

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

Andrea Barbon and Virginia Gianinazzi *

This draft : November 2017

Abstract

Since the introduction of its Quantitative and Qualitative Easing program in 2013, the Bank of Japan (BoJ) has been increasing its holdings of Japanese equity through large purchases of index-linked ETFs with the declared intention of “supporting asset prices and reducing the cost of capital” of domestic companies. This program represents an ideal setting in which to clarify the mechanism through which unconventional monetary policy impacts asset prices, thanks to the rich and exogenous cross-sectional dimension of the purchase schedule and the long-term commitment of the central bank. We exploit these unique features to develop our identification strategy, supported by an asset pricing model that accounts for the reduction in the free float induced by the program. We document a positive, sizeable and persistent impact on stock prices, suggesting that demand curves for stocks are downward-sloping in the long-run. We estimate 20 basis points increase in aggregate market valuation per trillion Yen invested into the program as a lower bound for the net effect of the policy. Finally, we show theoretically and empirically that the weighting scheme of the BoJ portfolio, tilted towards the price-weighted Nikkei 225 index, generates significant pricing distortions consistent with the portfolio balance channel.

JEL classification: Central Banking, BoJ, ETF, QE, QQE, Japan, Monetary Policy

Keywords: E52, E58, B26

*Andrea Barbon and Virginia Gianinazzi are with USI Lugano and the Swiss Finance Institute. We are grateful to Francesco Franzoni, Alberto Plazzi, Robin Greenwood, John Campbell, Suresh Sundaresan, Thierry Foucault, David Miles, Marco LiCalzi, Marco Di Maggio, Lorian Mancini, Andrea Buraschi, Jonathan Kearns, Miguel Antón, Ruediger Fahlenbrach, Jean-Charles Rochet, Harjoat S. Bhamra and Thomas Geelen for valuable comments and helpful discussions. All errors are our own.

1 Introduction

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

There is considerable evidence that central banks' asset purchases can have an economically significant impact on yields in the targeted markets, which has likely motivated central banks to continue these purchases over the past decade (see Gagnon et al. (2010), Neely et al. (2010), Krishnamurthy and Vissing-Jorgensen (2011), Swanson (2011), Hamilton and Wu (2012), Krishnamurthy and Vissing-Jorgensen (2013), D'Amico and King (2013), Eser and Schwaab (2016)). However, despite the widespread use of LSAP programs, the debate is still ongoing with regard to the mechanisms linking asset purchases to asset prices and the persistence of the impact. Unlike policy rate targeting, asset purchases are explicit decisions on quantities and are designed to have a noticeable impact on market prices. Even though the idea of *easing through quantity* relies on the view that large purchases by the central bank reduce assets' risk premia, there is not yet a clear theoretical foundation for how and under which conditions this is expected to work. In general, the relationship between the quantity of an asset and its price is not yet well understood.

Since 2013, the Bank of Japan (BoJ) has been engaging in what they have named *Quantitative and Qualitative Easing* (QQE) program as an attempt to fight against deflation. As part of its broader QQE agenda, the BoJ has been vigorously increasing its domestic equity holdings through purchases of index-linked ETFs. By the end of 2016, the BoJ owned approximately ¥14 trillion worth of Topix and Nikkei ETFs, which corresponds to more than 2.5% of the total market capitalization. This unprecedented equity operation has the declared objective of supporting asset prices and reducing the cost of equity capital of Japanese companies (BoJ, 2013). We argue that this intervention represents a unique laboratory in which to test how QE impacts equity prices and its implications for market efficiency, as well as a unique opportunity to shed some light on the long-debated question of 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, the basic mechanism of which was first discussed by Tobin (1969), Brunner and Meltzer (1973) and Frankel (1985). 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 investors' aggregate portfolio. For this to be an equilibrium, prices need to adjust to ensure market clearing.

We propose a simple structural asset pricing model that generalizes the idea of the portfolio balance channel to the case of equities.¹ The key implication of portfolio rebalancing that we derive from the model is that the change in the systematic risk for each firm is determined by: (i) the entire vector of purchases and (ii) the covariance matrix of cash flows.

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

We document that both announcements produced a highly heterogeneous response of equity prices at the company level. Figure 1 plots for the 2014 announcement the average cumulative raw returns of the two groups of firms that were expected to be most and least affected by the policy. 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 returns on longer horizons, we find no evidence of reversal over a one-year window after the policy announcements, which supports the time-series prediction of the model. Prices adjust only gradually to the announcement, which, according to our model is consistent with the decrease in the residual duration of the program over time. This may also suggest that investors do not fully believe in the commitment of the central bank, but rather gradually increase the probability they attach to the continuation of the policy.

We estimate the long-term net effect of the portfolio-balance channel at market level at about 22 basis points per trillion yen employed. With the total market capitalization of about ¥500 trillion, this implies an elasticity close to one since each yen invested translates into an increase in total market valuation of roughly one yen. As it is made clear by the model, this effect is a joint function of the covariance structure of the market and the vector of purchases.

We argue that the two extensions of the purchase program provide us with an ideal natural experiment to examine the net effect of a long-lasting change in supply on prices for at least 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, since roughly half of the capital of the central bank is

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

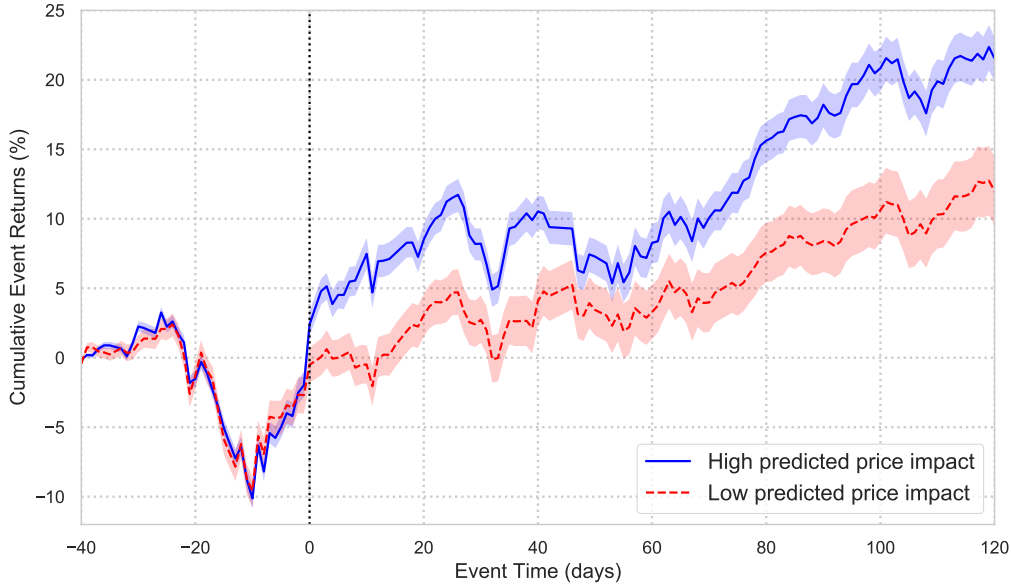


Figure 1: Heterogeneity of Event Returns. This figure shows the time series of the mean cumulative raw returns around the BoJ announcement of October 2014, of two groups of firms ranked by the model predicted price movement π . The blue line is the average for first quartile of the distribution (firms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

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 is a challenging task. However, the intervention of the BoJ provides us with an empirical framework that mitigates endogeneity concerns and improves the identification of the net effect of a change in supply, which crucially relies on the exogeneity of the shock.

The non-fundamental nature of the Nikkei weighting system was already exploited in Greenwood (2005, 2008) to establish a causal relationship between uninformed demand shocks and prices following a large redefinition of the Nikkei 225 index. A major difference between the LSAP setting and the one of index redefinitions lies in the nature of the shocks. As the central bank buys assets, it is effectively transferring a portion of fundamental risk from the private sector to its balance sheet, and holds it for an arguably long period of time. This is at least conceptually different from an index redefinition event, in which securities merely change hands from active investors to index funds. The central bank can be thought of as a buy-and-hold long-term investor whose portfolio holdings are not marked-to-market. Its long-term commitment to the policy induces a long-lasting change to supply, making our setting better suited than index redefinitions to identify price effects due to a movement along

investors' long-term demand curves.

The model that we propose extends the theoretical framework used in Greenwood (2005) to account for this difference in setting. As in Greenwood (2005), we consider an economy with multiple assets in finite supply and a CARA-utility representative agent that maximizes her wealth in each period. We introduce quantitative easing in the form of an exogenous shock to the supply of assets, which is first announced and then gradually carried out over a given policy horizon. The agent correctly understands that the QQE program will affect the market-clearing portfolio in each future period, which determines the new vector of equilibrium risk premia. Crucially, extending the model to an infinite horizon relaxes the assumption that uncertainty is resolved at a terminal date, which mechanically drives the reversal in Greenwood (2005). In our model, prices adjust to the change in supply to reflect the new risk composition of the market portfolio. It follows that we should not observe a reversal at any horizon, unless the central bank is expected to unwind its position. This prediction of the model is consistent with the observed time-series pattern. Market efficiency requires that today's asset prices reflect expectations of future returns, hence in this context, they should also reflect expectations about the future stock of assets. Therefore, a credible announcement by the central bank that it will purchase assets at a future date should reduce risk premia immediately. To allow for a sluggish price reaction after the announcement, we extend the model to feature imperfect credibility of the monetary policy commitment.

Finally, our cross-sectional 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. The model shows that the only way to achieve a cross-sectionally homogeneous shift in risk premia is for the BoJ to hold each stock proportionally to the company market capitalization. At the moment, however, still roughly a quarter of the BoJ capital is allocated to the price-weighted Nikkei index.

The rest of the paper is organized as follows. Section 2 describes the ETF purchase program of the BoJ. Section 3 reviews the relevant literature. Section 4 presents the model and discusses the effectiveness of different LSAP designs as implied by the model. Section 5 describes the data and the empirical results. In this section we also propose a way to quantify the net portfolio rebalancing effects of the policy. Section 6 concludes.

