

How unconventional is green monetary policy?*

Melina Papoutsis
European Central Bank

Monika Piazzesi
Stanford & NBER

Martin Schneider
Stanford & NBER

November 13, 2022

Abstract

This paper studies the environmental impact of unconventional monetary policy. Our theoretical framework is a multisector growth model with climate externalities and financial frictions. When central bank asset purchases have real effects on aggregate output, their sectoral composition typically affects the climate. Market neutrality of asset purchases does not follow from simple formulas used by policy makers, but depends on (i) the impact of central bank purchases on firms' cost of capital and (ii) the share of capital funded by bonds. We use micro data on individual bond holdings, firm characteristics and emissions to show that the ECB's corporate bond portfolio is tilted towards high emission sectors compared to the market portfolio of sectoral capital stocks.

*Preliminary, comments welcome. We thank Pete Klenow, Laura Veldkamp, Michael Woodford, Moto Yogo and seminar participants at the ASSA 2021, Bank of Spain, Chicago, Columbia, IMF, Imperial, NBER Summer Institute, and UCLA for helpful comments. The views expressed in this paper are the authors' and do not necessarily reflect those of the European Central Bank or the Eurosystem.

1 Introduction

Central bank purchases of corporate bonds have recently become an important policy tool. Their purpose is to lower firm's cost of capital and thereby stimulate investment. The implementation of these purchases requires selecting a bond portfolio. Which corporate bonds should central banks buy? The answer by policy makers emphasizes *market neutrality*, their goal to keep relative prices across assets the same. To implement their goal, central banks buy bonds in proportion to the amount of bonds outstanding. This formula helps central banks to select the maturity and country composition of their sovereign bond portfolios, and now also guides their choice of corporate bonds in a way that remains neutral on issues like firms' climate impact.

This paper develops a theoretical framework for thinking about the cross-sectional impact of unconventional monetary policy and uses it to interpret novel evidence on the ECB's bond portfolio. As in many macro studies of unconventional monetary policy, central bank bond purchases have real effects because they reduce liquidity and risk premia due to financial frictions, and hence firms' cost of capital. The focus of our paper, in contrast, is on how purchases affect the cross section of firms. We propose to define market neutrality as *keeping relative costs of capital across firms* the same. With this definition, a neutral program does not favor any one firm over another: it does not alter the market portfolio of firms' capital stocks, but has exclusively macro effects.

We show that a market neutral corporate bond purchase program typically does not exist in a modern financial system. Instead, by purchasing risky bonds and funding those purchases with safe debt, a central bank favors firms that are more bond-levered or more risky. Indeed, if eligibility for central bank purchases lowers a firm's bond spread by improving bond market liquidity, the cost of capital falls more when the firm is more bond-levered. At the same time, more safe debt lowers all risk premia in general equilibrium, which benefits riskier firms more. We show that some non-neutrality is desirable: if the central bank aims to minimize the cost of financial frictions in intermediation, firms that face larger distortions should in fact benefit more.

We then turn to a particular dimension of cross sectional firm heterogeneity: impact on the environment via greenhouse gas emissions. We first use new ECB bond holdings data together with data on emissions to show that, relative to the market portfolio, the ECB bond portfolio tilts towards brown, rather than green sectors. This is a direct consequence of the structure of the bond market and the simple bond market neutrality formula currently followed by the ECB. Firms in the energy and certain manufacturing sectors generate relatively more emissions and also issue relatively more bonds. Firms in

the service sector are cleaner, but are funded relatively more with equity and loans. We conclude that the current ECB portfolio is not market neutral according to our definition, but subsidizes firms with more tangible assets as well as riskier firms, especially those more exposed to climate risk.

Our theory has three main implications for the debate on green monetary policy. First, there is no simple formula that guarantees full market neutrality – the central bank always favors some firms, and it always takes a stance on the climate. In a world with a small number of risk factors and costly bond market access, purchase programs cannot fine-tune their cross sectional impact. Second, the central bank can nevertheless design bond portfolios that lower emissions. It can increase the relative costs of capital of dirty firms by making them ineligible for purchases, hence keeping liquidity premia for these firms high. More importantly, it can load negatively on a climate risk factor in its bond portfolio in order to increase dirty firms' risk premia.

Third, optimal policy depends on what other tools the government is using. If an optimal carbon tax is in place, an asset purchase program should focus only on minimizing financial frictions, and not take climate externalities explicitly into account. This result is reminiscent of the "principle of targeting" from optimal taxation, whereby taxing an externality directly is optimal regardless of other tax instruments. The difference here is that the other instrument used by the planner is an asset purchase program that alters the resource cost of financial intermediation in general equilibrium. Moreover, in the absence of an optimal carbon tax, green monetary policy can improve welfare.

The starting point for our theory is a multisector neoclassical growth model with heterogeneous firms that make different varieties of goods from capital and labor. We add a climate externality familiar from integrated assessment models: production generates firm-specific emissions that lower future TFP. We further assume that households invest in capital through intermediaries that hold bonds, loans, and equity issued by firms. Intermediaries must spend resources to hold these securities: holding costs are specific to the firms and securities, and capture variation in liquidity and risk. Firms choose how to optimally package their capital stock into bundles of securities. The model thus creates rich predictions for the cross section of security premia, firm capital structure, and investment.

Two market failures create a role for firm-specific policy intervention. One is the climate externality: in the absence of a carbon tax, dirty firms produce too much. The second is a financial friction: firms that are riskier or fund themselves in less liquid securities markets produce too little. Both market failures generate misallocation of capital and labor across firms: in particular, the financial friction makes aggregate TFP

too low. The central bank has a technology to improve liquidity by making firms eligible for purchase by a large player and to lower risk on private intermediary balance sheets by swapping risky private securities for safe central bank debt. This technology is not free as the central bank also incurs resource costs, but it is worth using as long as the size of the central bank is sufficiently small. The model thus delivers an optimal scale for the central bank purchase program together with its composition.

Intermediary holding costs provide a tractable, unified way to capture different forces that shape valuation in securities markets – in particular, risk and liquidity. We discipline their size and shape with data on security positions and returns. In our model, premia on securities – differences in average returns – reflect differences in marginal holding costs. In other words, premia compensate investors for holding securities that are less desirable, much like in models with explicit shocks or search frictions. In particular, while our model is formally deterministic, it nevertheless captures key effects of risk taking that are central in empirical asset pricing. Quasiconvexity of holding cost functions implies that securities are imperfect substitutes. Intermediaries' portfolio choice problems give rise to a demand system for securities that generates a factor structure in premia, as observed in the data.

The shape of holding costs describes (i) the scope for government action in alleviating financial frictions, (ii) frictions faced by firms in funding markets and (iii) the structure of the financial system. With respect to (i), the presence of holding costs does not per se imply a market failure: we provide conditions below under which holding costs simply penalize risk taking, regardless whether by the central bank or the private sector. Under those conditions, Ricardian equivalence holds and unconventional monetary policy is irrelevant. However, in the bulk of the paper we work with a cost function that is consistent with evidence that asset purchases lower return premia through liquidity and risk effects.

By making marginal holding costs depend on both the firm and the security, we incorporate heterogeneous tradeoffs between debt and equity at the firm level. This helps the model account for the cross section of security supply as well as its response to policy. In our data, dirty firms with more tangible assets, such as utility companies, are more levered. In the model, the holding cost for debt of such firms increases less with leverage so that these firms indeed issue more debt in equilibrium. The slope of the cost function thus captures how a larger recovery value of tangible assets lowers bankruptcy cost. In response to central bank purchases, firms tend to issue relatively more bonds, but also increase issuance of other securities. In the model, asset purchases lower intermediaries' relative cost of holding bonds, but also their overall cost of capital,

thus increasing overall issuance as well as bond leverage.

The second derivatives of the holding cost function describe how risk on intermediary balance sheets affects valuation and hence how the cross section of risk premia responds to policy. An important fact here is that asset purchase programs also affect yields on bonds that are not eligible for purchase, and more so if the bond is riskier. A natural explanation is that asset purchases replace risky securities by safe debt, so intermediaries are more willing to hold risky securities and risk premia decline. Key to this argument is that the same intermediary trade in eligible and ineligible securities. In our model, this effect is captured by nonseparability and quasi-convexity of holding costs across securities: more safe debt lowers the marginal cost of holding risky securities. In contrast, in a starkly segmented financial system where central bank debt, corporate bonds, loans and equity are all held by separate intermediaries, premia would not respond at all.¹

Our model clarifies a key difference between the liquidity and risk effects of a purchase program. Liquidity effects allow the central bank to target individual firms that issue bonds, but have limited impact on the cost of capital. Risk effects are larger in part because they generate spillovers across firms and securities, which makes asset purchases a blunter instrument. When eligibility for central bank purchase improves liquidity in an individual firm's bond market, this is a direct effect on the firm's cost of issuing bonds: it lowers yields controlling for risk characteristics and thus allows fine-tuning of policy impact at the firm level. However, bonds are a small share of funding: the cost of capital is a weighted average of returns on different funding sources. With a typical bond leverage of 10%, any bond market liquidity effect is muted by a factor of ten. Moreover, there is no effect for all those firms that do not have bond market access.

Risk effects of purchase programs work indirectly through the balance sheets of intermediaries. If asset demand induces a factor structure in risk premia, as the empirical literature documents, then the composition of the central bank portfolio in terms of individual securities is not important for policy impact. All that matters is how much the central bank purchases of each factor, and how much safe debt the central bank issues overall. In equilibrium, market prices of risk are determined by the risk appetite of private intermediaries given their portfolio composition. The effect on individual firms' premia is determined by their exposure to the risk factors. It is therefore generally true that when the central bank reduces market prices of risk, relatively more exposed firms benefit more.

