcement of October 2014, of two portfolios of firms ranked by the model predicted returns. The blue line represents the average for first quartile of the distribution (firms with the highest predicted price impact), and 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.

The nonfundamental 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 LSAPs and index redefinitions lies in the nature of the supply shock. When buying assets, a central bank is effectively transferring a portion of fundamental risk from the private sector to its balance sheet and holding 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 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. We consider an economy with multiple assets in finite supply and a constant absolute risk aversion (CARA) utility representative agent that maximizes her wealth in each period. We introduce QE 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 premiums. Crucially, we extend the model to an infinite horizon to relax the assumption that uncertainty about firms' fundamentals 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 nontrivial abnormal returns are still observed 1 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, a post-event drift can arise in equilibrium because of two reasons. First, purchases are scattered over various dates, so the model predicts that prices continue to adjust also after the announcement because of 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 quantitatively small relative to the initial price jump. 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 beliefs about the size of the purchase program are assumed to increase over time, consistent 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 from the BoJ may give rise to order imbalances, thus pushing prices upward 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 purchases rather than the change in the supply of assets, a channel quite distinct from portfolio balance. As

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

in D'Amico and King (2013), we will refer to this effect as the *flow effect* of the policy. 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, ending 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 reestimate 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 5% and 10% 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 premiums 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.

1. 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 LSAP program committing itself to buy large quantities of broad market equity ETFs, with the declared view of lowering risk premiums (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 quarterly ETF purchases, as shown in Figure 2. The time series of aggregate ETF purchases is publicly available at daily frequency on the BoJ Web site starting from December 2010.

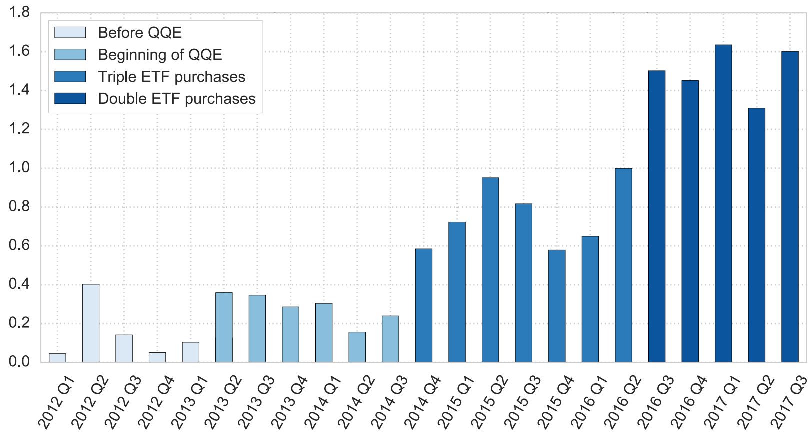

Figure 2
Quarterly ETF purchases of the Bank of Japan in trillion yen Changes in the bar color indicate changes in the policy target purchase amounts. In the first phase the target was set to ¥1 trillion per year; in the second phase it was tripled to ¥3 trillion; and in the third phase it was doubled to ¥6 trillion. Data come from the BoJ Web site.

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 growth is even more remarkable if we consider that the ETF industry in Japan almost tripled in value between 2013 and 2016. The ETF program is comparable in size 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

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

[4] On November 19, 2014, the BoJ also started buying ETFs tracking the JPX-Nikkei 400 Index, so that approximately 43% of the purchases were 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.

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 underweighting in the BoJ portfolio is clearly more severe for companies not included in the Nikkei index. However, the allocation of capital across Nikkei companies is highly heterogeneous 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 resultant 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 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 significantly deviates 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

Figure 3
BoJ purchases and ETF inflows The left axis plots 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. The right axis shows the cumulative return from 2010 of the TOPIX index (gray dashed line).

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. Figure A.2 in Appendix A visualizes this mechanism. Given the direct correspondence, in the rest of the paper we will consider ETF purchases

Table 1
BoJ Indirect shareholdings

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

Summary statistics on indirect ownership by the BoJ for the ten companies with the highest BoJ share. BoJ flow is the cumulative compounded BoJ purchases at the 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.

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 Web site of the Japan Exchange Group (JPX). We then get daily data on AUM for each ETF from Bloomberg.[5] 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. Moreover, this implies that the purchases by the BoJ have consistently triggered creation of new ETF shares.

It must be noted that the bias toward 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

[5] Figure A.3 in Appendix A illustrates the distribution of AUM by index and ETF provider.

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 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). However, a dearth of empirical evidence exists 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 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. 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)).[6] 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

[6] 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 proposed by Culbertson (1957) and Modigliani and Sutch (1966).

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; Schnitzler, 2016; Scholes, 1972; Shleifer, 1986b). The general finding is that large nonfundamental 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.[7]

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 (a) show that LSAPs by central banks have effects beyond those due to portfolio balance and (b) provide evidence of alternative transmission channels consistent with flat demand curves. For the case of purchases of long-term bonds, Krishnamurthy and Vissing-Jorgensen (2011) provide 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 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 premiums 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

[7] 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 her own assets and those held by the central bank.

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. Because in equilibrium 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 driven in the model by 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 $Q=(Q^1,\dots,Q^n)$, paying dividends in every time period. The dividend $D_{i,t}$ paid at time t is

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

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

$$\max_N \mathrm{E}_t(-\exp(-\gamma W_{t+1})) \tag{2}$$

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

where W_t is the total wealth, N_t' denotes the transpose of the vector N_t, and γ denotes the aggregate risk aversion. At date $t=1$ the central bank announces share purchases described by the vector $q=(q^1,\dots,q^n)$, 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 $t=1,\dots,M$ 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 nonstochastic 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.[8] 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, because the BoJ never announced such a stringent commitment. Still, given that the BoJ position has 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

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

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. Risk-free returns are uncorrelated with those of Japanese equities, so omitting the risk-free asset from the model does not change

[8] 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.

the predicted policy impact on stock prices. This assumption may be interpreted as a form of market segmentation, in that the representative agent cannot reinvest the proceeds in assets outside the Japanese equity universe. We discuss the implications of this assumption in the Internet Appendix.[9]

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

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

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 future asset supply, properly discounted by time, defined as

$$\Omega_t \equiv \frac{r}{1+r} \sum_{i=0}^{\infty} \frac{Q_{t,t+i}}{(1+r)^i}, \tag{6}$$

where $Q_{t,t+i}$ denotes the representative agent's time$-t$ belief about the asset supply available at time $t+i$. Notice that we can assume V constant, because in the model there is no uncertainty about future shocks to supply: assets are in fixed supply and the schedule of purchases by the central bank is deterministic. Equation (6) shows that at any point in time prices reflect the entire path of future asset supply. Given time-discounting, today's prices are less sensitive to quantities further into the future.

Staring at the vector of risk premiums $\gamma V \Omega_t$ 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.[10] This is easier to see in the absence of monetary policy shocks, in which case $V\Omega_t$ reduces to VQ.

Equation (5) shows 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, because what determines

[9] 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 the case of nonzero correlation with equities. We can easily 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.

[10] Although 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. Although the notation becomes messier, the intuition carries through.

systematic risk is the product $V\Omega_t$, and, in the model, central bank purchases affect only the latter term.