2 The ETF Program of the BoJ

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

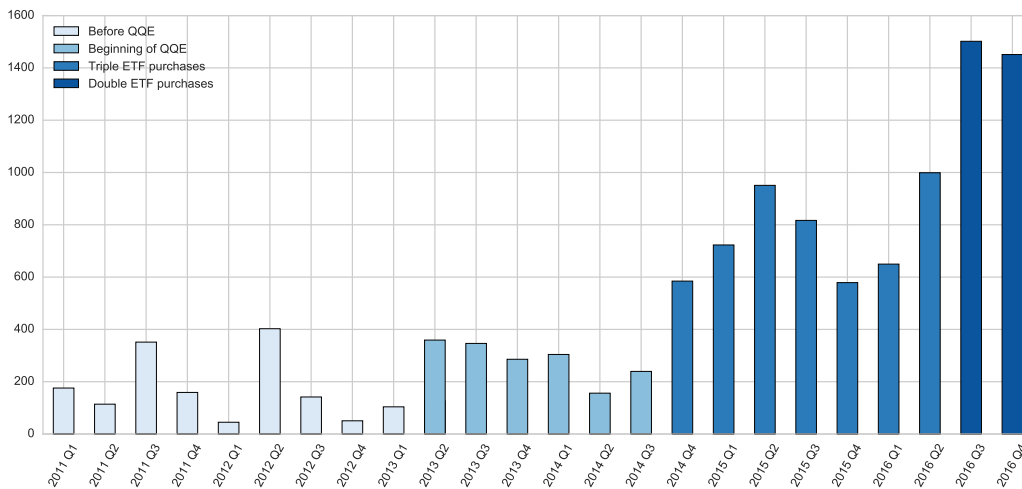


Figure 2: Quarterly ETF Purchases of the Bank of Japan in billion yen. Changes in the bar color indicate changes in the policy target purchase amounts. In the first phase the target was set to ¥1 trillion, in the second phase it was tripled to ¥3 trillion and in the third phase it was doubled to ¥6 trillion.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

Its holdings accumulated rapidly, and by the end of 2016, the BoJ owned more than ¥14 trillion worth of broad market ETFs. This corresponds to 2.5% of the total capitalization of the First Section of the Tokyo Stock Exchange, and around 3% of the Japanese GDP. The share of BoJ holdings to aggregate Assets Under Management (AUM) of broad market ETFs has grown from almost zero to more than 70% since the beginning of the program; this is even more remarkable if we consider that the ETF industry in Japan almost tripled in value between 2013 and 2016. For comparison, over the past 10 years the annual aggregate net flows into or out of the Japanese equity fund industry amounted on average to roughly ¥3 trillion

in absolute value.² The magnitude of the program is therefore economically important.

The policy targets two types of ETFs: those tracking the Tokyo Stock Price Index (Topix) and those replicating the return of the Nikkei 225 Stock Average.³ At inception of the program, the money allocated to each ETF was set to be proportional to its assets under management (AUM). The ratio of the aggregate AUM of ETFs tracking the Topix Index and those of ETFs tracking the Nikkei 225 Index is roughly 1 to 1.2. This approximately translates into half of the capital flowing into Nikkei ETFs and half into Topix ETFs. In turn, this then maps into a demand shock at the stock level that depends on each company's weight in the underlying index.

The Topix is a value-weighted index tracking the roughly 2000 companies listed on the First Section of the Tokyo Stock Exchange (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, liquid 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 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 misallocation 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 is clearly more severe for companies not included in the Nikkei index. However, Figure A3 shows that there is a high degree of heterogeneity in the allocation of capital across Nikkei companies as well. There we plot the distribution of the log of the ratio between the weight in the Nikkei and the weight in the Topix across Nikkei companies to measure the cross-sectional dispersion of the resulting allocation at the stock level.

Given the unusual weights of the BoJ purchase schedule, a sudden expansion of the policy budget produces a natural experiment where companies are hit by an uninformed demand shock that is highly heterogenous in the cross-section and orthogonal to firms fundamentals after controlling for market capitalization. In this paper we exploit the exogenous variation

²Data is from Thomson Reuters Lipper Global Fund Flows database.

³On November 19, 2014, the BoJ started buying also ETFs tracking the JPX-Nikkei 400 Index. This approximately corresponds to 43% of the purchases flowing to ETFs tracking the Topix, 53% to ETFs tracking the Nikkei 225 and the remaining 4% to ETFs tracking the JPX Nikkei 400. For simplicity, in the empirical analysis we round the share of both Topix and Nikkei ETFs to 50%, neglecting the JPX Nikkei 400. This simplification does not affect the results of our analysis.

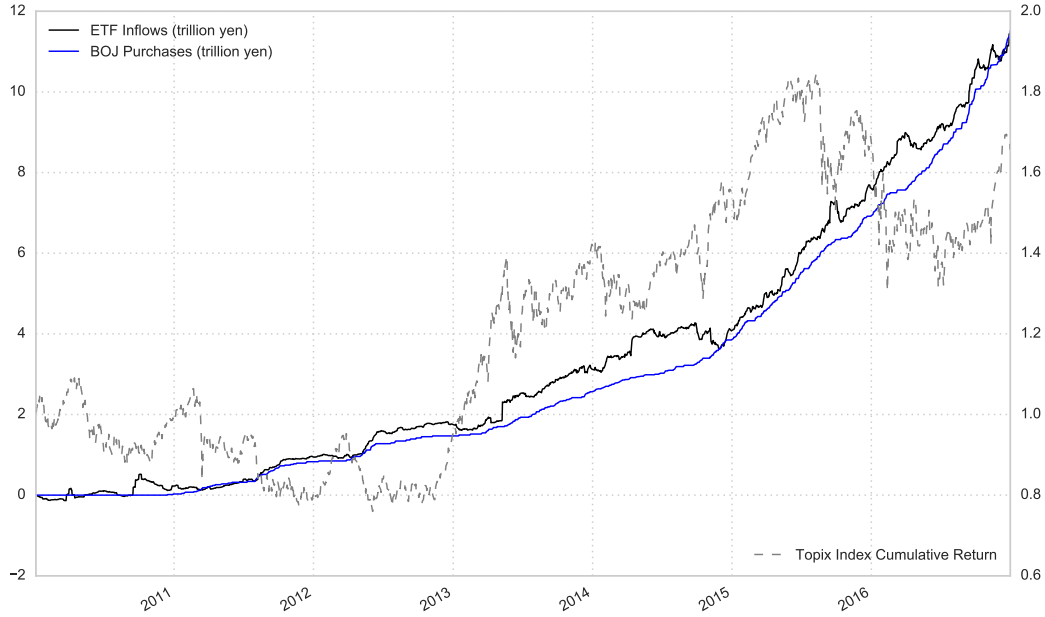


Figure 3: BoJ Purchases and ETF inflows. On the left axis we plot the daily cumulative purchases of ETFs by the BoJ (blue line) and the estimated daily cumulative inflows into ETFs tracking either the Nikkei or the Topix index (black line). Both are in trillion yen. On the right axis the figure shows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

in the cross-section of shocks to supply to identify the causal impact of the purchase program on equity prices. We argue with the help of a simple asset pricing model, that the deviation from a value-weighted allocation also allows us to isolate the portfolio-rebalancing channel of the policy impact.

Overall, the portfolio of the BoJ ends up deviating significantly from the allocation that market capitalization would dictate. To illustrate the extent of this distortion, take 3 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, however with very different weights, namely around 1.2% versus 9.5%, respectively. Nintendo, on the contrary, is not included in the Nikkei index. It follows that the BoJ allocates to Fast Retailing 4 times more capital than to Canon, and 19 times more than to Nintendo. The effects of the departure from a value-weighted allocation are reflected in the indirect ownership that the BoJ accumulated over time. According to estimates by the Financial Times, through these 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.

Given that the BoJ is purchasing ETFs, a question arises naturally about the link between the demand of ETFs and prices in the underlying stock market. Because of the creation

mechanism in the primary market, new ETF shares are issued if demand exceeds supply in the secondary market. In the case of physical ETFs, at creation the ETF sponsor is required to hold the physical basket of securities that composes the index that the ETF seeks to track for a value equal to the creation unit. This means that the creation of ETF shares necessarily involves a corresponding trade in the underlying market. The mechanism is visualized in Figure A1. Even though we cannot directly observe whether ETF shares have been created on demand of the central bank, we can infer it from the data on the AUM of ETF sponsors. First, from the website of the Japan Exchange Group (JPX) we obtain the list of the ETFs listed on the Tokyo Stock Exchange (TSE) that track either the Nikkei or the Topix index. From Bloomberg we then get daily data on AUM for each ETF. 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 the inflows into ETFs versus the amount purchased by the BoJ. It is apparent that the flows into these ETFs is almost completely due to the LSAP program. We can also infer from the plot that the purchases of the BoJ have consistently triggered creation of new ETF shares. In fact, trades that take place in the secondary ETF market do not affect the flows into an ETF since they do not directly involve the purchase or sale of underlying securities. We interpret this as evidence that the intervention of the BoJ in the ETF market constitutes an almost one to one intervention in the stock market.

It must be noted that the bias towards Nikkei companies did not go unnoticed among practitioners and the BoJ was blamed 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. From October 2016 the BoJ will allocate ¥2.7 trillion a year (US\$ 26.4 billion) to Topix ETFs, while the remaining ¥3 trillion will be spread out between the Topix, Nikkei 225 and JPX-Nikkei Index 400. For the Nikkei-ETFs this will mean a drop from 55% to about 25% of the annual purchases by the BoJ, which will bring 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.

3 Related Literature

With data on the effects of QE slowly becoming available following the actual implementation of unconventional policies by central banks, recent years have seen a surge in empirical and theoretical research on the effectiveness and the transmission mechanisms of unconventional monetary policy. Both Joyce et al. (2012) and, more recently, Buraschi and Whelan (2015) provide an excellent survey of the literature.

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

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

The empirical work generally finds that targeted asset purchases by central banks have affected rates of return. Gagnon et al. (2010), Neely et al. (2010), Swanson (2011), Hamilton and Wu (2012), Eser and Schwaab (2016) among others, all find evidence that the LSAPs by the Fed or the ECB did reduce longer term interest rates. Krishnamurthy and Vissing-Jorgensen (2011, 2013) find that the Fed’s purchase of long-term US Treasury bonds significantly raised Treasury bond prices, but had limited spillover effects for the private sector bond yields. This narrow effect seems to be consistent with clientele demand for safe assets, in the presence of which spillovers to other assets are expected to be limited to bonds that are considered as safe as Treasuries (scarcity or safety channel). Similarly, D’Amico and King (2013) find strong evidence of local effects with the impact being largest for the actual bonds purchased and for those of close maturities, which suggests limited substitutability across maturities. The concern that these local effects might be driven by the endogeneity of the choice of the Fed to target precisely those securities that it viewed as underpriced, is tackled in D’Amico and King (2013) with an instrumental variable approach. Krishnamurthy and Vissing-Jorgensen (2011) use intra-day data to support their claim that the effects documented in daily data

are indeed causal.

In order to keep up with the empirical evidence, theories on the effectiveness of LSAPs developed away from the notion 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 (2003) and crucially relies on the assumption of a rational infinitely lived agent with no credit restrictions, who sees no difference between its own assets and those held by the central bank.