¹In such a counterfactual system – described in our model by a separable holding cost function – policy would always be market neutral.

Liquidity and risk effects give rise to two distinct forces of non-neutrality. Liquidity effects imply that asset purchases favor bond-levered firms, whereas risk effects imply that asset purchases favor risky firms. Since the cross section of individual firm exposures is large relative to the number of factors, neutrality is generally impossible to achieve. Even if all firms had access to the bond market, keeping firm-level relative premia unchanged would mean that a large number of equations can be solved by choosing a small number of factor loadings that characterize the central bank portfolio. We emphasize that this does not mean that the central bank cannot target dirty firms. The empirical literature has documented a priced climate risk factor: dirty firms pay higher average returns. By shorting this factor – for example, purchasing green firms – the central bank can lower the risk premium on clean firms relatively more and make an impact on total emissions.

To derive optimal policy, we study a social planner who chooses a carbon tax as well as the scale and composition of a central bank purchase program. The optimal program minimizes combined private sector and central bank holding costs. Since the central bank can swap safe debt for any risk factor, the optimal portfolio should hold all risk factors in the same proportions as the market portfolio – the optimal policy is not neutral, but favors riskier firms. Unless the central bank has an advantage in managing risk, its overall risk exposure should also be lower than that of the market portfolio – the model thus provides a justification for central banks to buy low risk bonds rather than, say, junk bonds or equity. These principles apply also when there are no climate externalities – they follow from the desire to minimize financial frictions. Deviating from them makes sense only if a carbon tax is not available.

Our empirical exercise suggests that the current ECB portfolio overweighs dirty firms relative to the market portfolio and is therefore neither neutral nor optimal. The comparison is performed at the sector level. To determine the sectoral composition of the ECB's holdings, we start from confidential data on individual firm holdings that we merge to data on firm characteristics from Orbis as well as emissions data from Urgentem. A key step here is to handle the large share of financial sector bonds held by the ECB. A large majority of those bonds are issued by special purpose entities that provide funding for large nonfinancial corporations, including non-Eurozone corporations that have established SPEs to tap the European bond market. By hand-matching SPEs to their parent firms, we can reduce the share of financial sector bonds in the ECB's portfolio from 56% to 11%. We combine ECB holdings data aggregated to the NACE 2 sector level with emissions totals from Eurostat's air emission accounts.

We show that the carbon footprint of the ECB's corporate bond portfolio is larger than

that of the market portfolio. This result is robust across three alternative measures of the market portfolio: physical capital shares from Eurostat, book asset shares from Orbis and shares of market value plus debt for publicly traded firms from Orbis. In all cases, the ECB overweighs dirty sectors and underweighs cleaner sectors such as services. We show that the same weighting emerges if we consider eligible bonds. Through the lens of the model, this means that the ECB holds a portfolio that is safer than the market portfolio – consistent with optimality – but also more exposed to a climate risk factor. From our analysis of optimal policy, a central bank that either believes that the current emissions trading system resembles an optimal carbon tax or does not care about climate externalities would hold factors proportional to the market. A central bank concerned with an inadequate carbon tax would short the climate risk factor. There is thus no reason in our model for overweighing dirty firms.

Related Literature. In the low interest rate environment of the last two decades, monetary policy has adopted unconventional measures, most importantly central bank asset purchases. Initially, these purchases concentrated on government bonds, and then expanded to agency mortgage-backed securities and other agency bonds. Many papers have documented the price impact of these purchases (for example, Gagnon, Raskin, Remache and Sack, 2011, Krishnamurthy and Vissing-Jorgensen, 2011, D’Amico and King, 2013, and Krishnamurthy, Nagel and Vissing-Jorgensen, 2018). Most recently, central banks started buying corporate bonds (Mota and Papoutsis, 2020, Boyarchenko, Kovner and Shachar, 2020, Koijen, Koulischer, Nguyen and Yogo, 2021). Recent policy papers have argued that the eligibility criteria of corporate bond purchase programs by central banks favor bonds issued by companies with high emissions (Matikainen, Campiglio and Zenghelis, 2017, Schoenmaker, 2021). Our paper works with actual data on central bank bond holdings and benchmarks them to the market portfolio. Moreover, we study the market neutrality and optimality of these purchases.

Other papers study the economic mechanisms behind the price impact and, more generally, the real effects of unconventional monetary policy. The mechanism on aggregates in our paper works through the reduction of risk on financial intermediaries’ balance sheets, as in Gertler and Karadi (2011), Cúrdia and Woodford (2011), and Piazzesi, Rogers and Schneider (2020). Our paper focuses on the cross sectional impacts of these policies on heterogeneous firms. With only a few risk factors, asset purchases reduce the cost of capital of all firms, especially those with higher risk exposures. Moreover, corporate bond purchases by the central bank increase the liquidity of bonds by eligible firms. Recent work by Kurtzman and Zeke (2020) and Ferrari and Nispi Landi

(2021) also studies corporate bond purchases.² Our focus is on market neutrality and optimality of these purchases.

Our model starts from a multisector general equilibrium model and adds a climate externality: productivity depends on the climate, as in integrated assessment models by Nordhaus. However, unlike IAMs, we do not distinguish between energy goods and other goods. The climate is permanently affected by emissions which are generated by the production of various sectors in the economy. We assume linear dynamics for the climate (as in Hassler, Krusell and Smith (2016), Hambel, Kraft and van der Ploeg (2021)). The emissions may also increase firms' cost of capital, as documented in the empirical asset pricing literature (Baker, Bergstresser, Serafeim and Wurgler (2018) and Zerbib (2019)) for bonds, Hong and Kacperczyk (2009), Bolton and Kacperczyk (2021a), Bolton and Kacperczyk (2021b) for stocks.) More generally, our paper is related to the exploding literature on climate finance (for a survey, see Giglio, Kelly and Stroebele (2021); Engle, Giglio, Kelly, Lee and Stroebele (2020), Baldauf, Garlappi and Yannelis (2020), Choi, Gao and Jiang (2020)).

The rest of the paper is structured as follows. Section 2 contains the empirical results and Section 3 presents the model.

2 Empirical results

Our empirical work benchmarks the portfolio of the ECB against the market portfolio of capital holdings by sector in the euro area. We also compute the bond market portfolio, the sector shares of bonds outstanding which guide ECB asset purchases (because it wants to be market neutral.) Moreover, we measure the sector shares of eligible bonds. We then compare all these portfolios with the shares of emissions by sector.

2.1 Data

For the empirical analysis, we combine sector level data with micro-data from multiple sources, covering the period from 2016 to 2018. We restrict the empirical analysis to this period because the Corporate Sector Purchase Program (CSPP) was initiated in March 2016 and data availability does not allow us to extend our analysis after the end of 2018. First, to accurately approximate the market sector shares in the euro area we use information from Eurostat national accounts. In particular, we use gross value added, compensation of employees, and output by sector. These variables are reported at an

²Ottonello and Winberry (2020) and David and Zeke (2021) study conventional monetary with heterogeneous firms.

annual frequency based on a NACE sector code classification.

We complement the sector data with firm level data from Orbis by Bureau van Dijk. We restrict the sample to active firms in the euro area and use information on the firms' total assets, total debt, revenues, and market capitalization. In total, we use information from more than 3 million firms that operate in the euro area and construct the sectoral shares of total assets, total debt, revenues, and market capitalization. The advantage of using Orbis firm level data is that the sector distribution is representative of all firms that have accurate financial statements, regardless of their size, and we utilize more variables at the firm level. At the same time, micro level data could be more affected by outliers. For that reason, we examine the distribution of the market shares using both the Eurostat sectoral variables and the Orbis micro data.

We estimate ECB portfolio sector shares with confidential micro level data from the ECB Securities Holdings Statistics (SHS). This database reports the exact amount of each security held by the Eurosystem. In total, bonds from 355 companies were purchased under the Corporate Bond Purchase Program until the end of 2018. Specific criteria are in place to define the eligibility of a bond that can be purchased under the CSPP. These are related to the bond's original and remaining maturity, rating, and the issuing corporation.³ We construct the sector shares of the ECB portfolio based on the amount held by the Eurosystem for each issuer sector.

Similarly, to estimate the sector shares of the bond market, we use information from the ECB Centralized Securities Database (CSDB) that contains information for all securities held and transacted by euro area residents. We construct the sector shares of the bond market in the euro area using the amount outstanding at the end of each year of bonds issued by non-financial corporations incorporated in the euro area. In this construction, we exclude any privately placed bonds.

For our analysis, it is key that we match firms to their correct sector. A challenge with this match is that many large public companies issue bonds through a special purpose entity that is a financial company. For example, the oil manufacturer Royal Dutch Shell uses a subsidiary, Shell International Finance BV, to issue its bonds. The subsidiary is a company that belongs to the finance sector. If we used the raw sector classification of the bond issuer, we would find that the finance sector accounts for the majority of bonds outstanding. Moreover, we would find that finance has the largest sector weight in the ECB portfolio, accounting for approximately 55% of the portfolio. Obviously, this would

³To be eligible for purchase, securities must (a) have an investment grade rating, (b) be euro-denominated, (c) have minimum remaining maturity of six months and maximum remaining maturity of 30 years, and (d) be issued by a non-financial corporation established in the euro area.

be wrong—bonds issued by non-financial corporations are not even CSPP eligible, and so the ECB does not buy any bonds issued by financial corporations.