We now turn to the beliefs of the representative agent about future asset supply, which enter the pricing equation (5), thus determining the policy price impact. We assume that the central bank intervention is fully unexpected at $t=0$ before the announcement, that is, the representative agent's belief is $Q_{t,h}=Q$ 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. For simplicity we restrict to a family of beliefs parametrized by a scalar function of time λ_t. Formally, we let $\lambda_t\geq 0$ be real numbers such that

$$Q_{t,h}=Q-\lambda_t q_h, \qquad t\geq 1. \tag{7}$$

The function λ_t is assumed to change over time in a deterministic fashion and its path λ_t determines how the representative agent updates her beliefs regarding the size of the purchase program.[11] This formulation provides flexibility to our model in the time dimension, allowing the price impact of the policy to be nonfully contemporaneous to the announcement but, rather, to build up over time. Even though our main identification strategy does not depend on such a feature, this generality provides a framework which allows to better quantify the impact of the policy in Section 5.4. We first solve the model for a general mapping $t\mapsto\lambda_t$ and, then, we present 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.

The model's pricing equation predicts price changes given by

$$p_t-p_{t-1}=\frac{1}{r}(\varepsilon_t+\gamma\xi(t)Vq), \tag{8}$$

where the function $\xi(t)$ is defined piecewise as follows

$$\xi(t)=\begin{cases}0 & \text{if} \quad t\leq 0 \text{ or } t>M\\ \lambda_1(M-\varphi(1)) & \text{if} \quad t=1\\ \Delta\lambda_t M-(\lambda_t\varphi(t)-\lambda_{t-1}\varphi(t-1)) & \text{if} \quad 1<t\leq M\end{cases} \tag{9}$$

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 (9), $\xi(t)=0$ implies that both before the announcement ($t\leq 0$) and after the purchase program has been

[11] Even though we mainly think of λ_t as controlling the agent's beliefs about the central bank actions conditional on time-t information, this reduced form suits a number of nonmutually exclusive interpretations. For instance, 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.

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 if future supply changes are fully predictable in the model, after the policy announcement the shock to supply is immediately impounded into stock prices always only up to the term $\varphi(1)$. This is true also when the average path of future prices is perfectly anticipated by the representative agent ($\lambda_t\equiv 1$). 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.

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 belief λ_1 about future supply because 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 beliefs 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 beliefs 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 implies that the agent revises downward her beliefs on the 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. The Internet Appendix shows simulations of the price dynamics implied by different functional forms for λ_t.

In the special case in which $\lambda_t\equiv 1$, anticipated and realized purchases coincide at any point in time

$$Q_{t,h}=Q-q_h=Q-hq \tag{10}$$

It directly follows from Equation (8) that the price adjustment at $t=1$ is given by

$$p_1-p_0=\frac{1}{r}(\varepsilon_1+\gamma V(Mq-\varphi(1)q)) \tag{11}$$

and the post-event ($t\geq 1$) price changes are determined by

$$p_{t+1}-p_t=\frac{1}{r}(\varepsilon_{t+1}-\gamma V(\varphi(t+1)-\varphi(t))q), \qquad t=1,\dots M \tag{12}$$

Imposing $\lambda_t \equiv 1$ is equivalent to assuming that the representative agent believes the central bank to commit to the announced target exactly. This implies that she also believes the central bank will never unwind its positions neither will engage in additional purchase programs.

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. 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

$$u_i \equiv p_{i,t} q_i, \qquad \forall i \in 1\ldots,n \tag{13}$$

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, that is,

$$\Sigma_{i,j} \equiv \mathrm{Cov}(R_i, R_j), \qquad \forall i,j \in 1\ldots,n \tag{14}$$

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. Appendix B provides the proofs.

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

$$\sum_{s=1}^{t} \mathrm{E}[R_s] = \theta_t \pi \tag{15}$$

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 (9).

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, that is, 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 data.

As summarized in Proposition 2, the model predicts post-event returns in the same direction of event returns, that is, proportional to π, until the purchase target is met at $t = M$. This generates a post-event drift the magnitude of which depends on both the value of the risk-free rate and the beliefs dynamics parametrized by λ_t. The Internet Appendix shows analytically and from model simulations that for realistic values 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 beliefs 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 premiums, and hence of prices. Unless the central bank unwinds its positions, prices will not revert to the pre-event level.[12] 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 2,000 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

[12] Notice that a reversal would be observed as soon as investors update their beliefs about 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. We do not provide any empirical evidence of an exit from LSAP or formalize this scenario into a proposition.

from Japan Macro Advisors Inc. The time series of ETF purchases by the BoJ is publicly available at daily frequency on its Web site.

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

$$R_{i,e}^{H} = \alpha_e + \beta_e^H \pi_{i,e} + \delta_e' W_{i,e} + \eta_{i,e} \tag{16}$$

where R_i^H is the cumulative return of stock i computed over H days from the event day and W is a matrix of stock-level 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.

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. 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 (1 month, 3 months, 6 months, and 1 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 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.

The identifying assumption behind this strategy is that there is no transmission mechanism of monetary policy other than portfolio balance that would affect prices proportionally to π. If the central bank was buying according to market weights, π would be parallel to market betas, and this claim would be difficult to make, as we know from the literature that asset purchases can affect stock prices through multiple channels, possibly in proportion to their exposure to market risk. However, because the BoJ is tilting its purchases away from market capitalization, Proposition 1 predicts that the shock should leave a characteristic footprint in the cross-section of returns and abnormal returns, which makes the assumption more likely to be true.

4.3 Variable construction

This section 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.

4.3.1 BOJ purchases. Going to the data, we assume that the allocation rule between TOPIX and Nikkei ETFs 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

$$u_i = Tw_{i,T} + Nw_{i,N} \tag{17}$$

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.

Because $T \cong N$ 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 resultant 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. 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 resultant 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

$$R_{i,t} = \alpha_i + \beta_i^{mkt} R_{mkt,t} + e_{i,t} \tag{18}$$

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.

5. Empirical Results

5.1 Event study

As a preliminary test of Proposition 1 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 that might explain the heterogeneous response to the announcements and thus the divergence in returns. Consistent with the fact that the policy is heavily skewed toward Nikkei companies, which are on average larger than non-Nikkei ones, we find a positive correlation between π and market capitalization.