For QE to have an impact on equilibrium prices, any asset pricing model has to introduce some form of friction. Imperfect asset substitutability is key to the so called "portfolio-balance" channel of unconventional monetary policy, whose basic mechanism was first discussed by Tobin (1969), Brunner and Meltzer (1973) and Frankel (1985). According to the idea behind this channel, when the central bank buys a particular asset, it reduces the amount held by the private investors, effectively forcing them into a different portfolio. For this to be an equilibrium, prices need to adjust to ensure market clearing. In particular, through this channel asset purchases are expected to push up the price of the target asset and of its substitutes. In the case of purchases of government bonds, the policy reduces the amount of duration risk in the hands of investors and therefore the return required for holding it. In this context, the baseline portfolio rebalancing mechanism is therefore sometimes referred to as duration channel. Vayanos and Vila (2009) tries to reconcile the predictions of the portfolio rebalancing channel with the observed lack of spillovers across maturities, building on market segmentation and preferred-habitat theories as proposed by Culbertson (1957) and Modigliani and Sutch (1966).

A large number of papers show that large-scale asset purchases by central banks have effects beyond those due to portfolio balance. In the case of purchases of long-term bonds, Krishnamurty and Vissing-Jorgensen (2011) provide compelling empirical evidence that the so called signalling channel explains a significant fraction of the drop in bond yields observed after the Federal Reserve's QE announcements. The idea behind this channel is discussed in Eggertsson and Woodford (2003), 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. As argued in Clouse et al. (2000), the commitment to low rates is especially credible when the central bank exposes its balance sheet to considerable duration risk through purchases of long term bonds. Markets may also infer from the commitment of the central bank to engage in large-scale equity purchases an inclination towards monetary accommodation. According to Eggertsson and Woodford (2003), asset purchases by central banks may impact prices only indirectly via the signalling channel.

The literature has proposed various other channels to explain the impact of QE on asset

prices. For instance, many researchers attribute the beneficial effect of the Fed’s MBS purchases on risk premia during the financial crises to a capital constraints channel motivated by the distress in the financial intermediary sector (Curdia and Woodford, 2011, He and Krishnamurthy, 2013). Overall, 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 this 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.

Our analysis is also closely related to another strand of literature, namely to a question that has a long tradition in asset pricing: do demand curves for stocks slope down? Ultimately, the workings of the portfolio balance channel is a matter of elasticity of assets’ demand curves. This has usually been considered an empirical question and since Scholes (1972) and Shleifer (1986) researchers have been on the hunt for natural experiments that allow to investigate the effect of clearly uninformed trades on prices. The empirical evidence so far mostly relies on event studies around index redefinition announcements and fire sales by institutional investors. The general finding is that large non-fundamental trades have a significant but temporary price impact, even though the speed and extent of the reversal is still subject of debate. Greenwood (2005) provides rather convincing evidence of downward-sloping short-term demand curves exploiting a comprehensive redefinition of the Japanese Nikkei 225 index that triggered rebalancing by index tracking investors for a total trading amount of around US\$ 19 billion. The paper finds significant event returns and an almost complete reversal over the two-month period after the announcement, which is consistent with limits to arbitrage in the short-run and price efficiency in the long-run. These results are interpreted as evidence that short-term demand curves slope down, but long-term ones are instead *flat*. We document that the large-scale purchases of equity ETFs by the BoJ produce a persistent price impact in the underlying securities which runs counter the idea that uninformed supply shocks affect prices only temporarily due to short-lived limits to arbitrage and is at odds with previous evidence on flatness of long-run demand curves.

4 The Model

We develop a theoretical framework to describe the portfolio rebalancing channel as the transmission mechanism from LSAP to asset prices. The idea is that the asset purchases of the BoJ shift part of the fundamental risk from the market to the balance sheet of the central bank. The net effect on asset prices of this risk transfer is not simply proportional to the purchased amounts but, instead, it crucially depends on the correlation structure of firms fundamentals. Because the premium demanded by investors for a given security is

proportional to its marginal risk contribution in the market portfolio, the price effect of the monetary intervention is proportional to the implied change in this quantity.

Our model features the central bank only in reduced form: the policy rule is exogenous and set equal to what the BoJ is implementing in practice.

4.1 General Framework

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_t paid at time t is

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

where each $\varepsilon_{i,t}$ is revealed at time t . The fundamental innovations ε_i are modelled as zero mean jointly normal random variables, iid over time. Let Σ denote the fundamental covariance matrix of the steady state equilibrium and γ the aggregate risk-aversion

Investors optimally choose their time- t demand N_t to maximize their next period utility

$$\begin{aligned} \max_{N_t} \quad & E_t (-\exp(-\gamma W_{t+1})) \\ \text{s.t.} \quad & W_{t+1} = W_t(1+r) + N_t'(p_{t+1} + D_{t+1} - p_t(1+r)) \end{aligned}$$

where W is the total wealth and N_t' denotes the transpose of the vector N_t .

Appendix C shows that the equilibrium asset pricing equation in the model is

$$p_t = \frac{1}{r} (D_t - \gamma \Sigma \Omega_t) \quad (2)$$

where Ω_t is the vector of time- t expected future asset supply, properly discounted by time. The closed form expression for Ω_t is derived in Appendix C. Absent monetary policy shocks, the equity premium collapses to $\gamma \Sigma Q$, an increasing function of the covariance with the market portfolio and the risk aversion parameter γ .

At date $t = 1$ the central bank announces purchases described by the vector $u = (u^1, \dots, u^n)$, equally distributed over M periods after the announcement. We refer to M as the policy horizon. Let u_t denote the vector of cumulative purchases by the central bank up to period t . We assume for simplicity that $u_t = t \cdot u$ for $t = 1, \dots, M$ and that $u_t = M \cdot u$ for $t > M$. One can think of u_t as the active side of the balance sheet of the central bank at any time t . This implies that in our model the central bank's balance sheet evolves deterministically and grows linearly over time.

The realized demand shocks negatively affect the net supply of assets in each period so that, setting $Q_0 = Q$, we have

$$E_t[Q_{t+1}] = E_t[Q_0 - u_{t+1}] = Q - E_t[u_{t+1}]$$

In the following sections we consider the price impact of these supply shocks under different degrees of policy credibility. The central bank intervention is always assumed to be fully unexpected before the announcement date, i.e. $E_0[u_t] = 0$ for every $t \geq 0$.

4.2 Perfect Credibility

We first consider the case in which the commitment of the central bank to the target purchases is fully credible and understood, so that there is no uncertainty about the stream of demand shocks:

$$E_t[u_{t+1}] = u_{t+1} = (t + 1)u$$

As shown in Appendix C, the price adjustment at the announcement is

$$p_1 - p_0 = \frac{1}{r} (\varepsilon_1 + \gamma \Sigma(Mu - \varphi(1)u)) \quad (3)$$

Ignoring fundamental innovations, (3) thus predicts a price jump at $t = 1$ of magnitude $\gamma \Sigma(Mu - \varphi(1)u)$. This swing in prices is due to the fact that the policy is unexpected at $t = 0$ but it is fully impounded into prices as soon as it is revealed, correctly accounting for the delayed realizations of the actual BoJ purchases.

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

$$p_{t+1} - p_t = \frac{1}{r} (\varepsilon_{t+1} - \gamma \Sigma(\varphi(t+1) - \varphi(t))u), \quad t = 1, \dots, M \quad (4)$$

where $\varphi(t) < M$ is a decreasing scalar function of time defined in Appendix C.

Equation (4) shows that in the post-announcement period the non-stochastic component of price changes is $\gamma \Sigma(\varphi(t+1) - \varphi(t))u$, which accounts for the delay between the announcement of the supply shocks and their realizations. These adjustments are quantitatively small compared to the initial reaction to the BoJ disclosure at $t = 1$. Nevertheless, because their cross-sectional distribution is persistently parallel and of the same sign of the initial price impact, they add up to create a positive momentum continuation of the initial cross-sectional effect.

In short, if the central bank target is fully credible prices immediately adjust to incorporate the impact of the policy over the entire horizon, motivated by the expected decrease in the

asset supply. Subsequently, small adjustments follow due to the fact that the supply is *actually* being decreased as the central bank carries over its purchases.

We now formulate testable hypotheses on the cross-sectional distribution of the described price effects, which are the crucial predictions of our model. We define the empirical counterpart \hat{u} of u as the vector of dollars amount purchased by the BoJ for each security

$$\hat{u}_i \equiv p_{i,0} \cdot u_i, \quad \forall i \in 1 \dots, n$$

where u_i is the announced purchase of stock i in units of shares. Moreover let

$$\hat{\Sigma}_{i,j} \equiv \text{Cov}(R_i, R_j), \quad \forall i, j \in 1 \dots, n$$

be the stationary covariance matrix of stock returns.

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

Proposition 2. *The vectors of returns R_{t+1} following the announcement day ($t \geq 1$) are positively related to the vector $\pi = \hat{\Sigma}\hat{u}$ in the cross-section.*

Proofs are in Appendix C.

These testable predictions of our theoretical framework are important for two reasons. First, they allow to disentangle the portfolio rebalancing explanation from alternative channels based on improved market conditions or forward guidance – see e.g. Gagnon et al. (2010) and Woodford (2012) – since only the former can account for the predicted cross-sectional heterogeneity in the price effect of the asset purchase program. Second, they provide a practical instrument that central banks can use to map any purchase schedule u to the corresponding impact on equity prices and, consequently, to the cost of financing of public firms.

In Section 5 we show that Proposition 1 and 2 are born out in the data, thus supporting our theoretical framework and identifying the portfolio rebalancing channel as the key transmission mechanism of monetary policy interventions to asset prices in the equity market.

4.3 Imperfect Credibility

As we will describe in Section 5, the data shows a sizable momentum effect following the announcement day. From a quantitative perspective, this pattern can hardly be reconciled with the small trend induced in post event returns by the function $\varphi(t)$. We thus allow for the possibility of the effect of the announced purchase program to be *slowly* impounded into prices, by means of a parameter measuring the degree of credibility of the central bank. This

simple modification of our model can fully account for the observed time series pattern even from a quantitative point of view.

Specifically, we introduce a time-varying scalar λ_t controlling the degree of investors confidence in the continuation of the asset purchase program. Let λ_t be real numbers such that

$$E_t [Q_j] = Q - \lambda_t u_j, \quad t \geq 1$$

The following observations about investors beliefs at time t are an immediate consequence of the definition:

- (i) $\lambda_t > 1 \implies$ the size of the LSAP is expected to be larger than announced
- (ii) $\lambda_t = 1 \implies$ the LSAP will be in place, exactly as announced
- (iii) $\lambda_t < 1 \implies$ the size of the LSAP is expected to be smaller than announced

As we show in Appendix C and we report in the following Propositions, the size of both the initial price reaction and the subsequent adjustments depend on the degree of credibility. More specifically the price jump at $t = 1$ is increasing in the initial degree of credibility λ_1 , while the post-event returns are linked to the time series evolution of λ_t .