To obtain an accurate description of the bond market and the ECB portfolio, we re-search all financial companies individually and hand collect information about the right sector of their parent company. In particular, we hand collect the sectoral classification for all issuers that were characterized as financial companies in the ECB portfolio as well as for the 200 largest financial bond issuers in the euro area bond market. After we reassign these firms to their correct sector, we succeed to reduce the financial sector share in the ECB portfolio from 56% to 11%. For the remaining 11% we are unable to correctly identify the parent company as the issuer name was not informative enough and in most of the cases is a special purpose entity. Following a similar procedure for all bond issuers, we are able to reduce the share of the finance sector from 87% to 73%. It is not surprising that the bond market share of the finance sector is so large, because banks are the largest issuers of corporate bonds in the euro area. In our core analysis we only consider firms that are correctly classified as non-financial.

To measure the environmental impact of corporate bond purchases, we study data on emissions. These data come from two sources. First, the annual greenhouse gas Scope 1 air emissions by sector are reported in Eurostat national accounts. Scope 1 emissions capture the direct emissions associated with the production by the sector (for example, the emissions associated with car production in the automobile sector.) Our second source of emissions data is firm level information from Urgentem. The data provides measures of direct Scope 1 emissions and also indirect emissions of greenhouse gases by each firm. The firm's Scope 2 emissions include those generated by the production of energy that is used in the production (for example, Scope 2 emissions for Volkswagen include the emissions generated by the utility company that produces the electricity used by Volkswagen in its car production.) Scope 3 emissions include emissions that occur when the products of the company are used by their customers (such as the emissions associated with the driving of Volkswagen cars.) Scope 3 emissions are often the largest component of total emissions.

2.2 ECB portfolio holdings and other sector shares

We compare the ECB portfolio with various other sector shares. Importantly, we want a comprehensive measure of capital holdings, including holdings of intangible capital, by sector—the market portfolio. The ideal measure of a sector's capital would therefore be the sum of the market value of equity plus debt in the sector. Since the market value of equity is not available for private firms, we present results for several alternative

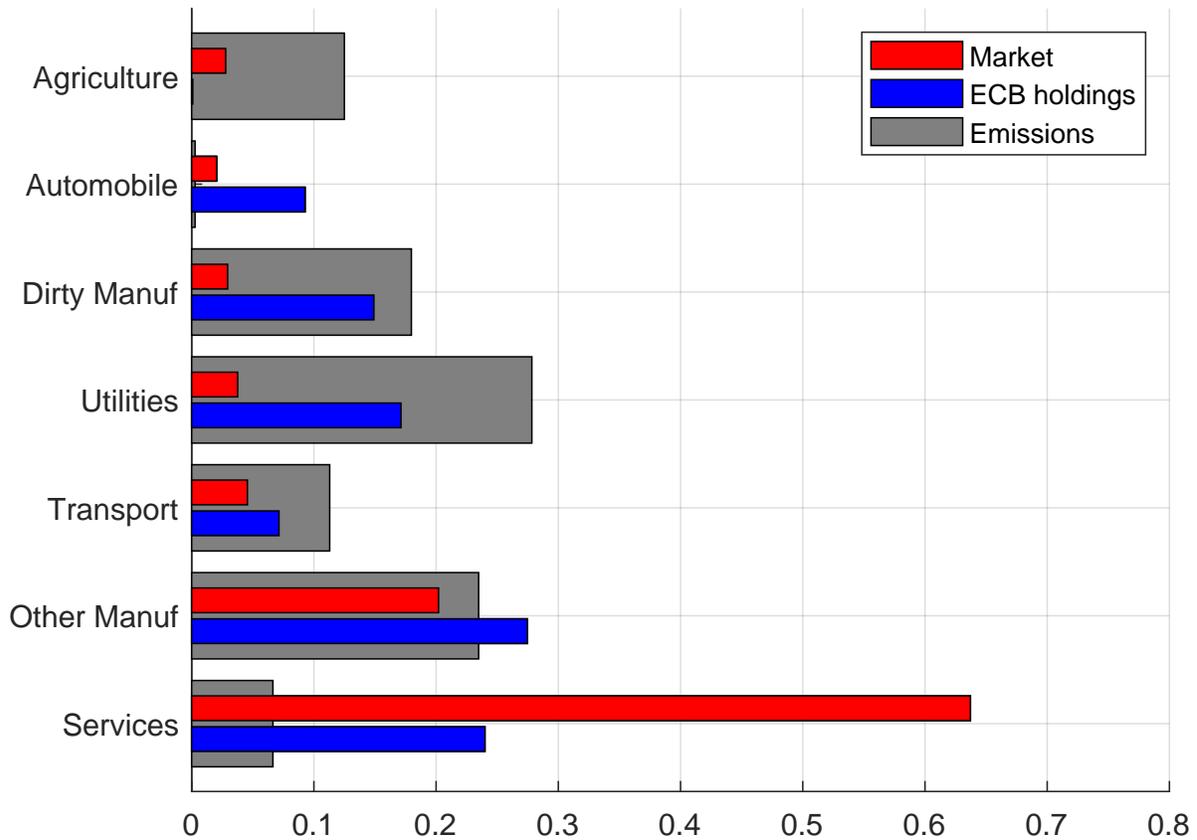


Figure 1: Sector shares of the market portfolio, ECB holdings, and emissions

This figure is constructed based on the year-end 2017 data. Market shares are measured as capital income by sector (capital income = value added – wages). Emissions are measured as Scope 1 air emissions by sector. The ECB portfolio includes only securities held under the corporate sector purchase program (CSPP) that was initiated in March 2016. By construction, all sector shares sum up to one. Data sources: SHS (ECB), Orbis, and Eurostat. Definition of sectors: Dirty Manufacturing includes: oil and coke, chemicals, basic metals, nonmetallic minerals manufacturing; Other Manufacturing includes: food, beverages, tobacco, textiles, leather, wood, paper, pharmaceuticals, electronics, electrical equipment, machinery, furniture, construction, and other manufacturing.

measures of capital holdings. Our main result is based on the share of capital income by sector. If returns on capital and depreciation rates of capital are the same across sectors, the share of capital income of a sector equals its market share. We measure capital income as value added subtracting wages by sector, using data from Eurostat. Further below, we compare our main result with two additional measures of market shares and show that our findings are consistent across these different measures. The first measure sums up book assets from Orbis within each sector, while the second measure aggregates the market value of equity plus debt with data on public firms in a sector.

Figure 1 presents our main empirical result. The red horizontal bars are the sec-

tor shares in the market portfolio, measured with data on capital income shares from Eurostat. The largest sector of the euro area economy is services. The other large sector, manufacturing, consists of many subsectors that greatly differ in their emissions. We therefore group these subsectors according to their emissions into dirty manufacturing (oil and coke, chemicals, basic metals, and nonmetallic manufacturing) and other manufacturing (food, beverages, tobacco, textiles, leather, wood, paper, pharmaceuticals, electronics, electrical equipment, machinery, furniture, construction, and other manufacturing.) The other sectors like agriculture, automobile, dirty manufacturing, utilities, and transport are small shares of the economy.

The blue bars in Figure 1 correspond to the ECB portfolio. Compared to the red bars with the market portfolio, the ECB holds a larger share of the automobile, dirty manufacturing, utilities, transport, and other manufacturing sectors and holds a smaller share of services. The grey bars in Figure 1 are the emission shares by sector. Here we show direct emissions, based on Scope 1 air emissions by sector. A simple comparison of the blue, red and grey bars reveals that the ECB overweighs sectors with more emissions (such as dirty manufacturing and utilities), while it underweighs sectors with low emissions (such as services). Specifically, when the ECB portfolio weight on a sector differs from the market portfolio, it tends to deviate in the direction of the sector's emission share. A more formal way to state the result is therefore to investigate the bias in the ECB's portfolio shares by sector, defined as the difference between the share of a sector in the ECB portfolio and its share in the market portfolio. Similarly, we define the sectoral bias in emissions as the difference between a sector's emissions share and its market share. Our main result is that the sectoral bias in the ECB's portfolio has a 95 percent correlation with the bias in emissions.

Figure 2 presents results based on alternative measures of market shares. Panel 2(a) uses book asset from Orbis. More specifically, we combine output by sector from Eurostat with the ratio of book assets to revenues from Orbis. While the red bars of other manufacturing and services now appear somewhat smaller, the market portfolio and the overall result are similar to those in our main Figure 1. More formally, when we define the ECB's portfolio bias based on this alternative measure of the market portfolio, the correlation of the ECB's bias with the bias in sectoral emission shares is 91 percent. Panel 2(b) aggregates total book assets of only those companies in a sector that are publicly traded and shows them as purple bars. For public companies, we can compute the market value of equity plus debt in Orbis by sector. The red bars in Panel 2(b) represent the market portfolio based on market values by sector. The sectoral distribution of these market shares based on only public companies is similar to the market shares in the other figures. In particular, the service sector has the largest total book assets and the largest