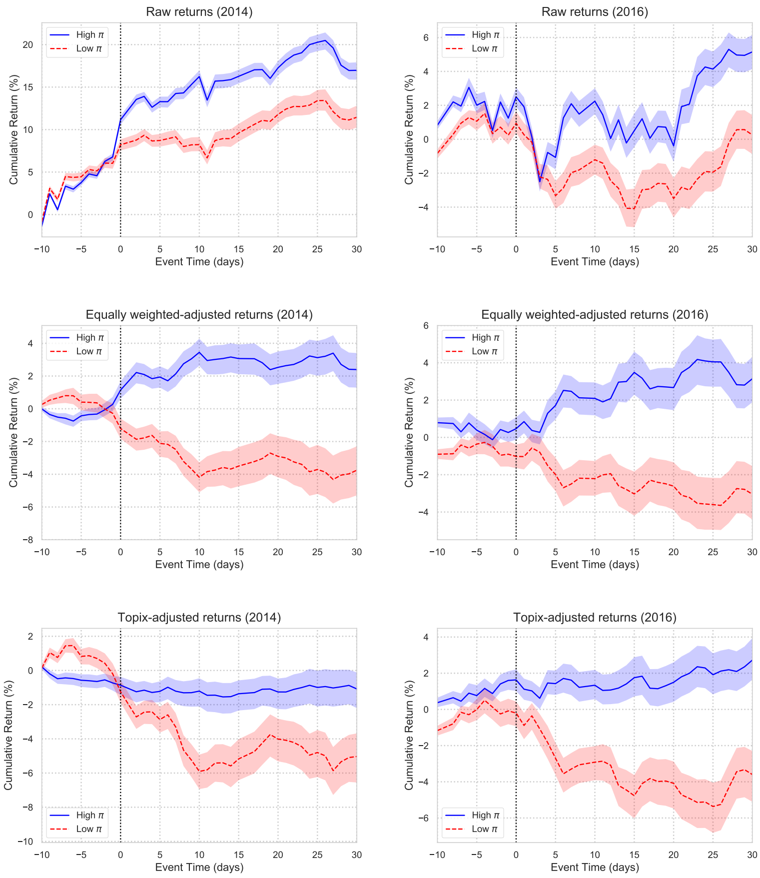

Figure 4
Cumulative returns of high versus low π stocks (in percentage) This figure shows the time series of the mean cumulative returns around the BoJ announcements of stocks with high predicted price impact π against that of low π stocks. The plots on the left refer to the announcement on October 31, 2014, whereas those on the right show the reaction to the announcement on July 29, 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 1 year. An equally weighted portfolio of stocks in the TOPIX universe is used a proxy for the market portfolio in the middle panels, and the return of the TOPIX index is used in the bottom panels. The blue line represents the average for the first quartile of the distribution (firms with the highest predicted price impact), and 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.

We also find a weak negative correlation between π and the exposure to the forex market. If the yen depreciated as a consequence of the announcement, we might spuriously observe returns proportional to π. Therefore, we estimate the model in Equation (16), including controls for market capitalization and forex beta. We also control for market beta, illiquidty (Amihud ratio) and companies' weight in the BoJ purchase schedule (u). The forex beta is the coefficient on the daily percentage change in the exchange rate from U.S. dollar to Japanese yen estimated from stock-by-stock time-series regressions where the dependent variable is the daily stock return and which additionally include daily market returns on the right-hand side. Both the forex beta and the market beta are estimated over a window of 1 year, ending 2 trading weeks before each BoJ announcement. The Internet Appendix provides summary statistics of the control variables by quartile of π.

Table 2 reports regression results. Panel A shows the estimated effect of the 2014 announcement, and panel B shows that of the 2016 announcement. The dependent variable is the cumulative return over 10 trading days after the announcement. In Columns 1–3 we use cumulative raw returns, and in Columns 4–6 we consider cumulative abnormal returns from a market model and 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 nonoverlapping periods to ensure that the empirical distribution is constructed from independent draws.[13]

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 Column 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, which might be due to the smaller change in the target purchase amount. Still, the coefficient

[13] We report empirical p-values computed using this methodology in all tables that show estimates of specifications of the model in Equation (16). For regressions involving short-horizon returns (up to 3 months), we exclude placebo event periods that do not include BoJ meetings with important monetary policy announcements. Specifically, we exclude the meetings on February 1, 2013, March 25, 2013, and June 18, 2012, and the announcement of the post-tsunami intervention on March 14, 2011.

Table 2
Cross-sectional regressions

Panel A: October 31st, 2014						
	Raw returns			Abnormal returns		
	(1)	(2)	(3)	(4)	(5)	(6)
π	57.86***	59.15***	31.92***	36.06***	37.75***	39.95***
	(8.59)	(8.00)	(4.40)	(4.94)	(4.81)	(5.54)
u		-0.00	-0.02***		-0.00	-0.02***
		(-0.95)	(-3.83)		(-1.34)	(-3.62)
Market beta			0.040			-0.05
			(0.75)			(-1.48)
Forex beta			0.040*			0.040*
			(1.87)			(1.96)
log(market cap)			0.007			0.005
			(1.17)			(0.97)
Amihud			0.000			0.000
			(0.53)			(0.22)
Obs	1,851	1,851	1,807	1,851	1,851	1,807
R^2	.106	.106	.162	.046	.047	.108
Industry FE	NO	NO	NO	NO	NO	NO

Panel B: July 29th, 2016						
	Raw returns			Abnormal returns		
	(1)	(2)	(3)	(4)	(5)	(6)
π	14.07*	14.30*	12.17*	15.33*	16.33*	17.49**
	(2.09)	(1.93)	(1.68)	(2.10)	(2.08)	(2.43)
u		-0.00	0.001		-0.00	0.004
		(-0.29)	(0.19)		(-1.33)	(0.56)
Market beta			0.006			0.002
			(0.12)			(0.06)
Forex beta			0.016			0.019
			(0.78)			(0.94)
log(market cap)			-0.00			-0.00
			(-0.10)			(-0.58)
Amihud			0.000			0.000
			(0.43)			(0.41)
Obs	1,905	1,905	1,839	1,905	1,905	1,839
R^2	.017	.017	.021	.019	.019	.028
Industry FE	NO	NO	NO	NO	NO	NO

The table reports 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, and in Columns 4–6 the dependent variable is the cumulative abnormal return with respect to the market model estimated in the pre-event window. Cumulative returns are computed over a 10-day horizon after the announcement date. t-statistics from placebo regressions are in parentheses. *$p<.1$; **$p<.05$; ***$p<.01$.

on π is positive and significant at the 10% confidence level in most specifications.

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 Columns 2 and 5 of each panel. The coefficient on u turns however negative and significant in Columns 3 and 6 of panel A. We find that this is due to the inclusion of the control for market capitalization, because u and market cap are highly correlated, as one would 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 3 of panel A especially, the coefficient on π drops significantly. This is, as expected, because 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 and 6 the number of observations drops slightly because of missing data on trading volume needed to estimate the Amihud ratio.

5.3 Time-series pattern

We now turn to the long-run predictions of the model summarized in Proposition 2. We estimate the cross-sectional model specified in Equation (16) at different horizons H over which cumulative returns are calculated. Table 3 reports the results. 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 the 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 1 year after the event. The absence of reversal is a key prediction of the portfolio-balance channel. Because the shock to supply induced by the BoJ has a unique cross-sectional shape, the observed effect is unlikely due to shocks other than the purchase program, suggesting a causal effect of the policy. Still, we cannot completely rule out other policy transmission mechanisms. In Section 6, we address the possibility that (part of) the effect might be due to continuous price pressure that prevents prices from reverting to the pre-announcement level.