Proposition 3. *The vector of price changes $p_1 - p_0$ during the announcement day is positively related to $\pi = \Sigma u$ in the cross-section. Moreover, the magnitude of $p_1 - p_0$ is directly proportional to λ_1 , the initial degree of the announced policy credibility.*

Proposition 4. *Assume $\Delta\lambda_{t+1} = \lambda_{t+1} - \lambda_t > 0$. Then the vector of price change $p_{t+1} - p_t$ following the announcement day ($t \geq 1$) is positively related to $\pi = \Sigma u$ in the cross-section. Moreover, the magnitude of $p_{t+1} - p_t$ is an increasing function of $\Delta\lambda_{t+1}$.*

Notice that $\lambda_t \equiv 1$ corresponds to the perfect credibility case, so that Propositions 3 and 4 collapse to Propositions 1 and 2, respectively.

Even though we think of λ_t as the degree of credibility of the central bank conditional on time- t information, the parameter suits a number of alternative interpretations, which are not mutually exclusive. For instance, a λ_t monotonically increasing to 1 over some trading days may reflect the slow movement of capital in the spirit of Duffie (2010). Or, as in Barberis and Thaler (2003), the slow reaction may be due to the bounded rationality of agents who fail to correctly process the consequences of the BoJ announced program.

5 Empirical Analysis

This section describes the data and reports the empirical findings.

5.1 A first look at the data

From Compustat Global we collect stock-level data on daily returns, volumes and shares outstanding for the roughly 2000 stocks of the Topix universe for the period 2013-2016. From the same source we also get returns and volume data for the Topix index during the period 2013-2016, and the time-series of dollar-yen exchange rates for the same period.

From Thomson Reuters Datastream we collect the Topix and Nikkei index weights for every stock in our sample, at monthly frequency. We get the daily time-series of ETF purchases carried out by the BoJ from the official website of the central bank, which is presented in Figure 2. The distribution of index weights is presented in Figures A4, together with the distribution of the weights of the BoJ purchase schedule. In the same Figures we plot the distributions of market betas and Forex betas, whose estimation is described in Section 5.3 and Section 5.4, respectively. Table A1 presents summary statistics on the companies in our sample.

5.2 Vector of expected purchases

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

Let A be the total Yen amount purchased by the BoJ, p_i be the price of security i at the announcement date and $w_{i,T}$, $w_{i,N}$ the weight of stock i in the Topix and Nikkei indices. Under the assumptions of the model the vector u of purchases by the BoJ is given by

$$u_i(\text{¥}) = Tw_{i,T} + Nw_{i,N}$$

where T and N indicate the amount BoJ capital allocated to Topix and Nikkei ETFs, respectively. Empirically, index weights $w_{i,T}$ and $w_{i,N}$ are calculated as of the end of the month preceding the announcement.

For the purpose of our analysis, we compute the vector u by $u_i = w_{i,T} + w_{i,N}$, so that the vector of model-implied price impacts $\pi = \Sigma u$ and its empirical counterpart are parallel.

5.3 Estimation procedure

We estimate the fundamental covariance matrix and market betas separately for the two events.

Over an estimation window of one year, ending four months before each BoJ announcement, we estimate a simple market model for each stock i

$$R_{i,t} = \alpha_i + \beta_i^M R_{M,t} + e_{i,t} \quad (5)$$

where $R_{i,t}$ are daily returns of stock i and $R_{M,t}$ is a proxy for daily returns of the market portfolio. We consider different proxies for the market portfolio, namely the return on the Topix Index and an equally weighted portfolio of the securities in our sample. Results are robust to the choice of the proxy.

The fitted residuals $\hat{e}_{i,t}$ from the estimation are used as the empirical proxy for the fundamental innovations $\varepsilon_{i,t}$ of the model as defined in (1). Consequently, we define the empirical counterpart $\hat{\Sigma}$ of the fundamental covariance matrix Σ as the cross-sectional covariance of the fundamental returns innovations⁴.

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

We estimate abnormal event returns of each stock i over a window of four months centered on the BoJ announcement date as $AR_{i,t} = R_{i,t} - \hat{R}_{i,t}$, where $\hat{R}_{i,t}$ is the predicted return of the estimated market model (5).

5.4 Event study

The model predicts a positive relationship between each security abnormal event return and the change in its marginal contribution to the risk of the aggregate portfolio. This testable cross-sectional hypothesis implied by the model is summarized in Proposition 4, of which Proposition 1 is a special case.

As a preliminary test of this relationship we rank stocks in the Topix universe by the predicted abnormal event return $\pi_i = (\Sigma u)_i$ into four equally-weighted portfolios. Figure 4 plots the cumulative abnormal returns of the low and high π portfolios. Plots on the left show the event returns around the first policy change in 2014, while those on the right present the effect of the second change in 2016. The reported bands represent bootstrapped 95% confidence intervals.

These plots provide a first evidence that the model is successful at predicting the cross-sectional variation in returns. The pattern of abnormal returns is similar for the two events,

⁴Estimating $\hat{\Sigma}$ on raw returns leads to qualitatively identical results in the empirical analysis.

with returns on the high π portfolio being significantly higher than those on the low π portfolio.

One might be concerned that the sorting on π is implicitly ranking stocks based on firms' characteristics such as size, export share, market beta, leverage or growth opportunities, which might explain the divergence in returns. Table A2 reports summary statistics for the firms in the four groups. Given that the policy is putting a lot of weight on Nikkei companies, which are on average larger than non-Nikkei ones (see Table A1), firms in the different quartiles differ in their market value, with π being positively correlated with market capitalization.

The policy announcement could affect equity prices through its impact on the foreign exchange market. We estimate stocks' sensitivities to changes in the exchange rate by running the following regression

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

for each stock i over the entire sample period. Here, F_t is the daily percentage change in the exchange rate from US Dollar to Japanese Yen. The summary statistics show that most Japanese companies are negatively affected by a depreciation of the yen, net of market movements. This is especially true the lower π . It is therefore important to control for β^F when testing the model: if the yen depreciated as a consequence of the announcement, we would spuriously observe returns proportional to π , which could be confused with evidence in favor of portfolio rebalancing.

Market leverage and market-to-book ratios are similar across groups. The market beta as well does not show a group specific pattern. To check that the explanatory power of our model is robust to firms' characteristics, in the following section we run cross-sectional regressions including different sets of control variables.

5.5 Cross-sectional regressions

We confirm the pattern in the plots by running security-level cross-sectional regressions of event returns on the predicted price impact π_i and a set of control variables:

$$R_i^H = a_0 + a_1 \pi_i + a_2 u_i + a_3 \log(\text{cap}_i) + a_4 \beta_i^M + a_5 \beta_i^F + a_6 \text{ILLIQ}_i + \eta_i \quad (6)$$

For the purpose of these regressions, event returns are defined as the cumulative returns computed over the 10 trading days following the announcement ($H=10$). We control for a security's weight in the purchase schedule of the BoJ (u), the natural logarithm of its market capitalization, its market beta, its Forex beta and its Amihud ratio as a proxy for illiquidity. The choice of the control variables is informed by the heterogeneity of companies across groups pointed out in the previous section. We include industry fixed effects to make sure that our results hold within industries.