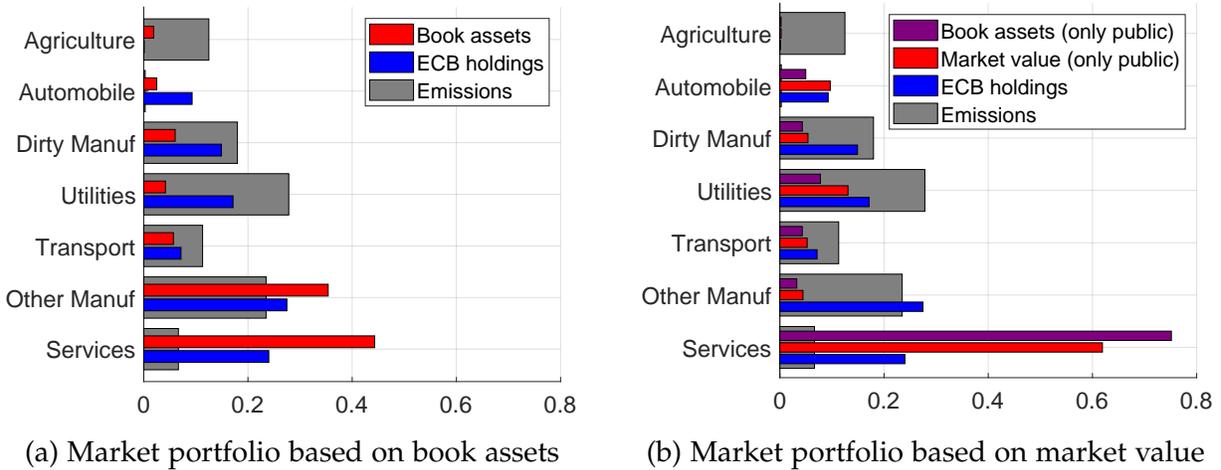


Figure 2: Sector shares of the market portfolio, ECB holdings, and emissions

This figure is constructed based on the year-end 2017 data. Panel 2(a) measures the market portfolio with book assets. We multiply output from Eurostat by the ratio of total book assets to revenues from Orbis by sector. Panel 2(b) measures the market portfolio with the total market value of public firms. We measure market capitalization plus debt for all public firms in a sector from Orbis. Total book assets for public firms are from Orbis. Emissions are measured as Scope 1 air emissions by sector. The ECB portfolio includes only securities held under the corporate sector purchase programme (CSPP) that was initiated in March 2016. By construction, all sector shares sum up to one. Data sources: SHS (ECB), Orbis, and Eurostat. Definition of sectors: Dirty Manufacturing includes: oil and coke, chemicals, basic metals, nonmetallic minerals manufacturing; Other Manufacturing includes: food, beverages, tobacco, textiles, leather, wood, paper, pharmaceuticals, electronics, electrical equipment, machinery, furniture, construction, and other manufacturing.

market value. Finally, the ECB’s bias based on the market value of public companies has a 96 percent correlation with the bias in sectoral emission shares.

To understand what explains the difference between the market portfolio and the ECB holdings, we compare the market share by sector to the bonds outstanding by sector, the bonds in the sector that are eligible for the CSPP, and the ECB holdings by sector. Figure 3 shows the four types of shares as horizontal bars. The red bars are the market portfolio. The dark red bars measure the bond market portfolio, the shares based on the total amount of bonds outstanding by sector. Similarly, the purple bars are the shares of CSPP eligible bonds based on the total amount of eligible bonds by sector. The figure shows that firms in the automobile, dirty manufacturing, utilities, transport, and other manufacturing sectors issue a relatively higher share of bonds, while firms in the service sector (and agriculture) do not issue many bonds. This striking difference is driven by the fact that large firms in carbon-intensive sectors have large holdings of fixed assets, such as machines and plants. These are long-term tangible assets that can serve as collateral.

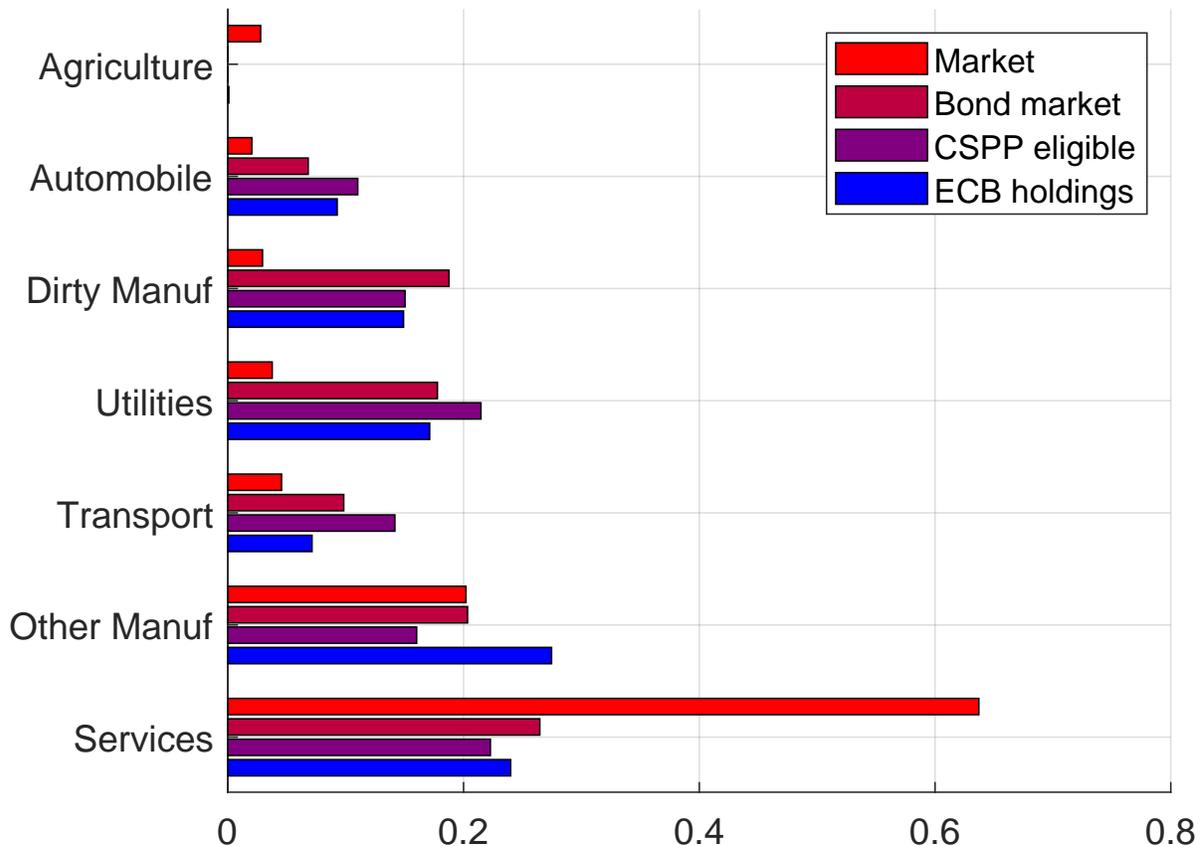


Figure 3: Sector shares of the market portfolio, bond market, CSPP eligible, and ECB holdings

This figure is constructed based on the year-end 2017 data. Market shares are measured as capital income by sector (capital income = value added – wages). The bond market shares are constructed based on the total amount of bonds outstanding by sector relative to the total. The CSPP eligible shares are constructed based on the total amount of eligible bonds by sector relative to the total. The ECB portfolio includes only securities held under the corporate sector purchase programme (CSPP) that was initiated in March 2016. By construction, all sector shares sum up to one. Data sources: SHS and CSDB (from the ECB), and Eurostat. Definition of sectors: Dirty Manufacturing includes: oil and coke, chemicals, basic metals, nonmetallic minerals manufacturing; Other Manufacturing includes: food, beverages, tobacco, textiles, leather, wood, paper, pharmaceuticals, electronics, electrical equipment, machinery, furniture, construction, and other manufacturing.

The big differences in bond issuance across sectors account for the ECB portfolio's bias towards high emission sectors. Since the ECB aims to be market neutral, its asset purchases are in proportion to the market value of outstanding bonds. The correlation between the ECB portfolio shares and the bond market shares is 92 percent. However, only few companies issue corporate bonds. Comparing the bond market portfolio in Figure 3 with the market value of public companies in Figure 2 illustrates that not even all public companies issue bonds. After all, the European corporate bond market is

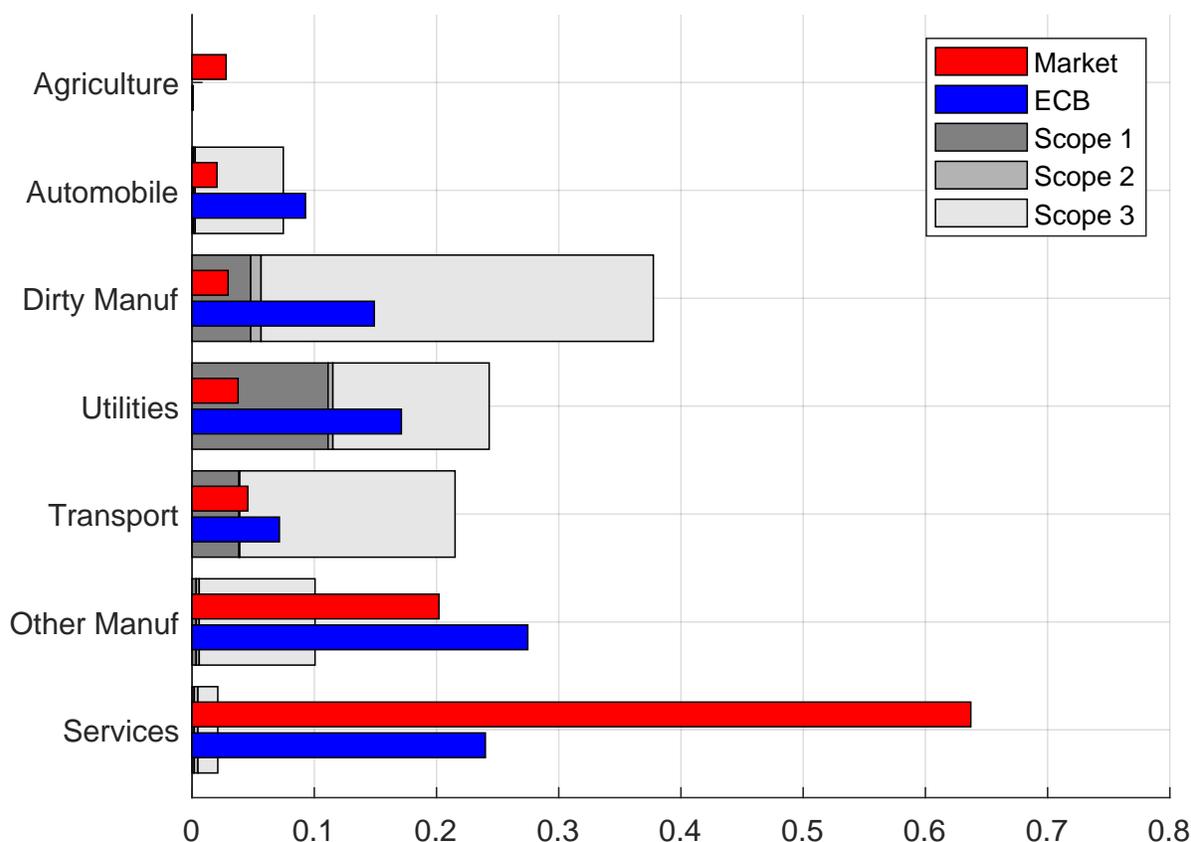


Figure 4: Market portfolio, ECB holdings and Scope 1, 2 and 3 emission intensities