Table 3
Cross-sectional regressions over different horizons

	Abnormal returns 2014						Abnormal returns 2016					
	5	10	21	63	126	252	5	10	21	63	126	252
π	17.78**	39.95***	28.31***	72.21***	171.9***	305.5***	17.43**	17.49**	33.70***	40.86**	93.52***	120.2***
	(3.09)	(5.54)	(2.89)	(4.52)	(8.97)	(11.91)	(3.03)	(2.43)	(3.44)	(2.56)	(4.88)	(4.69)
u	-0.00	-0.02***	-0.01	-0.03***	-0.03**	-0.07***	0.011**	0.004	0.026**	0.016	0.052***	0.118***
	(-0.30)	(-3.62)	(-1.21)	(-3.20)	(-2.46)	(-3.43)	(2.22)	(0.56)	(2.76)	(1.48)	(3.31)	(5.13)
Mkt beta	-0.02	-0.05	-0.07	-0.16*	-0.29**	-0.40***	0.026	0.002	0.025	0.025	0.041	0.068
	(-0.84)	(-1.48)	(-1.36)	(-2.02)	(-2.04)	(-2.51)	(0.86)	(0.06)	(0.47)	(0.32)	(0.28)	(0.43)
Forex beta	0.036**	0.040*	0.101***	0.097**	0.043	-0.13	-0.00	0.019	0.061*	0.070*	0.220***	0.216
	(2.17)	(1.96)	(3.83)	(2.29)	(0.61)	(-0.02)	(-0.23)	(0.94)	(2.30)	(1.67)	(3.10)	(0.04)
log(mkt cap)	0.001	0.005	0.001	0.001	0.005	-0.00	-0.00*	-0.00	-0.01*	-0.02	-0.07***	-0.11***
	(0.32)	(0.97)	(0.17)	(0.06)	(0.25)	(-0.27)	(-1.60)	(-0.58)	(-1.81)	(-1.80)	(-3.52)	(-6.82)
Amihud	0.001	0.000	-1.02	0.001	0.016***	0.018***	0.000	0.000	-0.00	-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)	(-2.16)
Obs	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^2	.055	.108	.073	.098	.153	.119	.051	.028	.079	.077	.178	.140

The table reports 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 parentheses. *p <.1; **p <.05; ***p <.01.

A second finding from Table 3 is that the estimated coefficients on π are generally increasing in H.[14] 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 model produces a similar post-event return pattern when beliefs about the size of the program (λ_t) are increasing over time. 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 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.

The Internet Appendix presents some robustness evidence. We show that the observed price impact cannot be explained by industry effects or by Nikkei stocks overperforming non-Nikkei stocks. 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 quantify the aggregate portfolio-balance effect of the BoJ intervention on equity market returns from the coefficient estimated in the cross-section. We then use this quantity to derive an estimate of the aggregate elasticity of equity demand curves. We are aware that the policy might have produced spillovers to other asset classes (this is indeed a prediction of the model), but we are not considering them in this paper.[15]

First, we run our main regression in Equation (16) over the pooled sample, to obtain the average effect of the policy across the two events. We include event fixed effects FE_e to allow for a different intercept across the two announcements. Formally, we estimate the following regression model for each daily horizon $h \in 1,\ldots,252$

$$R^h_{i,e} = \beta^h \pi_{i,e} + \gamma^h X_{i,e} + \delta^h \mathrm{FE}_e + \varepsilon_{i,e} \tag{19}$$

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.

[14] 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.

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

Because we consider both events together, we rescale π to account for the different magnitude of the announcements. In particular, we multiply π_{2014} by 3 and π_{2016} by 6. 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. The third model additionally includes the Amihud ratio as proxy for stock liquidity.

Given $\hat{\beta}^h$ from the estimation, the predicted net return through the portfolio balance channel for security i is $\hat{R}^h_{i,e} = \hat{\beta}^h \pi_{i,e}$.[16] To aggregate the effect at market level, we calculate for each event the predicted market return as the value-weighted sum of security level predicted returns at every horizon

$$\hat{R}^h_e = \hat{\beta}^h \sum_i w_{i,e} \pi_{i,e} \tag{20}$$

We then divide by the capital commitment by the central bank to obtain the induced market return per trillion yen. Considering the 2-year policy horizon, this amounts to ¥6 trillion for 2014 and ¥12 trillion for 2016. Thus, the per yen estimated average market return induced by the policy through the portfolio-balance channel is calculated as

$$\hat{R}^h = \frac{1}{2}\left(\hat{R}^h_{2014}/6 + \hat{R}^h_{2016}/12\right) \tag{21}$$

Table 4 reports results 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, because each yen invested translates into an increase of the market valuation by roughly 1 yen.

Figure 5 plots the time-series evolution of the point estimate for the third specification. Consistent with the qualitative prediction of our model, a momentum-like pattern is visible over the first 100 trading days following the announcement.

The quantification 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 reinvests the proceeds from the sale of stocks to the central bank at the constant risk-free rate.[17] While a violation of this assumption is not a concern for the identification of the portfolio-balance channel, it might bias our estimate

[16] 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.

[17] We thank an anonymous referee for pointing out this issue.

Table 4
Portfolio balance effects

	1 week	1 month	3 months	6 months	1 year
Baseline	3.54	10.10	10.28	25.17	22.32
+ market and forex	2.23	7.72	9.53	22.08	22.45
+ liquidity	1.56	7.17	9.56	19.80	22.05

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 already defined control variables.

of the aggregate effect. In the Internet Appendix, we extend the model to allow the representative agent to reinvest the proceed in a security correlated with the targeted assets and we derive an expression of the bias from incorrectly assuming zero correlation. We find that 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 reinvestment security and the stocks. The sign of the bias is ambiguous and depends on Σ, u, the market weights and the reinvestment security. Therefore, we run simulations of the model using the parameters estimated in the data and assuming different reinvestment securities, namely, S&P 500, 10-year U.S. Treasury bonds, and 10-year JGBs. The estimated bias is positive using the S&P 500 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 underestimating the true elasticity of Japanese equities.

Second, in our calculation we are assuming that the market is reacting to the announced size of the program. If the market believes that the program will be smaller, 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?

An alternative explanation for the observed persistence of the policy impact 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

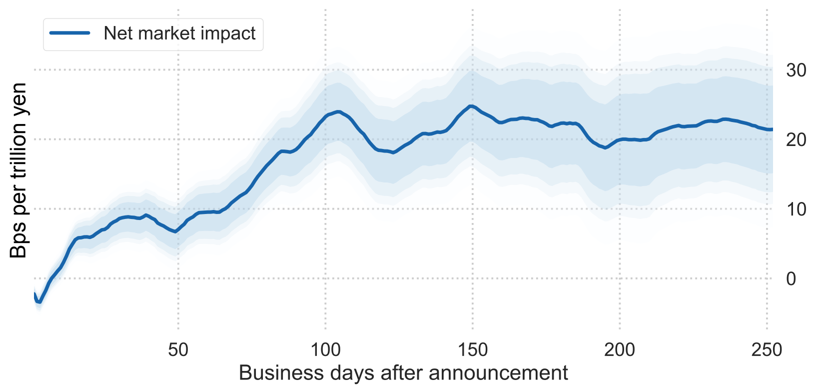

Figure 5
Portfolio balance effects This figure plots the time-series evolution of the estimated portfolio balance effect induced by the BoJ purchase program, expressed in basis points per trillion yen invested. The estimates are based on specification (3), which includes controls for stocks liquidity,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

fundamentals. If this was the case, the absence of reversal could not be interpreted as evidence for long-run, downward-sloping demand curves. In the spirit of D'Amico and King (2013), we will refer to this kind of dynamics as *flow effect* of the policy.