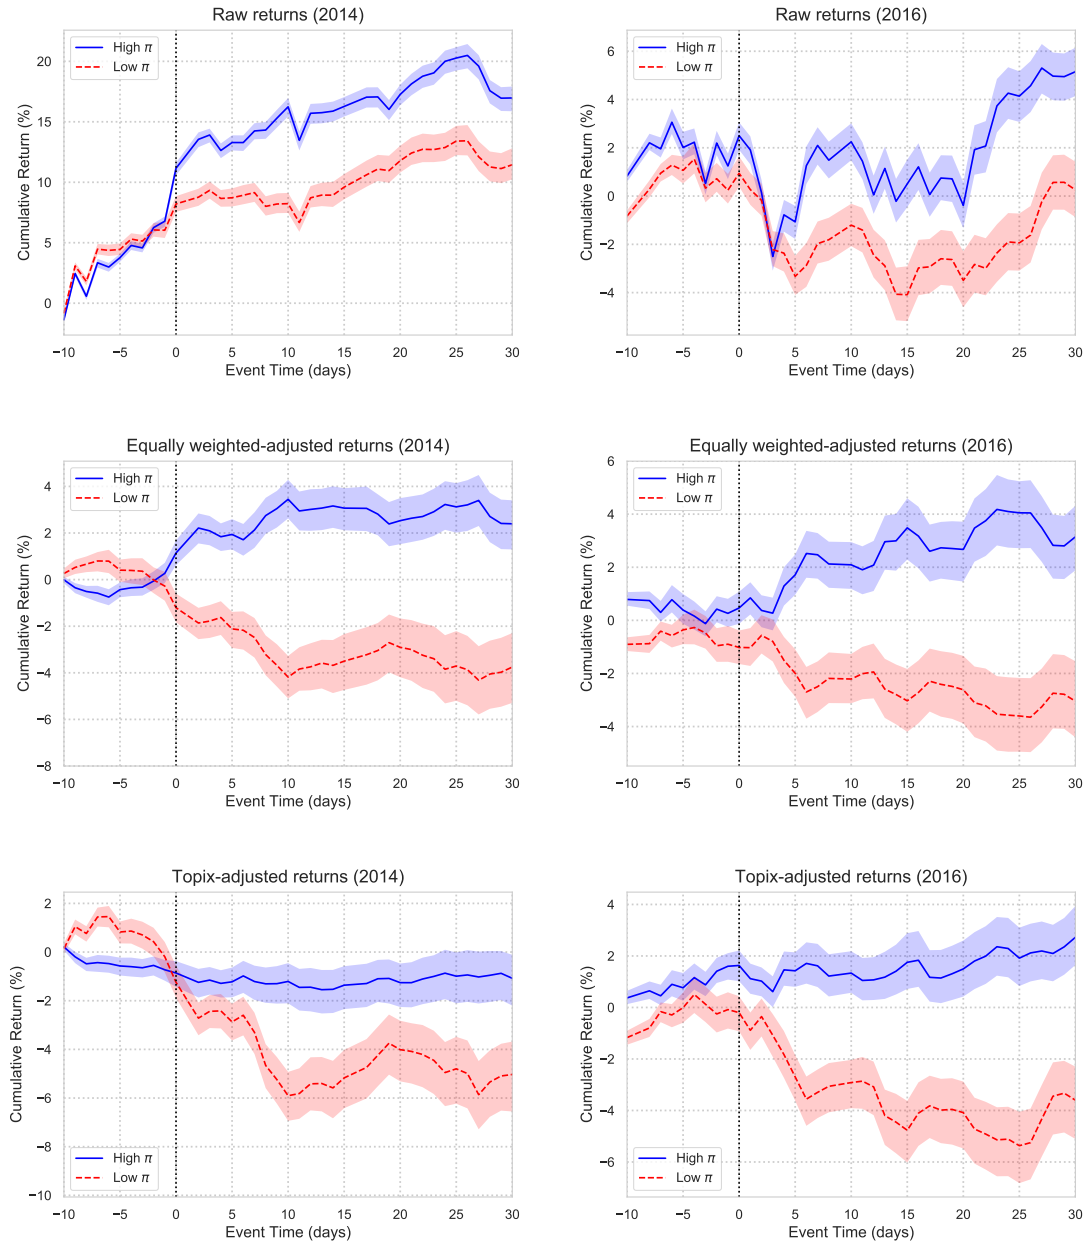


Figure 4: Cumulative returns of high versus low π stocks (in percentage). This figure shows the time series of the mean cumulative returns around the BoJ announcements of stocks with high predicted price impact π against that of low π stocks. The plots on the left refer to the announcement on October 31st, 2014, while those on the right show the reaction to the announcement on July 29th, 2016. The two top panels plot the unadjusted returns. In the four remaining panels returns are adjusted using a market model estimated in a window of one year, as described in Section 5.3. An equally-weighted portfolio of stocks in the Topix universe is used as a proxy for the market portfolio in the middle panels, while the return of the Topix index is used in the bottom panels. The blue line is the average for the first quartile of the distribution (firms with the highest predicted price impact), while the red dashed line corresponds to the average for the last quartile (firms with the lowest predicted price impact). Bands represent bootstrapped 95% confidence intervals.

We estimate this regression using data on the entire universe of Topix firms, for the two events separately. Panel A of Table 2 investigates the cross-sectional effect of the BoJ announcement on October 31, 2014 (when the target purchase amount of ETFs was tripled), while Panel B analyzes event returns following the announcement on July 29, 2016 (when the target was raised further, namely doubled). In either cases, no change was made to the weighting scheme of the purchases.

In the first three columns of both panels the dependent variable is the cumulative raw return, while in columns 4-6 the left-hand variable is the cumulative abnormal return with respect to the market model estimated in the pre-event window.

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 absence of policy shocks, which we use to compute robust standard errors. The placebo event days are chosen randomly on non-overlapping periods to ensure that the empirical distribution is constructed from independent draws. For regressions involving short-horizon returns (up to 3 months) we impose that placebo event periods do not include BoJ meetings on which bold monetary policy announcements were made. Namely, we exclude the meetings of February 1st 2013, March 25th 2013, June 18th 2012 and the announcement of the post-tsunami intervention in March 14th 2011. On all regression tables of this paper we report the empirical p-values computed using this methodology.

The coefficient on the predicted price impact π is positive and significant across specifications and events. This is consistent with our model, where the price adjustment of each security is driven by the change in its marginal risk contribution to the aggregate portfolio.

In the specifications 2 and 6 we include the vector of purchased amounts u in the regression, which we define as the sum of the weights in the Topix and the Nikkei index, for each security. Results show that the effect of π is robust to the inclusion of the index weights. This indicates significant spillovers as predicted by the portfolio-balance channel beyond local effects of the net purchases by the BoJ.

It is natural to expect any monetary policy announcement by the BoJ to have an impact on exchange rates. Figure A2 suggests that this is indeed the case for both announcements, even though the Forex reaction is of opposite sign across the two events. To control for this channel, we include in the regression the estimated exposure of each security on the JP-US rate (β_{Forex}). The estimation procedure is explained in the previous section.

Results reported in Table 2 show that this control does not impair the significance of the coefficient on π . The observed heterogeneity in the price response to the announcement cannot thus be explained by the indirect effect of the policy on the foreign exchange rate.

Panel A: October 31st, 2014

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

Panel B: July 29th, 2016

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

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

The coefficient on β_{Forex} is positive and significant in 2014, when the BoJ announcement was followed by a sharp rise in the Forex. In 2016, on the other hand, the coefficient on β_{Forex} is not significant, consistent with the fact that the Forex did not move significantly (see Figure A2).

5.6 Time-series pattern

In this section we turn to the analysis of post-event returns.

According to our model, the time-series pattern of returns following the BoJ announcement depends crucially on the credibility of the policy. In particular according to the baseline version of our model – with perfect credibility – the effect of the policy should be reflected by stock prices immediately after the BoJ announcement. On the other hand, the extended version of the model – with imperfect credibility – produces a wider range of time-series patterns depending on the dynamics of the policy uncertainty parameter λ_t , as stated in Proposition 4.

For instance, if the BoJ proves commitment to the program and investors augment their confidence (λ_t increasing in t), then the impact of the policy on asset prices should grow over time. We should therefore observe a momentum continuation of the initial price adjustments.

Intuitively, if investors are fully confident that the announced purchase program will be carried out as promised, then the entire effect of the policy should be incorporated in stock prices immediately after the announcement. If, instead, investors are not fully confident about the BoJ commitment to the program, the immediate effect on stock prices should only partially reflect the risk-related consequences of the program. Abnormal price movements in the days after the announcement depend on the dynamics of investors confidence in the policy.

Figure 4 suggests that the cross-sectional effect of the BoJ announcements is permanent and weakly increasing over time. To systematically assess this pattern we run a number of cross-sectional regressions, specified as in (6), varying the horizon H over which event returns are calculated.

Results of this exercise are reported in Table 3, showing that the coefficients on π are positive and significant for every horizon H . Still 1 year after the event there is no evidence of reversal of the initial price impact. Moreover, the coefficients on π are generally increasing in H . This suggests that the effect of the BoJ policy is not immediately reflected on prices, but rather it is increasingly impounded over the following weeks.

The observed pattern of post-event returns is consistent with a dynamics of the uncertainty parameter λ_t which is increasing over time. For instance one could think of an affine λ_t or, alternatively, an exponential functional form like $\lambda_t = \exp(t/M - 1)$.

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

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

5.7 Quantification of portfolio-balance effects

In this section we are interested in quantifying the net aggregate portfolio balance effect of the implemented program.

We start by estimating the following regression model for each daily horizon $h \in 1, \dots, 252$

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

where X is a vector of control variables that depends on the regression specification and $e \in (2014, 2016)$. We include market capitalization in each specification to control for the size factor, which is expected to become more relevant as the horizon increases. Unlike the regressions in the previous sections, here we pool the data from the two events in order to estimate the average effect of the policy. We include event fixed-effects FE_e to allow for a different intercept across the two announcements. Since we are considering both events together, we need to rescale the π vectors to take into account the different magnitude of the announcements. Therefore we multiply π_{2014} by 3 and π_{2016} by 6 to reflect the magnitude of the target amount announced by the BoJ in the two events, respectively.

Given $\hat{\beta}^h$ from the estimation, the predicted net return through the portfolio balance channel for security i is $\hat{R}_{i,e}^h = \hat{\beta}^h \pi_{i,e}$.⁵ For each event e , the predicted market return is the value-weighted sum of security level predicted returns at every horizon

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

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

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

Results of this exercise, reported in Table 4 for the three specifications, can be summarized as a long-term impact of about 22 basis points per trillion yen employed. With the total market capitalization of about ¥500 trillion, this implies an elasticity close to one since each yen invested translates into an increase of the market valuation by roughly one yen.

⁵Notice that the estimated $\hat{\beta}^h$ allows us to compare the impact of alternative purchase schedules u' that the central bank could have implemented, conditional on the same covariance matrix Σ .

⁶If the extension of the program was partially anticipated by the market the estimated impact per yen would increase, therefore we can consider our figures as a lower bound for the policy effect.

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

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

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

5.8 Policy Implications

In light of our empirical findings, in this section we discuss the policy implications of the theoretical framework, focusing on the effectiveness of different LSAP designs.

To formalize the discussion on the price distortions induced by the LSAP program, recall that in our model the cost of capital of each firm is proportional to its marginal risk contribution to the market portfolio (systematic risk). Formally, the vector of risk premia prior to the BoJ intervention is proportional to ΣQ , where Σ 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 $u \in \mathbb{R}^N$, the cost of capital is affected and converges to $\Sigma(Q - Mu)$. In particular, firm i experiences a percentage shift in its perceived cost of capital equal to

$$\Delta k_i = \frac{(\Sigma(Q - Mu))_i}{(\Sigma Q)_i} - 1$$

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

It follows that, from a theoretical point of view, an homogeneous impact on risk premia can be achieved only with a vector of purchases proportional to Q . In fact, if the purchase schedule is $u^* = aQ$ for $a \in \mathbb{R}$, the effect on firm i is

$$\Delta k_i^* = \frac{(\Sigma(Q - Mu^*))_i}{(\Sigma Q)_i} - 1 = \frac{((1 - Ma)\Sigma Q)_i}{(\Sigma Q)_i} - 1 = \frac{((1 - Ma)\Sigma Q)_i}{(\Sigma Q)_i} - 1 = -Ma$$

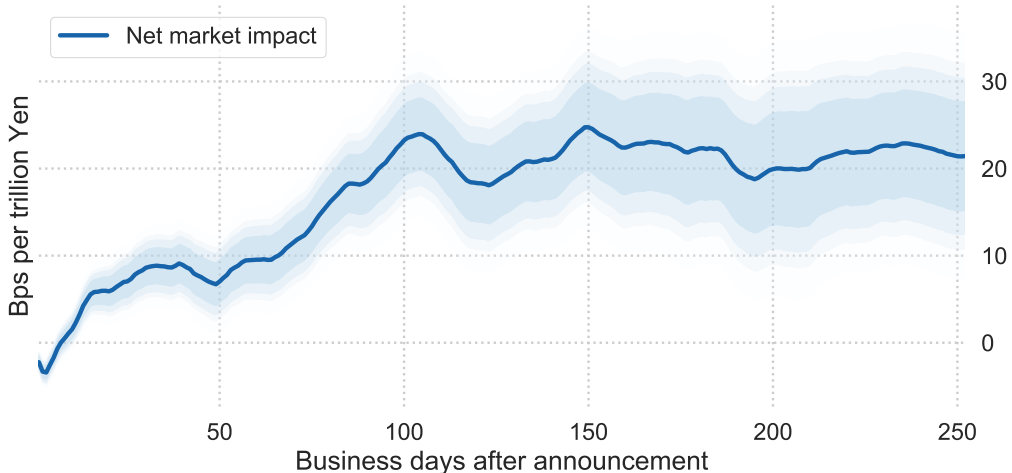


Figure 5: Portfolio Balance Effects. This figure plots the time-series evolution of the estimated portfolio balance effect induced by the BoJ purchase program, expressed in basis points per trillion Yen invested. The estimates are based on specification (3), which includes controls for stocks liquidity,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

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

In the case of Japan, a purchase schedule u parallel to Q corresponds to the BoJ limiting its purchases of ETFs to those tracking the value-weighted Topix Index ⁷. Buying ETFs tracking the price-weighted Nikkei 225, on the other hand, introduces a component in u 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 A6 shows that the distortion is evident also at industry level.