This figure is constructed based on the year-end 2017 data. Market shares are measured as capital income by sector. The ECB portfolio includes only securities held under the corporate sector purchase programme (CSPP) that was initiated in March 2016. Emission intensities are defined as tons of CO₂ emissions per 10bn Euro revenue. Scope 1, 2, and 3 emission intensities as characterized in the greenhouse gas protocol (GHG) are included. The sector shares of emission intensities are constructed as before. By construction, all sector shares sum up to one. Data sources: SHS (from the ECB), Urgentem, and Eurostat. Definition of sectors: Dirty Manufacturing includes: oil and coke, chemicals, basic metals, nonmetallic minerals manufacturing; Other Manufacturing includes: food, beverages, tobacco, textiles, leather, wood, paper, pharmaceuticals, electronics, electrical equipment, machinery, furniture, construction, and other manufacturing.

relatively underdeveloped. Moreover, companies with relatively more tangible assets issue more corporate bonds. All of these reasons lead to a bond market portfolio in Figure 3 that does not look like the market portfolio. Importantly, the CSPP eligibility criteria are not distorting the ECB portfolio. CSPP eligible shares are similar to the bond market shares and to the ECB portfolio; the correlations are 90 percent and 79 percent, respectively. The bottom line is that the ECB buys in proportion to the bond market portfolio and therefore ends up with a portfolio that is biased towards sectors with more tangible assets, which tend to be more carbon-intensive.

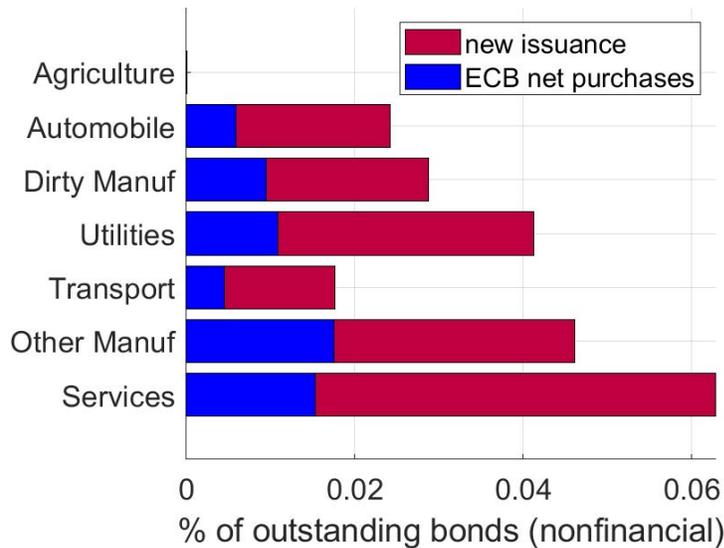


Figure 5: ECB holdings and new bond issuance

This figure is constructed based on the year-end 2017 data. The ECB net purchases include only securities held under the corporate sector purchase programme (CSPP) that was initiated in March 2016. The percent of new issuance relative to total bonds outstanding is estimated using the year-end 2017 data from the CSDB. Data sources: SHS and CSDB (from the ECB). Definition of sectors: Dirty Manufacturing includes: oil and coke, chemicals, basic metals, nonmetallic minerals manufacturing; Other Manufacturing includes: food, beverages, tobacco, textiles, leather, wood, paper, pharmaceuticals, electronics, electrical equipment, machinery, furniture, construction, and other manufacturing.

The emission shares used in the figures discussed above rely only on the direct emissions produced by each sector (Scope 1 air emissions as characterized in the greenhouse gas protocol). We also look at indirect emissions produced by firms (Scope 2 and 3 emissions based on the GHG protocol). To do so, we use firm level data from Urgentem on emission intensities by firm. Here, intensity is defined as the level of emissions per revenue. Scope 1 emissions correspond to the emissions produced directly by the firm while it operates, while Scope 2 and 3 emissions include indirect emissions linked to the firm’s usage of electricity and the use of the end product, respectively. Reassuringly, when we sum up the Scope 1 emissions for all firms in a sector using the Urgentem data, the sectoral aggregate is close to the direct emissions by sector in the Eurostat data.⁴

Figure 4 shows our result with broader measures of carbon footprint. Scope 2 emissions by sector are similar to Scope 1 emissions, suggesting that the use of electricity is not changing the sectoral distribution of emissions. However, Scope 3 emissions play a significant role as these are significantly higher for automobile, dirty manufacturing,

⁴An exception is the agriculture sector which contains many small firms that are not public and do not issue bonds. Many of these firms are not covered by the Urgentem database.

utilities, and transport, while there is only a small increase for services. Overall, Figure 4 implies that the ECB portfolio appears to overweigh carbon-intensive sectors even more when we include indirect emissions.

2.3 How important is the ECB in the corporate bond market?

To investigate whether the ECB is an important investor in the market for corporate bonds, we measure the amount of bonds that each sector newly issues as a percent of bonds outstanding. The percentage for each sector is plotted as red horizontal bar in Figure 5 and ranges from close to zero for agriculture to above 6 percent for services. The blue-shaded portions of these horizontal bars represent the newly issued bonds bought by the ECB. The figure shows that the ECB buys between 20 and 40 percent of the newly issued bonds of each sector. We conclude that the ECB is a big player in the market for corporate bonds.

3 Model

3.1 Preferences & technology

The representative household cares about consumption of a single final good. Preferences take the form

$$\sum_{t=0}^{\infty} \beta^t u(C_t), \quad (1)$$

where the discount factor β is less than one and u is increasing and strictly concave. The household also inelastically supplies one unit of labor every period.

The final good is made from intermediate goods that belong to N sectors. Varieties of intermediate goods are indexed by $i \in I := [0, 1]$. Intermediate goods in sector n are produced by a large number of firms with indices $i \in I_n \subset I$. Aggregate output is

$$Y_t = \prod_{n=1}^N Y_{t,n}^{\theta_n}, \quad \text{where } Y_{t,n} = \left(\int_{I_n} y_{t,i}^{1-\frac{1}{\sigma}} di \right)^{\frac{1}{1-\frac{1}{\sigma}}} \quad \text{and} \quad \sum_{n=1}^N \theta_n = 1. \quad (2)$$

The final good is produced with a Cobb-Douglas technology from the output $Y_{t,n}$ of various sectors n , which are themselves CES aggregates of the output $y_{t,i}$ of individual firms i . We make the standard assumption that the elasticity of substitution $\sigma > 1$ within sectors is larger than the elasticity of substitution across sectors, which is equal to one.

Production of intermediate goods generates emissions, and is in turn affected by the quality of the environment. Making $y_{t,i}$ units of good i creates $\varepsilon_{t,i} y_{t,i}$ units of emissions

(measured in tons of CO2 equivalent). The quality of the environment is described by a one dimensional index η_t . Its evolution reflects the accumulation of emissions and is described recursively by

$$\eta_{t+1} = \eta_t + \int_I \varepsilon_{t,i} y_{t,i} di. \quad (3)$$

The quality of the environment is persistent and more emissions make the planet hotter.

Intermediate good firm i produces at $t + 1$ with capital $k_{t,i}$ installed one period in advance and labor $l_{t+1,i}$ according to the production function

$$y_{t+1,i} = z_{t+1,i}(\eta_{t+1}) k_{t,i}^{\alpha_n} l_{t+1,i}^{1-\alpha_n}, \quad (4)$$

where $z_{t+1,i}(\eta_{t+1})$ is total factor productivity and n is the sector that firm i belongs to. Capital shares are the same for all firms in a sector. TFP may depend on the quality of the environment which is how emissions generate economic and hence welfare costs. Capital is made one-for-one from final goods and fully depreciates after one period.

Securities, intermediaries and holding costs. Households own capital through intermediaries that invest in securities. There is a finite number of security types $s \in S$ that differ in the holding cost that they imply for intermediaries. In our leading example below, the set S consists of equity, loans, and bonds that differ in risk and liquidity. Firm i pays for capital $k_{t,i}$ purchased at date t by issuing securities $\bar{a}_{t,i,s}$ such that its financing constraint is satisfied:

$$k_{t,i} = \sum_{s \in S} \bar{a}_{t,i,s}. \quad (5)$$

Outstanding securities are represented by a function \bar{a}_t defined on the set of firms I , so $\bar{a}_{t,i}$ is the $\#S \times 1$ vector of all outstanding securities issued by firm i . An intermediary's holdings of private securities are similarly summarized by a function a_t on I .