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 rerun 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

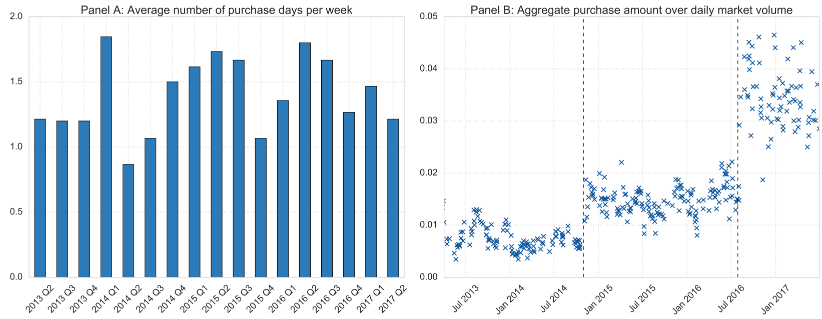

Figure 6
Purchase frequency and volume Panel A plots the average number of purchase days per week at quarterly frequency. Panel B plots the ratio between the yen amount purchased by the BoJ on a given day and the aggregate trading volume in yen on that day. The aggregate trading volume is computed as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

aggregate impact of the policy is unaffected by the timing of the purchases. In the extreme case in which 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 BoJ does not commit 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 represent 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.

Because the purchase frequency remained stable over the policy horizon, the upward revisions of the annual target in October 2014 and 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.

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

$$AR_{i,t} = a + b_0 AV_{i,t} + b_1 AV_{i,t-1} + b_2 \left(\sum_{k=2}^{K} \rho^{k-2} AV_{i,t-k} \right) + e_{i,t} \qquad (22)$$

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

$$AV_{i,t} := \frac{\text{BoJ Flow}_{i,t}}{\text{E}[\text{Volume}_{i,t}]} \qquad (23)$$

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 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.[18] The average daily volume E[Volume$_{i,t}$] is estimated over a backward-looking window of 6 months excluding days in which the BoJ is intervening. On nonpurchase 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

$$F = b_0 + b_1 + b_2 \left(\sum_{k=2}^{K} \rho^{k-2} \right) \qquad (24)$$

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

[18] Underlying securities inherit shocks that occur in the ETF market, both through primary market arbitrage and through the arbitrage that continuously takes place in the secondary market and is carried out by hedge funds and high-frequency traders (Ben-David et al., 2018). Secondary market arbitrageurs profit from 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. They show that most of the price impact seems to take place on the event day and the day after.

effect is exhausted after 2 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, we have reason 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 nonintervention 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 not significantly different from zero in both intervention and nonintervention days supports our claim.

We do not include an announcement dummy in the specification, because no ETF purchases were made by the BoJ on those days. Looking at the time series of BoJ purchases, we see that the bank intervened 2 weeks before and 1 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 4 days after.

We estimate five different specifications of model (22). We start considering only contemporaneous volumes ($K = 0$), then we augment the specification to K equal to 1,2,5 or 10. Panel A of Table 5 reports the estimated parameters, together with the implied long-run impact. The positive coefficients on b_0 and b_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 b_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

Table 5
Flow Effect

	Panel A					Panel B		
K	b_0	b_1	b_2	ρ	F	$\tilde{\beta}$	Flow	Balance
0	0.011				0.011	138.644	5.63%	94.37%
	(9.356)					(11.175)		
1	0.004	0.015			0.019	132.288	9.96%	90.04%
	(3.314)	(11.358)				(10.662)		
2	0.004	0.015	-0.001		0.018	132.608	9.74%	90.26%
	(3.358)	(10.979)	(-0.679)			(10.688)		
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)		
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)		

The table reports results from the estimation of the dynamic model described in Equation (22), 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 $\tilde{\beta}$ 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 $\tilde{\beta}$ with the coefficient $\hat{\beta}$ from Section 5.2 based on standard cumulative absolute returns (CARs).

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

$$\widehat{AR}_{i,t}^{Flow} = \hat{b}_0 AV_{i,t} + \hat{b}_1 AV_{i,t-1} + \hat{b}_2 \left(\sum_{k=2}^{K} \hat{\rho}^{k-2} AV_{i,t-k} \right) \tag{25}$$

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

$$\widetilde{AR}_{i,t} = AR_{i,t} - \widehat{AR}_{i,t}^{Flow} \tag{26}$$

We then estimate our main regression (16) using $\widetilde{AR}$ instead of AR, computing the cumulative abnormal returns over a 1-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 $\tilde{\beta}$ of the price impact of the policy through the portfolio balance channel. The ratio between $\tilde{\beta}$ and its counterpart $\hat{\beta}$ 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. These figures notably represent upper bounds for the persistent flow effect of the policy, an effect that might be amplified by belief 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

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. Doing so 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 premiums 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

$$\Delta k_i = \frac{(V(Q-Mq))_i}{(VQ)_i} - 1 \tag{27}$$

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 premiums 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

$$\Delta k_i^* = \frac{(V(Q-Mq^*))_i}{(VQ)_i} - 1 = \frac{((1-Ma)VQ)_i}{(VQ)_i} - 1 = -Ma \tag{28}$$

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.[19] Buying ETFs tracking the price-weighted

[19] A purchase of $u' = aW_{\text{Topix}}$ in yen corresponds to $u = aQ$ in shares, because the TOPIX is value 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 III in the Internet Appendix shows that the distortion is also evident 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 reshaping its equity portfolio in a value-weighted fashion.

8. Conclusion

In this paper we study asset pricing implications of the ETF purchase program undertaken by the BoJ since April 2013. We build a simple asset pricing model that features multiple assets whose supply is time-varying due to open market operations by a central bank, from which we derive testable predictions about the policy impact on the cross-section and time series of stock returns. We test these predictions in the data around two dates when the BoJ announced an expansion of the purchase target amount. The identification of the portfolio-balance effect exploits the arbitrary nature of the purchase schedule of the BoJ, which is heavily skewed toward stocks in the price-weighted Nikkei 225 index. This creates variation in the shock to supply that is orthogonal to market capitalization and as good as random in the cross-section, which allows us to rule out the effect of confounding factors.

We find that the intervention has a positive and persistent effect on domestic equity prices, thus reducing the cost of equity capital of domestic companies. We show that the response of equity returns to the policy is consistent with a portfolio-balance channel, both in the cross-section and in the time series. The cross-section of returns is significantly explained by our model implied measure of systematic risk change induced by the central bank's purchases. In the time series, we find no evidence of reversal over a 1-year window after each announcement, which is consistent with the prediction of the model that a permanent change in supply should produce a persistent change in risk premiums.

Overall, our results suggest 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.

We also show that the outright purchases of the BoJ generate positive and persistent pressure on prices. This alternative channel might explain at most a minor fraction of the estimated portfolio-balance effect, which remains significant after accounting for the possible flow effect of the policy.

References

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

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

Barberis, N., and R. Thaler. 2003. A survey of behavioral finance. *Handbook of the Economics of Finance* 1:1053–128.