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

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

⁷A purchase of $u' = aW_{\text{Topix}}$ in Yen corresponds to $u = aQ$ in shares, since the Topix is value-weighted.

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

6 Conclusion

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

Exploiting the exogeneity and the cross-sectional dimension of the BoJ's purchase schedule, we overcome the endogeneity problems of previous studies and clearly identify the portfolio balance channel at work. We show that the intervention is having a long-lasting effect on Japanese equity prices, thus reducing the cost of capital of domestic companies. We document economically significant portfolio balance effects on the equity market, with an estimated magnitude of 20 basis points per trillion Yen invested into the program.

The long-term reduction in asset supply induced by the commitment of the central bank has a long-lasting effect on prices, thus implying that long-term demand curves for stocks are not flat. We rationalize this empirical finding by arguing that the intervention induces a change in the structure of the systematic risk borne by the private sector, thus leading to a new discount factor. The slope of securities demand curves hence depends on the entire covariance structure of securities' fundamentals.

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

References

- Barberis, N. and Thaler, R. (2003). A survey of behavioral finance. *Handbook of the Economics of Finance*, 1:1053–1128.
- BoJ (2013). Introduction of the "quantitative and qualitative monetary easing, april 4th".

- Brunner, K. and Meltzer, A. H. (1973). Mr. Hicks and the "monetarists". *Economica*, 40(157):44–59.
- Buraschi, A. and Whelan, P. (2015). Bond markets and unconventional monetary policy.
- Culbertson, J. M. (1957). The term structure of interest rates. *The Quarterly Journal of Economics*, 71(4):485–517.
- Curdia, V. and Woodford, M. (2011). The central-bank balance sheet as an instrument of monetary policy. *Journal of Monetary Economics*, 58(1):54–79.
- D’Amico, S. and King, T. B. (2013). Flow and stock effects of large-scale treasury purchases: Evidence on the importance of local supply. *Journal of Financial Economics*, 108(2):425–448.
- Duffie, D. (2010). Presidential address: Asset price dynamics with slow-moving capital. *The Journal of Finance*, 65(4):1237–1267.
- Eggertsson, G. B. and Woodford, M. (2003). Optimal monetary policy in a liquidity trap. Technical report, National Bureau of Economic Research.
- Eser, F. and Schwaab, B. (2016). Evaluating the impact of unconventional monetary policy measures: Empirical evidence from the ecbs securities markets programme. *Journal of Financial Economics*, 119(1):147–167.
- Frankel, J. A. (1985). Portfolio crowding-out, empirically estimated. *The Quarterly Journal of Economics*, 100(Supplement):1041–1065.
- Gagnon, J., Raskin, M., Remache, J., and Sack, B. P. (2010). Large-scale asset purchases by the federal reserve: did they work?
- Greenwood, R. (2005). Short- and long-term demand curves for stocks- theory and evidence on the dynamics of arbitrage. *Journal of Financial Economics*.
- Greenwood, R. (2008). Excess comovement of stock returns- evidence from cross-sectional variation in nikkei 225 weights.pdf. *The Review Of Financial Studies*.
- Hamilton, J. D. and Wu, J. C. (2012). The effectiveness of alternative monetary policy tools in a zero lower bound environment. *Journal of Money, Credit and Banking*, 44(s1):3–46.
- He, Z. and Krishnamurthy, A. (2013). Intermediary asset pricing. *The American Economic Review*, 103(2):732–770.
- Joyce, M., Miles, D., Scott, A., and Vayanos, D. (2012). Quantitative easing and unconventional monetary policy—an introduction. *The Economic Journal*, 122(564).
- Krishnamurthy, A. and Vissing-Jorgensen, A. (2011). The effects of quantitative easing on interest rates: channels and implications for policy. Technical report, National Bureau of Economic Research.
- Krishnamurthy, A. and Vissing-Jorgensen, A. (2013). The ins and outs of lsaps. In *Federal Reserve Bank of Kansas City’s Jackson Hole Symposium on the Global Dimensions of Unconventional Monetary Policy, Jackson Hole, Wyoming, August*.

- Ledoit, O. and Wolf, M. (2004). A well-conditioned estimator for large-dimensional covariance matrices. *Journal of multivariate analysis*, 88(2):365–411.
- Modigliani, F. and Sutch, R. (1966). Innovations in interest rate policy. *The American Economic Review*, 56(1/2):178–197.
- Neely, C. J. et al. (2010). *The large scale asset purchases had large international effects*. Federal Reserve Bank of St. Louis, Research Division.
- Scholes, M. S. (1972). The market for securities: Substitution versus price pressure and the effects of information on share prices. *The Journal of Business*.
- Shleifer, A. (1986). Do demand curves for stocks slope down? *The Journal of Finance*.
- Swanson, E. T. (2011). Let’s twist again: a high-frequency event-study analysis of operation twist and its implications for qe2. *Brookings Papers on Economic Activity*, 2011(1):151–188.
- Tobin, J. (1969). A general equilibrium approach to monetary theory. *Journal of money, credit and banking*, 1(1):15–29.
- Vayanos, D. and Vila, J.-L. (2009). A preferred-habitat model of the term structure of interest rates. Technical report, National Bureau of Economic Research.
- Woodford, M. (2012). Methods of policy accommodation at the interest-rate lower bound. *The Changing Policy Landscape*, pages 185–288.

Appendix A : Figures

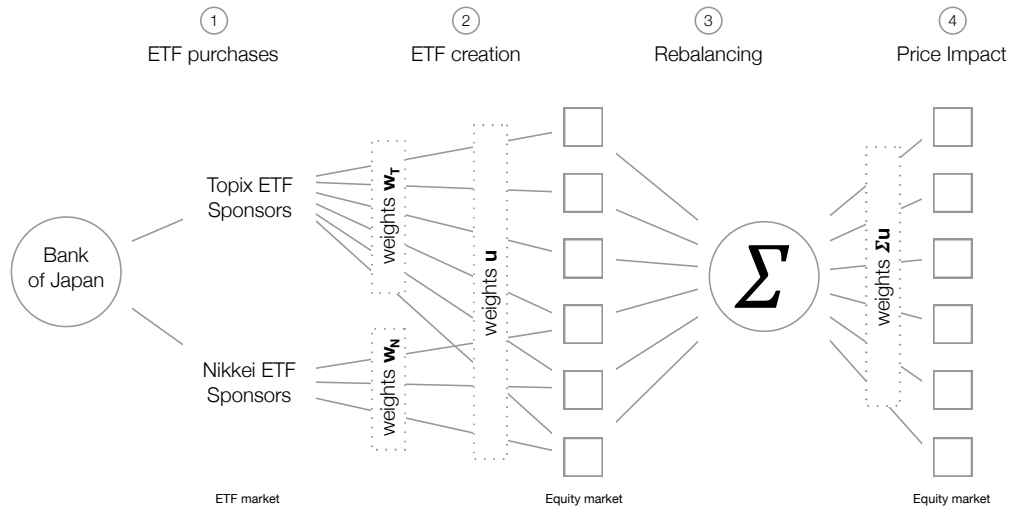


Figure A1: From ETFs to equity This figure describes the channel through which ETF purchases of the central bank may have an impact on equity prices. As the BoJ buys Topix- and Nikkei-linked ETFs, these are created by ETF sponsors and/or authorized participants. The securities needed to form the ETF basket are collected by these intermediaries in the equity market, thus effectively reducing the supply of equity shares available to private investors.

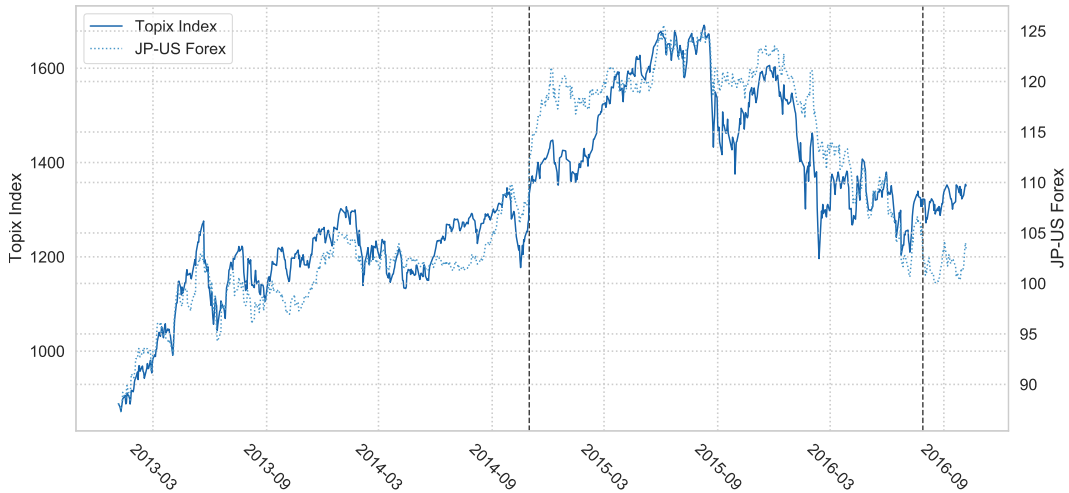


Figure A2: Topix Index and JP-US Exchange Rate. This figure shows the time-series of the Topix Index over our sample period (green solid line, left axis) and of the exchange rate from US Dollar to Japanese Yen (purple dotted line, right axis).