We distinguish between private intermediaries that issue claims held by households and the central bank that issues central bank debt held by private intermediaries. We denote the portfolio of private securities held by the central bank by \bar{a}_t . The central bank funds this portfolio by issuing central bank debt

$$\bar{d}_t = \sum_{s \in S} \int_I \bar{a}_{t,i,s} di. \quad (6)$$

A private intermediary portfolio (a_t, d_t) has private securities and central bank debt. The positions of the central bank portfolio come from a set of eligible securities: we summarize eligibility by a collection of sets $\tilde{I}_t = (\tilde{I}_{t,s})_{s \in S}$ where $\tilde{I}_{t,s}$ is the set of firms

such that the central bank is allowed to buy security s issued by firm i at date t . In our leading example below, we will restrict central bank purchases to high quality bonds.

All intermediary holdings of securities from date t to date $t + 1$ consume resources at date t in the form of final goods. In particular, a private intermediary with portfolio (a_t, d_t) incurs holding cost $h(a_t, d_t; \bar{a}_t, \tilde{I}_t)$, while the central bank with portfolio \bar{a}_t incurs holding cost $\tilde{h}(\bar{a}_t; \bar{a}_t, \tilde{I}_t)$. We assume that the functions h and \tilde{h} are homogeneous of degree one and quasi-convex in portfolio holdings (a_t, d_t) and \bar{a}_t , respectively. In other words, intermediaries operate under constant returns to scale. Moreover, isocost curves are concave which captures the idea that holding costs are lower for diversified portfolios. We further allow for an effect of central bank eligibility \tilde{I}_t on private and central bank holding costs. For example, markets may become more liquid when the central bank is a potential participant. Importantly, eligibility can lower intermediaries' cost of holding claims on firm i even if the central bank has not actually purchased such claims.⁵

Finally, we allow intermediary holding costs to depend on firm capital structure, as described by outstanding securities \bar{a}_t . The idea is that leverage entails bankruptcy costs for intermediaries. To make this concrete, we impose two assumptions on the derivatives of h . First, for every firm i , the derivatives $\partial h / \partial a_{t,i,s}$ and $\partial \tilde{h} / \partial a_{t,i,s}$ are homogeneous of degree zero in that firm's outstanding securities $\bar{a}_{t,i}$. In other words, the contribution to cost of an additional dollar in claims on firm i may depend on *ratios* of outstanding security totals (such as the firm's leverage ratio) but not on the scale of the firm.

Second, the derivatives $\partial h / \partial \bar{a}_{t,i,s}$ and $\partial \tilde{h} / \partial \bar{a}_{t,i,s}$ are homogeneous of degree one in $a_{t,i}$ and $\bar{a}_{t,i}$ respectively, but their dependence on other asset positions $a_{t,k}$ for $k \neq i$ and d_t is negligible. In other words, an intermediary that doubles its holdings of firm i claims is affected twice as much by an increase in cost due to a change in firm i 's capital structure. At the same time, a change in capital structure for one small firm k does not affect the cost of holding other assets. In our leading example below, this is because the marginal holding cost depends on exposure to a small number of risk factors, and another firm's leverage has a negligible effect on this exposure. The setup is analogous to monopolistic competition in goods markets in that demand for firm i 's securities depends on other firms' prices through a general price index, but the effect of any individual firm's price is negligible.

⁵We allow only a liquidity effect of eligibility, and not of the size of the central bank investment. More generally, it is possible that, for example, the amount of a security purchased by the central bank matters for liquidity. We omit this effect since existing evidence has focused on eligibility. It is, however, straightforward to extend the model to accommodate it. As we will see, the size of the central bank investment does matter for risk premia in our model.

Resource constraints. All securities are held either by the central bank or through private intermediaries, so $a_t + \tilde{a}_t = \bar{a}_t$. Since private intermediaries hold all debt issued by the central bank (from equation (6)), the financing constraint (5) implies that equilibrium assets of the private intermediary sector always add up to the total capital stock, denoted $K_t = \int k_{t,i} di$. A central bank purchase program thus only changes the composition of the private portfolio. The resource constraints for final goods and labor are

$$C_t + K_t + h(\bar{a}_t - \tilde{a}_t, \bar{d}_t; \bar{a}_t, \tilde{I}_t) + \tilde{h}(\tilde{a}_t; \bar{a}_t) = Y_t, \quad (7)$$

$$\int_I l_{t,i} di = 1. \quad (8)$$

The resource constraint for final goods takes into account that balance sheet costs use up final output.

Discussion. Intermediary holding costs provide a tractable way to capture how costs of liquidity and risk-taking affect lenders to firms. In particular, we would like our model to be consistent with evidence that central bank purchases lower firms' cost of capital by lowering liquidity premia on bonds and risk premia on bonds as well as other securities. Such effects have been widely studied in explicit models of search and risky borrowing. However, explicit models often include a lot of detail that is not important for the issues we study. In particular, they require specific assumptions on, for example, the nature of search in OTC markets, the protocol for default resolution and the distributions of shocks encountered by firms. Many such sets of assumptions will lead to a similar nexus between central bank policy and premia—this is what we focus on in this paper.

In principle, holding costs in our model need not reflect financial frictions. They can alternatively be interpreted as a stand-in for risk aversion of households. In this interpretation, output in the resource constraint (7) should be viewed not as measured output, but as incorporating costs of bearing risk. Measured output is then output net of holding costs. Of course, if risk aversion is the only source of holding cost, it makes sense to further assume that central bank purchases cannot add value, that is, Ricardian equivalence holds. Formally, we would assume that the equilibrium sum of costs $h + \tilde{h}$ in equation (7) is independent of the split of securities between private holdings a and central bank holdings \tilde{a} and that the eligible set \tilde{I} has no effect on costs.

More generally, our model allows for real effects of unconventional monetary policy through the financial system, as emphasized in the literature. The central bank has two distinct instruments. First, expanding the eligible set \tilde{I} makes markets more liquid and thereby lowers private intermediary holding costs. This channel is firm-specific, and hence allows narrow targeting of firms. Second, the central bank may affect private

intermediaries' holding cost by absorbing risky securities and replacing them with safe central bank debt. Quasi-convexity of h means that private holding cost may decline when a portfolio of private securities is held jointly with safe central bank debt. Intuitively, safe debt provides a buffer against risks that harm bank balance sheets, and the central bank is special as an intermediary because it can issue safe debt and therefore provide more such buffer to other intermediaries.⁶ The strength of this effect depends on the structure of the financial system, discussed further below.

We can use the shape of holding costs to clarify the role of the financial system for policy impact. As a stark example, suppose that h is separable in all security types $s \in S$. A separable cost function describes a segmented financial system where distinct intermediaries hold bonds, loans, and equity, for example. As we will see, segmentation means that spillover effects of policy across different types of securities are absent. The opposite extreme is an integrated system where, say, universal banks hold all types securities. More generally, we can allow for hybrid forms of intermediaries that specialize in a subset of securities and specify h as the sum of holdings through these technologies. The key assumption is only that intermediary technology exhibits constant returns to scale.

We can further use the shape of holding costs to interpret differences in return premia and responses to policy across firms and sectors, as well as over time. For the cross section, our empirical work confirms the familiar finding that sectors with more tangible assets issue more bonds. We can capture this in our model by making the derivative of holding costs $\partial h / \partial \bar{a}_{i,s}$ with respect to bonds of these firms flatter; when intermediaries face lower cost they are willing to lend more. Similarly, we can capture the extensive margin of bond issuance by making bond holding costs for some firms infinite. Moreover, holding costs can vary with the health of the financial system – temporary distress of financial intermediaries might lower their risk appetite, or increase their holding costs. To keep notation simple, we suppress dependence of the cost functions h and \tilde{h} on calendar time, but adding such dependence would not change the tractability of the model, as we will see.

⁶It is not essential for this effect that the central bank is a better lender than private intermediaries. In particular, consider the benchmark $\tilde{h}(\bar{a}; \bar{a}, \bar{I}) = h(\bar{a}, 0; \bar{a}, \bar{I})$, that is, holding private securities in the central bank has the same cost as holding them in private intermediaries. With quasi-convexity of h in (a, d) , we can still have that the overall cost for society to hold some private securities in the central bank is lower than holding them all in private intermediaries:

$$h(\bar{a} - \tilde{a}, \tilde{d}; \bar{a}, \bar{I}) + \tilde{h}(\tilde{a}; \bar{a}, \bar{I}) < h(\bar{a}, 0; \bar{a}, \bar{I}) + \tilde{h}(0; \bar{a}, \bar{I}).$$

3.2 Market equilibrium

We now describe a system of goods and securities markets and define equilibrium. We normalize prices so the final good is the numéraire.

Firms and production. An intermediate goods firm exists for two periods and produces good of variety i . In the first period, at date t , say, firm i issues securities $\bar{a}_{t,i}$ in order to fund purchases of capital. We denote by $R_{t+1,i,s}$ the return earned on security s issued by firm i between dates t and $t + 1$. Firms are monopolistic competitors in markets for their own securities. We choose to work with returns rather than prices: we assume that firm i faces a system of demand functions $\tilde{R}_{t+1,i,s}(\bar{a}_{t,i})$ for its securities $s = 1, \dots, S$ that we derive below. For a given vector of issued securities $\bar{a}_{t,i}$ as well as date $t + 1$ aggregates, $\tilde{R}_{t+1,i,s}(\bar{a}_{t,i})$ is the return that investors require in order to hold security s between dates t and $t + 1$. In the second period of the firm's life at date $t + 1$, capital is in place and intermediate good firms hire labor and sell goods in competitive spot markets. We write $p_{t+1,i}$ for the date $t + 1$ price of good i and w_{t+1} for the date $t + 1$ wage. An intermediate goods firm also pays a carbon tax of τ_{t+1} per unit of emissions.