Ben-David, I., F. Franzoni, and R. Moussawi. 2018. Do ETFs increase volatility? *Journal of Finance* 73:2471–535.

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.

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

Brunner, K., and A. H. Meltzer. 1973. Mr. Hicks and the ”Monetarists”. 40:44–59.

Buraschi, A., and P. Whelan. 2015. Bond markets and unconventional monetary policy. *Handbook of fixed-income securities*, first edition.

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

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

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

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

D’Amico, S., and T. B. King. 2013. Flow and stock effects of large-scale treasury purchases: Evidence on the importance of local supply. *Journal of Financial Economics* 108:425–48.

Eggertsson, G. B., and M. Woodford. 2004. Optimal monetary policy in a liquidity trap. Working Paper.

Eser, F., and B. Schwaab. 2016. Evaluating the impact of unconventional monetary policy measures: Empirical evidence from the ECBs securities markets programme. *Journal of Financial Economics* 119:147–67.

Frankel, J. A. 1985. Portfolio crowding-out, empirically estimated. *Quarterly Journal of Economics* 100:1041–65.

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

Greenwood, R. 2005. Short-and long-term demand curves for stocks: Theory and evidence on the dynamics of arbitrage. *Journal of Financial Economics* 75:607–49.

———. 2007. Excess comovement of stock returns: Evidence from cross-sectional variation in Nikkei 225 weights. *Review of Financial Studies* 21:1153–86.

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

Hanson, S. G., and J. C. Stein. 2015. Monetary policy and long-term real rates. *Journal of Financial Economics* 115:429–48.

Harris, L., and E. Gurel. 1986. Price and volume effects associated with changes in the S&P 500 list: New evidence for the existence of price pressures. *Journal of Finance* 41:815–29.

Hau, H., M. Massa, and J. Peress. 2009. Do demand curves for currencies slope down? Evidence from the MSCI global index change. *Review of Financial Studies* 23:1681–717.

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

Joyce, M., A. Lasaosa, I. Stevens, and M. Tong. 2011. The financial market impact of quantitative easing in the United Kingdom. *International Journal of Central Banking* 7:113–61.

Joyce, M., D. Miles, A. Scott, and D. Vayanos. 2012. Quantitative easing and unconventional monetary policy – an introduction. *Economic Journal* 122:F271–F288.

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

Krishnamurthy, A., and A. Vissing-Jorgensen. 2011. The effects of quantitative easing on interest rates: Channels and implications for policy. Working Paper.

———. 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 M. Wolf. 2004. A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis* 88:365–411.

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

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

Modigliani, F., and R. Sutch. 1966. Innovations in interest rate policy. *American Economic Review* 56:178–97.

Nakamura, E., and J. Steinsson. 2018. High-frequency identification of monetary non-neutrality: The information effect. *Quarterly Journal of Economics* 133:1283–30.

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* 44:1013–44.

———. 2011. The index premium and its hidden cost for index funds. *Journal of Empirical Finance* 18:271–88.

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. *Journal of Business* 45:179–211.

Shleifer, A. 1986a. Do demand curves for stocks slope down? *Journal of Finance* 41:579–90.

———. 1986b. Do demand curves for stocks slope down? *Journal of Finance* 41:579–90.

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* 2011:151–88.

Tobin, J. 1969. A general equilibrium approach to monetary theory. *Journal of Money, Credit and Banking* 1:15–29.

Vayanos, D., and J.-L. Vila. 2009. A preferred-habitat model of the term structure of interest rates. Working Paper.

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

Appendix.

A Additional Material: Figures and Tables

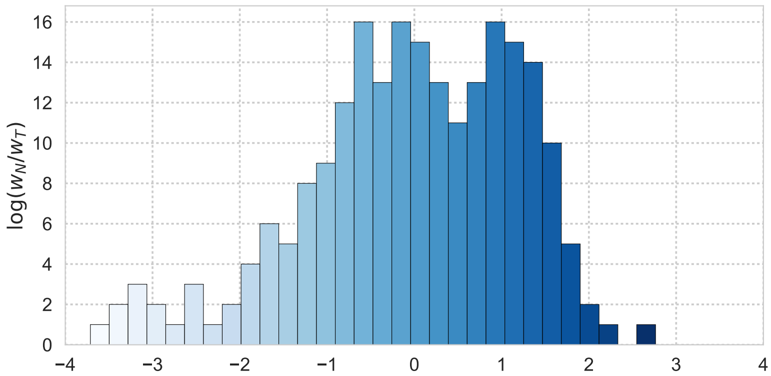

Figure A.1
Distortion The figure plots the distribution of the log ratio between the Nikkei weight w_N and the TOPIX weight w_T for Nikkei firms only. The histogram shows a significant dispersion, confirming that Nikkei weights induce significant cross-sectional variation of purchased quantities relative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

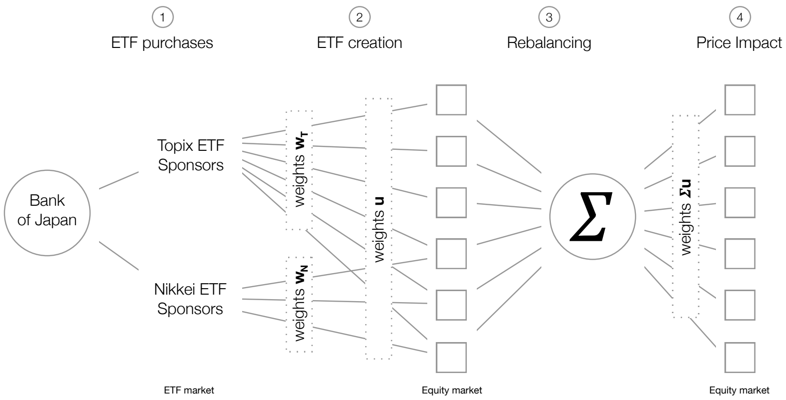

Figure A.2
From ETFs to equity This figure describes the channel though which ETF purchases of the central bank may affect equity prices. As the BoJ buys TOPIX- and Nikkei-linked ETFs, the ETFs are created by ETF sponsors and/or authorized participants. The securities needed to form the ETF basket are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

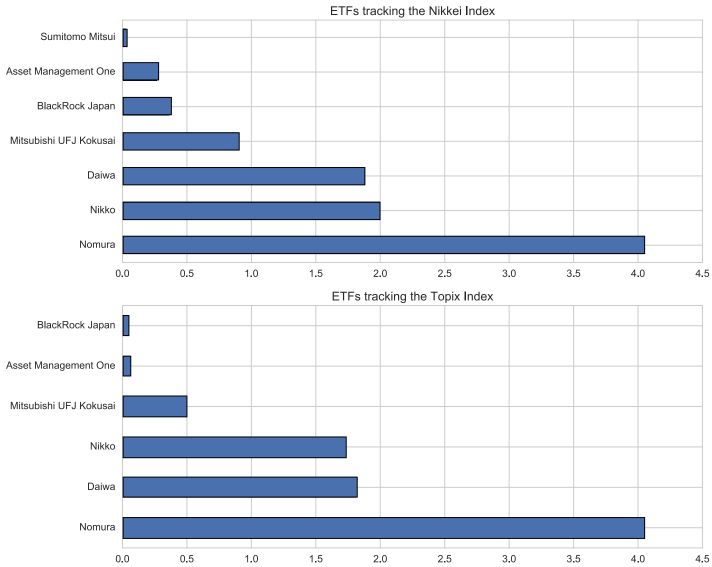

Figure A.3
AUM by provider (in trillion yen) This figure shows the assets under management (AUM) of ETFs aggregated at the provider level. The values are computed as of December 30, 2016.