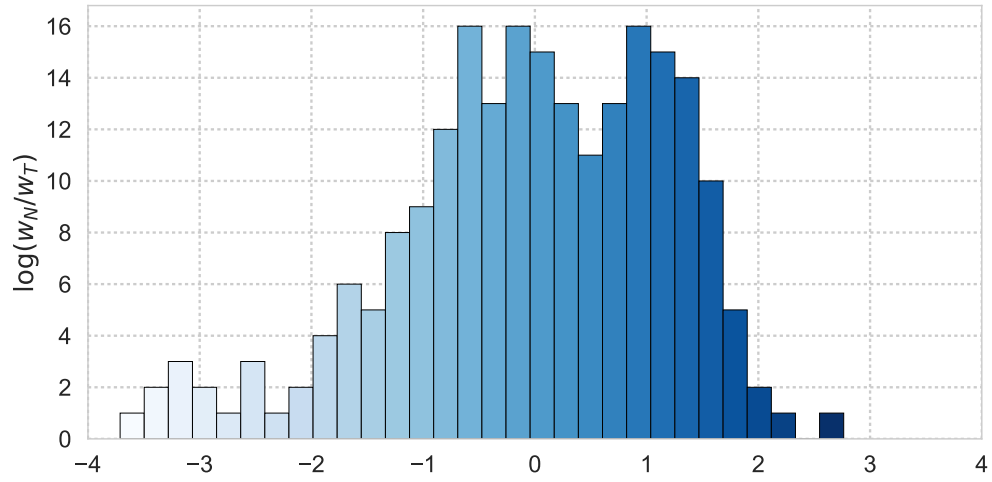
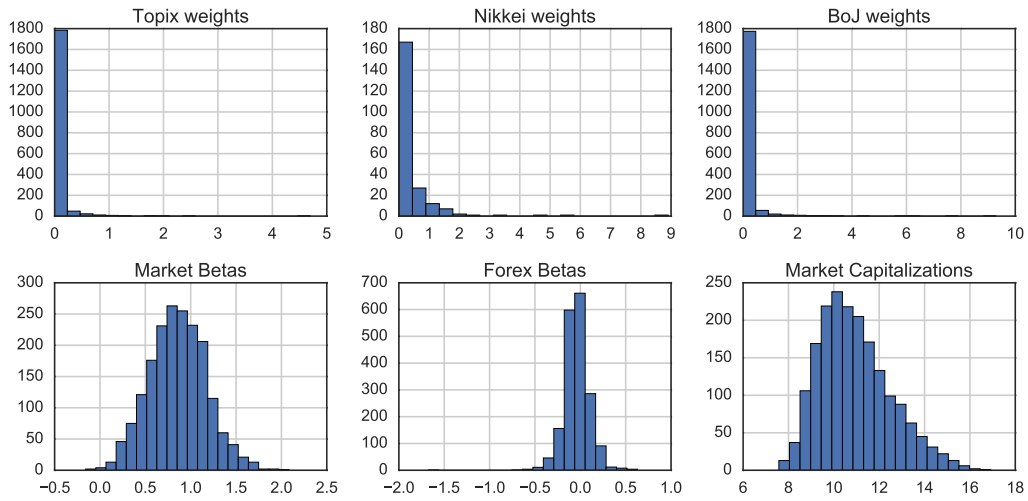


Figure A3: Distortion. The figure plots the distribution of the log ratio between the Nikkei weight and the Topix weight for Nikkei firms only. The histogram shows a significant dispersion, confirming that Nikkei weights induce significant cross-sectional variation of purchased quantities relative to market capitalization.

Panel A



Panel B

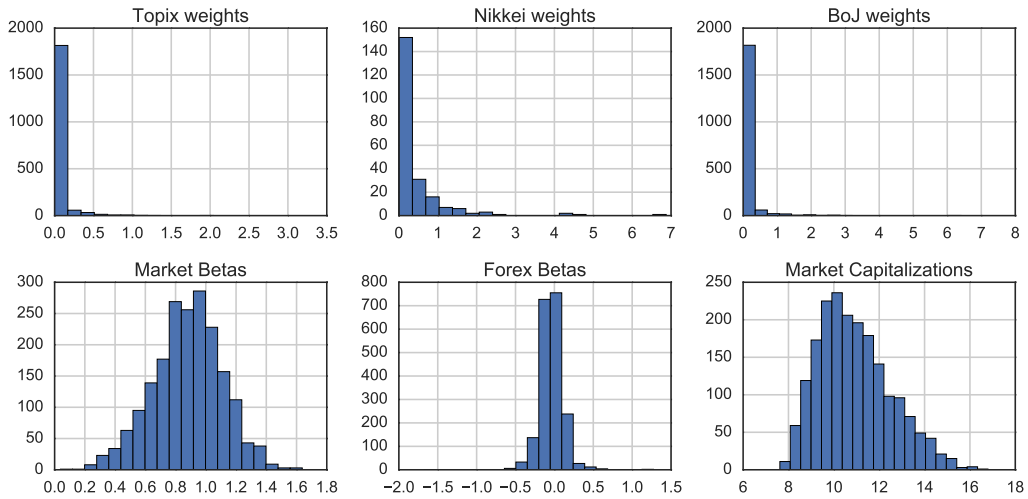


Figure A4: Weights, betas and market capitalizations. The histograms display cross-sectional distributions at the time of the BoJ announcements. Panel A refers to the announcement in 2014, Panel B to the announcement in 2016. The first two top panels show the distribution of the weights in the Topix index (ω_T) and in the Nikkei 225 index (ω_N). BoJ weights are computed as $\omega_T + \omega_N$ and correspond to the elements of the vector u in the model. Market betas and Forex betas are estimated following the procedure explained in section 5.3. Companies Market Capitalizations are in logarithmic scale.

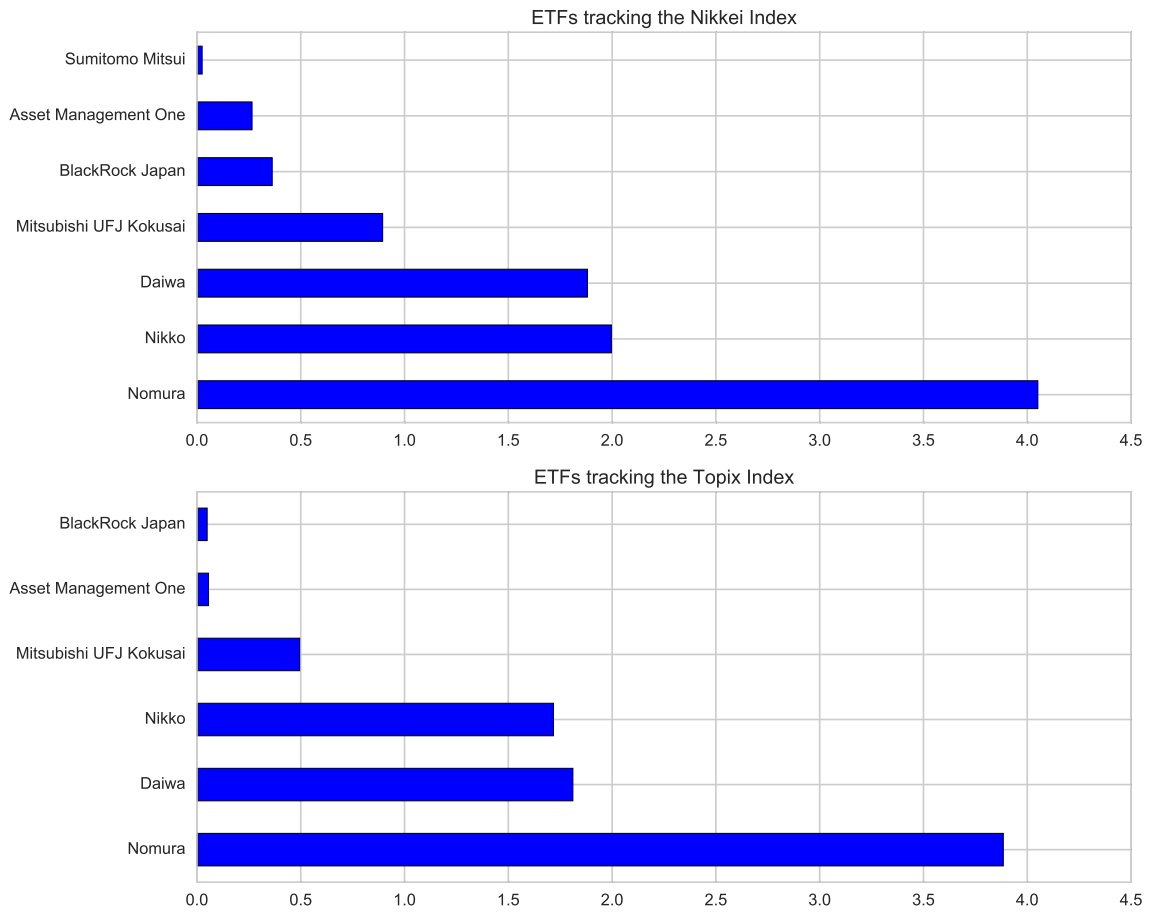


Figure A5: Assets Under Management (AUM) by Provider (in trillion yen). This figure shows the Assets Under Management of ETFs aggregated at Provider level. The values are computed as of December 30th, 2016.

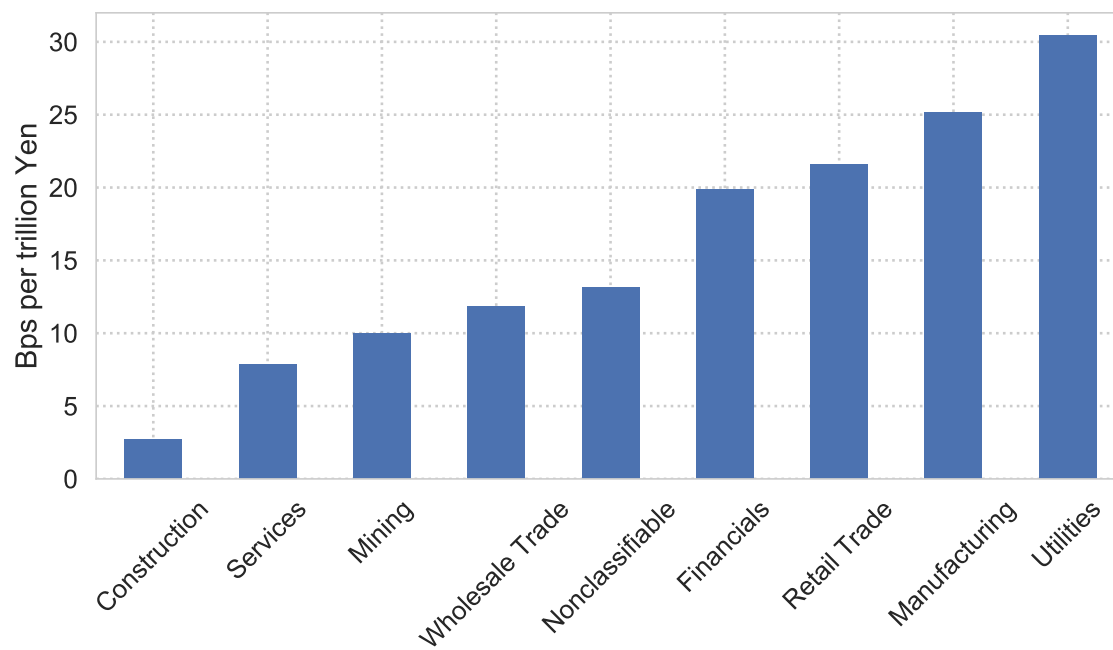


Figure A6: Portfolio Balance Effect across Industries This figure shows the estimated portfolio balance impact of the policy, expressed in basis points per trillion Yen, computed separately for each sector.