Firms take as given prices of goods, wages, the quality of the environment and demand functions $\tilde{R}_{t+1,i,s}(\bar{a}_{t,i})$ for their securities. Firm i in sector n chooses capital and labor inputs as well as securities issuance to maximize

$$(p_{t+1,i} - \tau_{t+1} \varepsilon_{t+1,i}) z_{t+1,i} (\eta_{t+1}) k_{t,i}^{\alpha_n} l_{t+1,i}^{1-\alpha_n} - w_{t+1} l_{t+1} - \sum_{s \in S} \bar{a}_{t,i,s} \tilde{R}_{t+1,i,s}(\bar{a}_{t,i}) \quad (9)$$

subject to the financing constraint (5). The final goods firm buys intermediate goods in date $t + 1$ spot markets at prices $p_{t+1,n}$ and sells final goods at the price of one.

Consider an intermediate goods firm's capital structure decision. The cost of capital R_{t+1}^i of firm i is the Lagrange multiplier on its financing constraint (5).⁷ The firm's first order condition with respect to issuance of security s is then

$$R_{t+1}^i = \tilde{R}_{t+1,i,s}(\bar{a}_{t,i}) + \sum_{s' \in S} \bar{a}_{t,i,s'} \frac{\partial \tilde{R}_{t+1,i,s'}}{\partial \bar{a}_{t,i,s}}(\bar{a}_{t,i}). \quad (10)$$

Issuing security s generally affects not only the return investors require in order to hold for security s itself, but also required returns on all other securities $s' \neq s$. For example, suppose that bond issuance makes the firm riskier and hence increases the required return on bonds, as well as that on equity and other debt such as loans. When the

⁷To see this, maximize the firm's revenues minus the cost of inputs, evaluated with the cost of capital R_{t+1}^i . The firms' FOC for capital is the same as for profits (9).

security s is a bond, the last term in the first order condition (10) is thus positive. Bonds are then issued only if their return is low relative to the cost of capital, $\tilde{R}_{t+1,i,s}(\bar{a}_{t,i}) < R_{t+1}^i$. In contrast, if issuing equity makes the firm's securities less risky, the last term in (10) is negative for s being equity, and the return on equity can be relatively high.

The tradeoffs considered by the firm can also be understood using familiar intuition from multiproduct monopoly. The firm produces multiple securities that all share the same marginal cost R_{t+1}^i ; all securities package capital. Promising a higher return on security s is equivalent to charging a lower price. Continuing the example above, securities that make the firm more risky, such as bonds, work like normal goods: producing more of the security lowers demand. As a result, the firm prices such securities above marginal cost: it promises a lower return than the cost of capital. The effect is more pronounced if issuing bonds further lowers demand on other securities, much like a multiproduct monopolist charges higher prices when goods are substitutes. In contrast, producing equity stimulates demand, so equity is priced as a "loss leader" below marginal cost: it promises a higher return than the cost of capital.

Intermediate goods firm i in sector n further equates the marginal product of capital net of the carbon tax with the cost of capital, and the marginal product of labor net of the carbon tax with the wage. The final good firm equates its marginal revenues and costs. The first order conditions are

$$\begin{aligned}
(p_{t+1,i} - \tau_{t+1} \varepsilon_{t+1,i}) \alpha_n \frac{y_{t+1,i}}{k_{t,i}} &= R_{t+1}^i, \\
(p_{t+1,i} - \tau_{t+1} \varepsilon_{t+1,i}) (1 - \alpha_n) \frac{y_{t+1,i}}{l_{t+1,i}} &= w_{t+1}, \\
\theta_n \frac{Y_{t+1}}{Y_{t+1,n}} \left(\frac{Y_{t+1,n}}{y_{t+1,i}} \right)^{\frac{1}{\sigma}} &= p_{t+1,i}.
\end{aligned} \tag{11}$$

While expenditure shares on broad sectors are constant, individual firms within a sector can increase market share by charging a lower price, for example when they have lower cost due to cheaper funding sources.

Private intermediaries. Competitive private intermediaries owned by households invest in private bonds and equity as well as in central bank debt. Private intermediaries exist for two periods and issue shares, that is, claims to their asset portfolio. We write R_{t+1}^d for the rate of return on central bank debt from t to $t + 1$ and denote by M_{t+1} the price of a one period riskfree asset, that is, the discount factor used to value all assets held by

the household. Intermediaries choose assets to maximize shareholder value

$$\max_{a_t, d_t} \sum_{s \in S} \int_I (M_{t+1} R_{t+1, i, s} - 1) a_{t, i, s} di + \left(M_{t+1} R_{t+1}^d - 1 \right) d_t - h(a_t, d_t; \bar{a}_t, \tilde{I}_t). \quad (12)$$

For every security position s in firm i , intermediary shareholders pay in $a_{t, i, s}$ at date t and receive returns at date $t + 1$ that are weighted by the discount factor. In addition, they pay holding costs h on the portfolio (a_t, d_t) . Since asset pricing is linear and holding costs are homogeneous of degree one in asset holdings, the intermediary problem exhibits constant returns to scale. We can therefore work with a representative intermediary.

The intermediary first order conditions imply that returns on central bank debt as well as private securities satisfy "Intermediary Euler equations":

$$M_{t+1} R_{t+1}^d = 1 + \frac{\partial h}{\partial d_t} (a_t, d_t; \bar{a}_t, \tilde{I}_t), \quad (13)$$

$$M_{t+1} R_{t+1, i, s} = 1 + \frac{\partial h}{\partial a_{t, i, s}} (a_t, d_t; \bar{a}_t, \tilde{I}_t). \quad (14)$$

Without holding cost ($h = 0$), all premia are zero as all returns including the rate on central bank debt are equated to the riskfree rate $R_{t+1}^f = 1/M_{t+1}$, as is standard in deterministic models. More generally premia reflect marginal holding costs: intermediaries demand compensation for holding securities which increases firms' cost of capital. At the same time, intermediaries earn a convenience yield on central bank debt $\partial h / \partial d_t$ which lowers the interest rate R_{t+1}^d , possibly below the rate R_{t+1}^f on a safe asset held directly by households.

Intermediary optimization implies demand functions for individual firm securities. The demand function for security S issued by firm i is defined as

$$\tilde{R}_{t+1, i, s}(\bar{a}_{t, i}) := M_{t+1}^{-1} \left(1 + \frac{\partial h}{\partial a_{t, i, s}} (a_t, d_t; \bar{a}_t, \tilde{I}_t) \right). \quad (15)$$

The firm anticipates that when it changes its own capital structure $\bar{a}_{t, i}$, there is only a negligible effect on required returns of other firms. Moreover, it takes as given the portfolio a_t of the representative intermediary.⁸

Households and government. Households invest in private intermediary shares with

⁸The firm does so even though the portfolio a_t contains firm i securities – the firm believes that its own scale does not change intermediaries' marginal cost of holding its securities. This belief is sensible when intermediary costs depend on systematic risk or sectoral exposure, and the firm is small relative to the financial system. Our leading example below illustrates this point.

riskfree return R_t^f and receive (possibly negative) lump sum transfers T_t from the government. The government raises carbon taxes and uses those, together with the returns on its investments, to pay for its balance sheet costs, interest rate payments on past debt issues, and transfers. The budget constraints of households and the government are then

$$C_t + S_t = R_t^f S_{t-1} + w_t + T_t, \quad (16)$$

$$\tau_t \int_I \varepsilon_{t,i} y_{t,i} di + \sum_{s=1}^S \int_I R_{t,i,s} \tilde{a}_{t-1,i,s} di = \tilde{h}(\tilde{a}_t; \bar{a}_t, \tilde{I}_t) + R_t^d \bar{d}_{t-1} + T_t, \quad (17)$$

where S_t is household savings in intermediary shares.

Household optimization equates the price of a risk-free asset $M_{t+1} = 1/R_{t+1}^f$ to the household's marginal rate of substitution:

$$M_{t+1} = \beta \frac{u'(C_{t+1})}{u'(C_t)}. \quad (18)$$

This is the standard Euler equation in representative agent models. We emphasize that intertemporal substitution as captured by the Euler equation is only one force that affects the return on capital and therefore investment in our model. Holding costs further introduce a return premium through the first-order conditions of intermediaries (14). Therefore firms face a higher cost of capital R_{t+1}^i in their first order conditions (11).

Equilibrium. Given a sequence of central bank portfolios \tilde{a}_t and central bank debt \bar{d}_t that satisfy the central balance sheet condition (6), a *feasible allocation* consists of sequences for consumption C_t , firm capital stocks k_t , securities issues \bar{a}_t and a private intermediary investment portfolio (a_t, d_t) together with firms' labor inputs, outputs and emissions such that the resource constraints (2)-(8) hold and securities markets clear, that is $a_t + \tilde{a}_t = \bar{a}_t$ and $d_t = \bar{d}_t$. Given a carbon tax τ_t , an *equilibrium* is a feasible allocation as well as sequences for prices $\{p_{i,t}\}_{i \in I}$, asset returns R_t^f, R_t^d and $\{R_t^i\}_{i \in I}$, and transfers T_t such that firms maximize profits (9) subject to the financing constraint (5), private intermediaries maximize shareholder value (12), households maximize utility (1) subject to their budget constraints (16), and the government budget constraint (17) holds.