(A) 2014 event

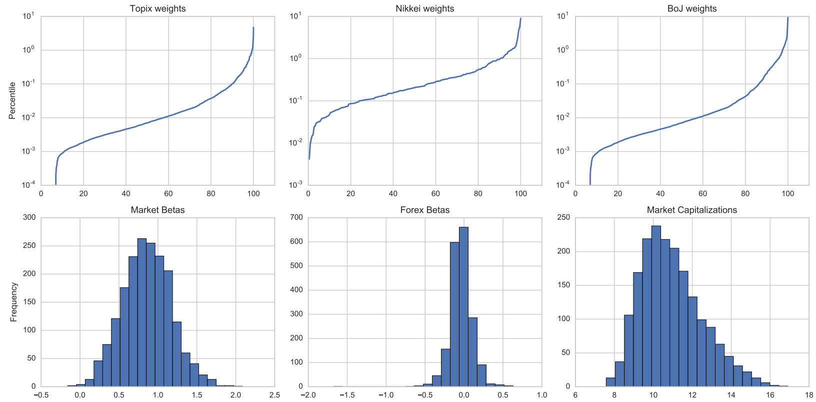

(B) 2016 event

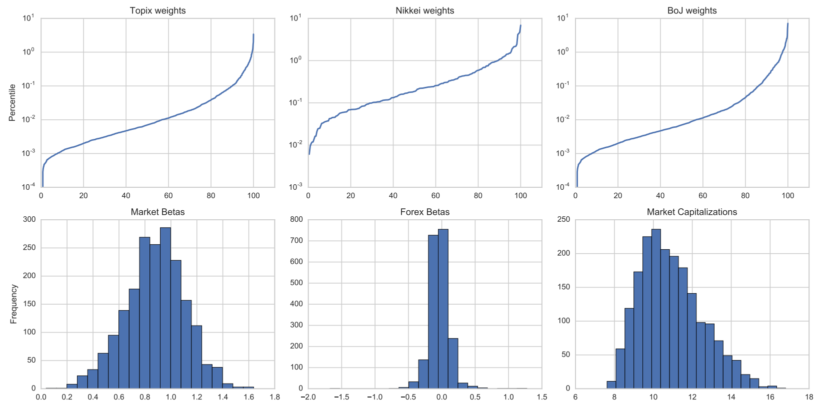

Figure A.4
Weights, betas, and market capitalizations The plots display cross-sectional heterogeneity of the variables of interest at the time of the BoJ announcements. Panel A refers to the announcement in 2014, and panel B refers to the announcement in 2016. The first row of each panel plots the percentile 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. The forex beta is the stock-level return sensitivity to the USD/JPY exchange rate. Market betas and forex betas are estimated over a window of 1 year, ending 2 trading weeks before each BoJ announcement. Companies market capitalizations are in logs.

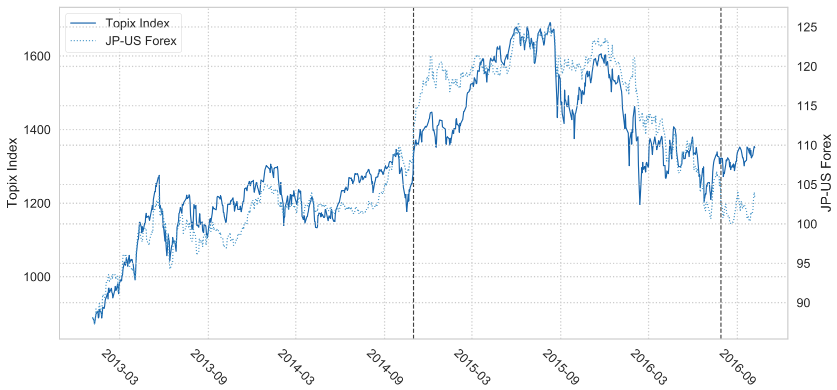

Figure A.5
TOPIX index and JP-US exchange rate This figure shows the time series of the TOPIX Index over our sample period (green solid line, left axis) and of the exchange rate from U.S. dollar to Japanese yen (purple dotted line, right axis).

Table A.1
Summary statistics

	Mean	SD	Min	25%	50%	75%	Max	Obs
Mkt cap (bn ¥)								
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
Mkt 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 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

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 on the Nikkei 225 index, and Not Nikkei stocks are those that only appear on the TOPIX index. All Nikkei companies also belong to the TOPIX index. All statistics are computed using pre-event information and by pooling both events together.

Table A.2
Robustness: Alternative covariance matrix estimation

Panel A: October 31st, 2014

	Raw returns			Abnormal returns		
	(1)	(2)	(3)	(4)	(5)	(6)
π	95.77***	91.16***	39.82***	18.13***	13.14**	47.69***
	(16.84)	(15.62)	(4.49)	(3.17)	(2.24)	(5.50)
u		0.01***	-0.00		0.01***	-0.00
		(3.32)	(-0.86)		(3.58)	(-0.99)
Mkt beta			0.02***			-0.08***
			(2.70)			(-11.04)
Forex beta			0.07***			0.07***
			(7.58)			(8.30)
log(mkt cap)			0.01***			0.01***
			(6.87)			(7.13)
Amihud			0.00			0.00
			(0.96)			(0.94)
Obs	1,851	1,851	1,807	1,851	1,851	1,807
R^2	.13	.14	.19	.01	.01	.11
Industry FE	NO	NO	NO	NO	NO	NO

Panel B: July 29th, 2016

	Raw returns			Abnormal returns		
	(1)	(2)	(3)	(4)	(5)	(6)
π	18.00***	16.27***	15.80***	18.00***	16.64***	18.59***
	(6.29)	(5.57)	(3.59)	(6.11)	(5.52)	(4.10)
u		0.01***	0.01*		0.01**	0.01*
		(2.75)	(1.82)		(2.10)	(1.80)
Mkt beta			-0.00			-0.01
			(-0.57)			(-0.99)
Forex beta			0.02**			0.02***
			(2.06)			(2.63)
log(mkt cap)			0.00			-0.00
			(0.99)			(-0.02)
Amihud			0.00			0.00*
			(1.48)			(1.75)
Obs	1,905	1,905	1,839	1,905	1,905	1,839
R^2	.02	.02	.03	.02	.02	0.03
Industry FE	NO	NO	NO	NO	NO	NO

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, and 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 parentheses. *p <.1; **p <.05; ***p <.01.