Appendix B : Tables

	Mean	Std Deviation	Min	25%	50%	75%	Max	Obs
Market Cap (Billions Yen)								
Topix	254	876	2	18	45	145	22210	1879
Nikkei	1388	2171	28	293	693	1496	22210	220
Not Nikkei	104	231	2	16	35	90	3191	1659
Forex Beta								
Topix	-0.04	0.14	-1.68	-0.12	-0.04	0.04	0.63	1879
Nikkei	0.02	0.13	-0.39	-0.07	0.02	0.10	0.46	220
Not Nikkei	-0.05	0.14	-1.68	-0.12	-0.05	0.03	0.63	1659
Market Beta								
Topix	0.86	0.31	-0.16	0.64	0.86	1.07	2.08	1879
Nikkei	1.08	0.22	0.46	0.91	1.08	1.23	1.71	220
Not Nikkei	0.83	0.31	-0.16	0.61	0.81	1.03	2.08	1659
BoJ Weight								
Topix	0.05	0.23	0.00	0.00	0.00	0.01	4.65	1879
Nikkei	0.38	0.56	0.01	0.11	0.21	0.41	4.65	220
Not Nikkei	0.01	0.02	0.00	0.00	0.00	0.01	0.34	1659
Nikkei Weight								
Topix	0.05	0.33	0.00	0.00	0.00	0.00	8.91	1879
Nikkei	0.45	0.86	0.00	0.10	0.21	0.44	8.91	220
Not Nikkei	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1659
Topix Weight								
Topix	0.05	0.19	0.00	0.00	0.01	0.03	4.70	1879
Nikkei	0.31	0.47	0.01	0.06	0.16	0.37	4.70	220
Not Nikkei	0.02	0.05	0.00	0.00	0.01	0.02	0.69	1659

Table A1: Summary Statistics. This table provides summary statistics for various stock characteristics by index membership. Topix stocks represent our entire sample of stocks. Nikkei stocks are those included in the Nikkei 225 index, while Not Nikkei stocks are those that only appear in the Topix index. All Nikkei companies also belong to the Topix index.

	Market Cap (bn JPY)			Market beta			Forex beta			Market Leverage (%)			Market to Book		
	mean	std	p50	mean	std	p50	mean	std	p50	mean	std	p50	mean	std	p50
As of 2013:															
High Pi	808.4	1643.5	297.8	1.04	0.24	1.03	0.00	0.12	-0.01	26.4	22.4	21.6	1.2	0.5	1.1
Medium-high Pi	118.3	195.4	53.2	0.79	0.27	0.78	-0.03	0.11	-0.04	24.1	22.8	17.2	1.2	2.0	1.0
Medium-low Pi	46.3	69.7	25.4	0.75	0.31	0.72	-0.06	0.14	-0.07	25.9	23.3	20.0	1.1	0.8	1.0
Low Pi	30.8	55.7	17.7	0.90	0.32	0.90	-0.07	0.16	-0.06	25.7	22.6	21.1	1.4	1.3	1.0
As of 2015:															
High Pi	943.1	1822.3	391.3	1.00	0.18	1.01	0.00	0.13	-0.01	23.9	22.7	18.1	1.4	0.9	1.1
Medium-high Pi	134.8	219.2	53.7	0.84	0.22	0.85	-0.04	0.12	-0.05	23.3	22.6	17.9	1.2	0.8	1.0
Medium-low Pi	66.2	110.1	30.0	0.81	0.23	0.82	-0.05	0.14	-0.06	24.5	22.6	18.2	1.2	0.9	1.0
Low Pi	44.6	58.9	22.7	0.89	0.22	0.89	-0.05	0.17	-0.05	24.1	22.0	18.2	1.5	1.3	1.1

Table A2: Summary Statistics by π -quartile. Market beta and Forex beta are estimated as explained in the main text. Market leverage is defined as $(DLTT + DLC) / (DLTT + DLC + \text{Market Cap})$. Market to book is the ratio of market assets to book assets and is computed as $(LT + PSTK - TXDITC + \text{Market Cap}) / AT$. Variables indicated with capital letters are from Compustat Global. Market capitalization is from Bloomberg.

Appendix C : 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

$$\begin{aligned} \max_N E_t (-\exp(-\gamma W_{t+1})) \\ \text{s.t.} \quad W_{t+1} = W_t(1+r) + N_t'(p_{t+1} + D_{t+1} - p_t(1+r)) \end{aligned}$$

From the first order condition it follows that

$$N_t = \frac{1}{\gamma} [\text{Var}_t(p_{t+1} + D_{t+1})]^{-1} (E_t[p_{t+1} + D_{t+1} - p_t])$$

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

$$(1+r)p_t = E_t[p_{t+1} + D_{t+1}] - \gamma \Sigma Q_t \quad (7)$$

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

$$(1+r)p_t = E_t \left[\frac{p_T}{(1+r)^{T-1}} \right] + \sum_{i=0}^{T-1} \frac{D_t}{(1+r)^i} - \gamma \Sigma \sum_{i=0}^{T-1} \frac{E_t[Q_{t+i}]}{(1+r)^i}$$

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

$$p_t = \frac{D_t}{r} - \frac{\gamma\Sigma}{(1+r)} \left(\sum_{i=0}^{\infty} \frac{\mathbb{E}_t[Q_{t+i}]}{(1+r)^i} \right) = \frac{1}{r} (D_t - \gamma\Sigma\Omega_t) \quad (8)$$

where we introduced the notation

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

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

$$p_t = \frac{1}{r} (D_t - \gamma\Sigma Q)$$

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

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

Perfect Credibility

Assuming that the central bank's commitment to the policy is fully credible means that

$$\begin{cases} \mathbb{E}_t(Q_{t+i}) = Q & \text{for } i \geq 0 \text{ and } t < 1 \\ \mathbb{E}_t(Q_{t+i}) = Q - (t+i)u & \text{for } i \geq 0 \text{ and } t = 1, \dots, M \\ \mathbb{E}_t(Q_{t+i}) = Q - Mu & \text{for } i \geq 0 \text{ and } t \geq M \end{cases}$$

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

$$\begin{aligned}
\Omega_t &= \frac{r}{1+r} \left(\sum_{i=0}^{M-t-1} \frac{Q - (t+i)u}{(1+r)^i} + \sum_{i=M-t}^{\infty} \frac{Q - Mu}{(1+r)^i} \right) \\
&= Q - \frac{r}{1+r} \left(\sum_{i=0}^{M-t-1} \frac{(t+i)u}{(1+r)^i} + \sum_{i=M-t}^{\infty} \frac{Mu}{(1+r)^i} \right) \\
&= Q - \frac{r}{1+r} \left(\sum_{i=0}^{M-t-1} \frac{(t+i-M)u}{(1+r)^i} + \sum_{i=0}^{\infty} \frac{Mu}{(1+r)^i} \right) \\
&= Q - Mu + \frac{r}{1+r} \sum_{i=0}^{M-t-1} \frac{(M-t-i)u}{(1+r)^i} \\
&= Q - Mu + \varphi(t)u
\end{aligned}$$

where we introduced the real-valued function $\varphi(t) = \frac{r}{1+r} \sum_{i=0}^{M-t-1} \frac{(M-t-i)}{(1+r)^i}$, defined for $t \geq 1$ and enjoying the following properties

- (i) $\varphi(t+1) - \varphi(t) < 0$
- (ii) $\varphi(t) < M$ for $t \geq 1$
- (iii) $\varphi(t) = 0$ for $t \geq M$

Therefore the pricing equation takes the form

$$\begin{cases} p_t = \frac{1}{r} (D_t - \gamma \Sigma Q) & \text{for } t < 1 \\ p_t = \frac{1}{r} (D_t - \gamma \Sigma (Q - Mu + \varphi(t)u)) & \text{for } t = 1, \dots, M \\ p_t = \frac{1}{r} (D_t - \gamma \Sigma (Q - Mu)) & \text{for } t \geq M \end{cases}$$

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

$$p_1 - p_0 = \frac{1}{r} (\varepsilon_1 + \gamma \Sigma (Mu - \varphi(1)u)) \quad (9)$$

and is positively related to Σu in the cross-section, since $\varphi(1) < M$. Dividing (9) by p_0 component-wise proves Proposition 1.

In the days following the announcement, as the supply is gradually reduced by the central bank open market operations, the price changes are given by

$$p_{t+1} - p_t = \frac{1}{r} (\varepsilon_{t+1} - \gamma \Sigma (\varphi(t+1) - \varphi(t))u), \quad t = 1, \dots, M \quad (10)$$

and are also positive related to Σu , since $\varphi(t)$ is a decreasing function of t . Dividing (10) by p_t component-wise proves Proposition 2.

We thus expect to observe a momentum-like effect extending the initial price impact during the days following the announcement by the central bank. This continuation effect is quantitatively smaller in size relative to the initial price adjustment.

Imperfect Credibility

We now explore the possibility that agents expectations on future supply only partially account for the announced asset purchase program. We assume that for each $t \geq 1$ there exist a scalar $\lambda_t \geq 0$ such that the time- t expectation is

$$\begin{cases} E_t(Q_{t+i}) = Q & \text{for } i \geq 0 \text{ and } t < 1 \\ E_t(Q_{t+i}) = Q - \lambda_t(t+i)u & \text{for } i \geq 0 \text{ and } t = 1, \dots, M \\ E_t(Q_{t+i}) = Q - \lambda_t Mu & \text{for } i \geq 0 \text{ and } t \geq 1 \end{cases}$$

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. Alternately, an increasing λ_t can be interpreted as a lagged reaction of investors, to be attributed to slow-moving capital, bounded rationality, slow information percolation or similar frictions.

Building on the previous section is easy to see that the expected supply is

$$\Omega_t = Q - \lambda_t Mu + \lambda_t \varphi(t)u$$

Plugging this in equation (8) we find that the price change at the announcement date is

$$p_1 - p_0 = \frac{1}{r} (\varepsilon_1 + \lambda_1 \gamma \Sigma (Mu - \varphi(1)u))$$

showing that the size of the price *jump* is increasing in the initial policy credibility λ_1 and proving Proposition 3.

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

$$p_{t+1} - p_t = \frac{1}{r} (\varepsilon_{t+1} - \gamma \Sigma (\Delta \lambda_{t+1} M - (\lambda_{t+1} \varphi(t+1) - \lambda_t \varphi(t)))u), \quad t = 1, \dots, M$$

In particular if $\Delta \lambda_{t+1} > 0$ we have a momentum effect similar to the perfect credibility case, since we have the following inequalities

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

Therefore we conclude that if $\Delta \lambda_{t+1} > 0$ for every $t = 1, \dots, M$, then we should observe a momentum effect proportional to $\pi = \Sigma u$ in the cross-section, proving Proposition 4.