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

$$\max_N \mathrm{E}_t(-\exp(-\gamma W_{t+1})) \tag{B.1}$$

$$\text{s.t.} \qquad W_{t+1} = W_t(1+r) + N_t'(p_{t+1} + D_{t+1} - p_t(1+r)). \tag{B.2}$$

From the first-order condition it follows that

$$N_t = \frac{1}{\gamma}[\mathrm{Var}_t(p_{t+1} + D_{t+1})]^{-1}(\mathrm{E}_t[p_{t+1} + D_{t+1} - p_t(1+r)]). \tag{B.3}$$

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

$$(1+r)p_t = \mathrm{E}_t[p_{t+1} + D_{t+1}] - \gamma V Q_t \tag{B.4}$$

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

$$(1+r)p_t = \mathrm{E}_t\left[\frac{p_T}{(1+r)^{T-t-1}}\right] + \sum_{i=0}^{T-t-1}\frac{D_t}{(1+r)^i} - \gamma V \sum_{i=0}^{T-t-1}\frac{Q_{t,t+i}}{(1+r)^i}. \tag{B.5}$$

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

$$p_t = \frac{D_t}{r} - \frac{\gamma V}{(1+r)}\left(\sum_{i=0}^{\infty}\frac{Q_{t,t+i}}{(1+r)^i}\right) = \frac{1}{r}(D_t - \gamma V \Omega_t), \tag{B.6}$$

where we introduced the notation

$$\Omega_t = \frac{r}{1+r}\sum_{i=0}^{\infty}\frac{Q_{t,t+i}}{(1+r)^i}, \tag{B.7}$$

which can be interpreted as the discounted time-t belief on future supply. This term is crucial for our analysis, representing the channel through which the central bank is able to affect risk premiums. Under no expectation of monetary policy intervention, the resultant pricing equation collapses to

$$p_t = \frac{1}{r}(D_t - \gamma V Q), \tag{B.8}$$

where the vector $\gamma V Q$ can be interpreted as the cross-sectional vector of risk premiums 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 consider the effect on prices in case the central bank unexpectedly commits itself to a large-scale purchase of assets over a defined period, thus affecting the path of future supply Ω_t. We solve the model in its most general form, allowing agents beliefs on future supply to change over time, assuming that for each $t \geq 1$ there exist a scalar $\lambda_t \geq 0$ such that the time-t belief is

$$\begin{cases} Q_{t,t+i} = Q & \text{for } i \geq 0 \text{ and } t < 1 \\ Q_{t,t+i} = Q - \lambda_t(t+i)q & \text{for } i \geq 0 \text{ and } t = 1,\ldots,M \\ Q_{t,t+i} = Q - \lambda_t M q & \text{for } i \geq M - t \text{ and } t \geq 1 \end{cases} \tag{B.9}$$

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 future supply can be written as

$$\Omega_t = \frac{r}{1+r}\sum_{i=0}^{\infty}\frac{Q_{t,t+i}}{(1+r)^i} \tag{B.10}$$

$$= \frac{r}{1+r}\left(\sum_{i=0}^{M-t-1}\frac{Q-\lambda_t(t+i)q}{(1+r)^i}+\sum_{i=M-t}^{\infty}\frac{Q-\lambda_t Mq}{(1+r)^i}\right) \tag{B.11}$$

$$= Q-\frac{\lambda_t r}{1+r}\left(\sum_{i=0}^{M-t-1}\frac{(t+i)q}{(1+r)^i}+\sum_{i=M-t}^{\infty}\frac{Mq}{(1+r)^i}\right) \tag{B.12}$$

$$= Q-\frac{\lambda_t r}{1+r}\left(\sum_{i=0}^{M-t-1}\frac{(t+i-M)q}{(1+r)^i}+\sum_{i=0}^{\infty}\frac{Mq}{(1+r)^i}\right) \tag{B.13}$$

$$= Q-\lambda_t Mq+\frac{\lambda_t r}{1+r}\sum_{i=0}^{M-t-1}\frac{(M-t-i)}{(1+r)^i}q \tag{B.14}$$

$$= Q-\lambda_t Mq+\lambda_t\varphi(t)q \tag{B.15}$$

where we introduced the real-valued function

$$\varphi(t)=\frac{r}{1+r}\sum_{i=0}^{M-t-1}\frac{(M-t-i)}{(1+r)^i}, \qquad t\geq 1 \tag{B.16}$$

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 afterward it will not engage in further market operations, it follows that the pricing Equation (B.8) takes the following form:

$$\begin{cases} p_t=\frac{1}{r}(D_t-\gamma VQ) & \text{for } t<1 \\ p_t=\frac{1}{r}(D_t-\gamma V(Q-\lambda_t Mq+\lambda_t\varphi(t)q)) & \text{for } t=1,\dots M \\ p_t=\frac{1}{r}(D_t-\gamma V(Q-Mq)) & \text{for } t\geq M \end{cases} \tag{B.17}$$

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

$$p_1-p_0=\frac{1}{r}(\varepsilon_1+\lambda_1\gamma V(Mq-\varphi(1)q)) \tag{B.18}$$

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

$$p_{t+1} - p_t = \frac{1}{r}(\varepsilon_{t+1} - \gamma V((\lambda_{t+1}\varphi(t+1) - \lambda_t\varphi(t)) - \Delta\lambda_{t+1}M)q) \quad \text{(B.19)}$$

$$= \frac{1}{r}(\varepsilon_{t+1} + \gamma\xi(t+1)Vq), \qquad t = 1,\dots M \quad \text{(B.20)}$$

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

$$\lambda_{t+1}\varphi(t+1) - \lambda_t\varphi(t) < \lambda_{t+1}\varphi(t) - \lambda_t\varphi(t)) = \Delta\lambda_{t+1}\varphi(t) < \Delta\lambda_{t+1}M \quad \text{(B.21)}$$

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

$$R_{i,t+1} = \frac{1}{r}(\varepsilon_{i,t+1}/p_{i,t} + \gamma\xi(t+1)(Vq)_i/p_{i,t}) = \frac{1}{r}(\varepsilon_{i,t+1}/p_{i,t} + \gamma\xi(t+1)\pi_i) \quad \text{(B.22)}$$

Finally taking the expectation of the cumulative returns we get

$$\sum_{s=1}^{t} \mathrm{E}[R_s] = \sum_{s=1}^{t} \frac{1}{r}(\gamma\xi(s)\pi) = \theta_t\pi \quad \text{(B.23)}$$

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

$$\beta(P)_i = \mathrm{Cov}(\Delta MP, \Delta p_i) \propto \mathrm{Cov}(\varepsilon_t' Q, \varepsilon_{i,t}) = (VQ)_i \quad \text{(C.24)}$$

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 $\beta(R)_i$

$$\beta(R)_i = \frac{\mathrm{Cov}(R^{mkt}, R_i)}{\mathrm{Var}(R^{mkt})} \propto \mathrm{Cov}(R'W, R_i) = \sum_{i,j} \mathrm{Cov}(R_j, R_i)W_j = (\Sigma W)_i \quad \text{(C.25)}$$

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

$$\beta(P)_i^{\text{post}} - \beta(P)_i \propto (VQ^{\text{post}} - VQ)_i \propto -(Vq)_i \tag{C.26}$$

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

$$\beta(R)_i^{\text{post}} - \beta(R)_i \propto (\Sigma W^{\text{post}} - \Sigma W)_i \propto -(\Sigma u)_i = -\pi_i \tag{C.27}$$

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, whereas the fundamental covariance structure of returns is unchanged.