3.3 Characterizing equilibrium without a carbon tax

In this section we characterize equilibrium without a carbon tax, that is, $\tau_t = 0$. The only government policy is therefore the central bank portfolio \tilde{a}_t that is funded by central bank debt \bar{d}_t . We use the notation $\kappa_{t,i} = k_{t,i}/K_t$ to refer to the share of firm i 's capital in firm i , so $\kappa_t = k_t/K_t$ represents the *market portfolio*. For a given government policy,

equilibrium can be conveniently characterized in two steps. The first shows how banks and firms allocate resources within a period, choosing a market portfolio κ_t and how it is packaged into individual securities. The second then shows the macro and climate dynamics. The system of equations is block recursive in the sense that the effect of policy on firm shares can be determined independently of the macro and climate dynamics of the economy.

Consider first how capital is packaged into securities in equilibrium. Firm capital structure is determinate only if the cost of raising funds $\sum_{s \in S} \bar{a}_{t,i,s} \tilde{R}_{t+1,i,s}(\bar{a}_{t,i})$ is not independent of firm security shares $\bar{a}_{t,i}/k_{t,i}$. From our definition of securities demand (15), this is a restriction on private intermediaries' holding cost h . For example, if h were completely independent of its third argument \bar{a} , then any split of capital into individual securities costs the same. As a result, the Modigliani-Miller theorem holds.

We focus on the case where the shape of h guarantees interior solutions for capital structure. Substituting the demand for firm i 's securities (15) into its cost of capital (10), firm i 's optimal funding $\bar{a}_{t,i}$ satisfies

$$M_{t+1}R_{t+1}^i = 1 + \frac{\partial h}{\partial a_{t,i,s}} (\bar{a}_t - \tilde{a}, \bar{d}_t; \bar{a}_t, \tilde{I}_t) + \sum_{s' \in S} \bar{a}_{t,i,s'} \frac{\partial^2 h}{\partial a_{t,i,s'} \partial \bar{a}_{t,i,s}} (\bar{a}_t - \tilde{a}, \bar{d}_t; \bar{a}_t, \tilde{I}_t). \quad (19)$$

Firm i issues securities $\bar{a}_{t,i}$ not simply to equate returns (or marginal holding costs $\partial h / \partial a_{t,i,s}$) but it also takes into account that issuance of one security alters required returns on other securities along the securities demand curves (15). This price impact is captured by the last term, with cross-elasticities of demand derived from the second derivatives of the holding cost function; it is zero if $\bar{a}_{t,i}$ does not affect cost.

Intuitively, optimal capital structure thus trades off direct differences in premia, or marginal holding cost, against effects of issuance on cost. As a simple example, suppose that there are only two securities, debt and equity. A direct difference is, say, that regulation makes equity more costly to hold than debt, so $\partial h / \partial a_{t,i,s}$ is larger for equity than for debt. At the same time, leverage generates default risk and hence bankruptcy costs: issuing more debt thus makes both debt itself and equity more costly to hold – the sum of cross partials in (19) is positive – whereas issuing more equity reduces costs of holding both securities and cross partials are negative. The firm then chooses leverage as to balance extra bankruptcy cost to the regulatory benefit.

When we multiply the cost of capital (19) of firm i by the share of its capital financed

by each security $\bar{a}_{t,i,s}/k_{t,i}$ and sum over all securities, we get

$$M_{t+1}R_{t+1}^i = 1 + \sum_{s \in S} \frac{\bar{a}_{t,i,s}}{k_{t,i}} \frac{\partial h}{\partial a_{t,i,s}} (\bar{a}_t - \tilde{a}_t, \bar{d}_t; \bar{a}_t, \tilde{I}_t). \quad (20)$$

This result uses that the sum of all securities (5) equals firm i ' capital, $\sum_{s \in S} \bar{a}_{t,i,s} = k_{t,i}$. Moreover, marginal holding costs are scale independent: $\partial h / \partial a_{t,i,s}$ are homogeneous of degree zero in $\bar{a}_{t,i}$. Therefore, issuing more securities does not affect the firm's cost of capital: the sum of cross-partials in (19) weighted by the share of capital financed by security s are zero when summed over all s . Without holding costs, $h = 0$, the cost of capital is the same for all firms and equal to the real interest rate $R_{t+1}^f = 1/M_{t+1}$. More generally, the cost of capital may be higher or lower with marginal holding costs, depending on whether holding more securities of firm i increases or lowers intermediaries' overall holding cost of securities.

Consider now the allocation of capital across firms. Combining firms' first order conditions (11), firm capital shares κ_t equate marginal products and costs of capital.⁹ Within sector n , capital responds to firm-level TFP and the cost of capital

$$k_{t,i} \propto z_{i,t+1} (\eta_{t+1})^{\sigma-1} \left(R_{t+1}^i \right)^{-(\alpha_n(\sigma-1)+1)} \quad \text{for } i \in I_n. \quad (22)$$

Since the elasticity of substitution σ between goods within a sector is larger than one, firms that can fund investment more cheaply increase their market share. This effect is stronger the larger is the capital share in the sector. If a firm in sector n enjoys a 1% lower cost of capital than other firms in the sector, then the firm buys $\alpha_n(\sigma - 1) + 1$ percent more capital, hires $\alpha_n(\sigma - 1)$ percent more labor and produces $\alpha_n\sigma$ percent more output. The allocation of capital across sectors ensures that expenditure on capital in sector n is

⁹Inserting the price p_i from the FOC of the final goods producer into the FOC for capital of firm i provides an expression for the ratio Y_n/y_i of output in sector n to output of firm i . Together with the FOC for labor to substitute out for labor l_i , we get that the capital share $\kappa_{t,i}$ of firm i in sector n solves the equation

$$\frac{\alpha_n \theta_n}{\kappa_{t,i}} \frac{Y_{t+1}}{K_t} \frac{\left(\kappa_{t,i} z_{t+1,i} (\eta_{t+1}) \left(R_{t+1}^i \right)^{1-\alpha_n} \right)^{1-\frac{1}{\sigma}}}{\int_{I_n} \left(\kappa_{t,j} z_{t+1,j} (\eta_{t+1}) \left(R_{t+1}^j \right)^{1-\alpha_n} \right)^{1-\frac{1}{\sigma}} dj} = R_{t+1}^i. \quad (21)$$

Without holding costs, capital differs across firms only because of technology and preferences on the left hand side of (21). More generally, we obtain deviations from this benchmark so firms with lower cost of capital (lower R_{t+1}^i) have more capital, and firms with higher cost of capital have less. In the simple Cobb-Douglas case with $\sigma = 1$ capital shares are proportional to $\alpha_n \theta_n / R_{t+1}^i$: firms buy more capital if their cost of capital is lower. With higher substitution elasticity $\sigma > 1$, firms that are more productive or have lower cost of capital can compete away market share from other firms, which amplifies the effect of the cost of capital on investment.

a share $\alpha_n \theta_n$ of aggregate output

$$\alpha_n \theta_n Y_{t+1} = \int_{J_n} k_{t,i} R_{t+1}^i di. \quad (23)$$

Equations (22)-(23) determine firm level capital shares $\kappa_{t,i} = k_{t,i}/K_t$. Together with the optimal issuance of securities (19) and the financing constraint (5) we have $2\#S + 1$ equations per firm to determine shares $\bar{a}_{t,i,s}/k_{t,i}$ of individual securities and costs of capital R_{t+1}^i , all independently of aggregate capital and output.

We emphasize that intermediary holding costs can lead to either over- or underinvestment in capital at the firm level. In typical macroeconomic models with financial frictions in firms, the marginal funding cost is always higher than the rate of return on savings, so the return on capital is too high relative to the first best, and there is an underinvestment problem. The same is true here if there is only one firm and no central bank debt. With heterogeneous firms, however, some firms may have a marginal funding cost below the interest rate on savings. As a result, those firms invest more than in the first best. This is because low cost sectors may make it cheaper for the financial system to hold capital of higher cost sectors.

Aggregates. For given sector weights and total capital K_t , the bank and firm first-order conditions further give rise to an aggregate production function

$$Y_{t+1} = Z_{t+1} K_t^{\bar{\alpha}} L_{t+1}^{1-\bar{\alpha}}, \quad Z_{t+1} = \prod_{n=1}^N \frac{((1-\alpha_n)\theta_n)^{(1-\alpha_n)\theta_n}}{(1-\bar{\alpha})^{1-\bar{\alpha}}} \left(\int_{I_n} \left(\kappa_{t,i}^{\alpha_n} z_{t+1,i} \right)^{\frac{\sigma-1}{\alpha_n(\sigma-1)+1}} di \right)^{\frac{\alpha_n(\sigma-1)+1}{\sigma-1} \theta_n}, \quad (24)$$

where $\bar{\alpha} = \sum_{n=1}^N \alpha_n \theta_n$ is the (output-weighted) average capital share. Without risk or financial frictions, costs of capital $R_{t+1,i}$ are equated across firms and measured aggregate TFP Z_{t+1} reflects only firm-level TFPs $z_{t+1,i}$. More generally, holding costs matter because the capital share $\kappa_{t,j}$ depends on firm-specific costs of capital by (22). The weighted average of firm-level TFPs $\kappa_{t,i}^{\alpha_n} z_{t+1,i}$ is TFP in sector n . It captures both risk adjustment – efficiently shifting capital towards safer sectors – and misallocation due to financial frictions in intermediation.

The macro dynamics now follow from combining

$$\beta \frac{u'(C_{t+1})}{u'(C_t)} \bar{\alpha} \frac{Y_{t+1}}{K_t} = 1 + \sum_{n=1}^N \sum_{s \in S} \int_{I_n} \frac{\bar{a}_{t,i,s}}{K_t} \frac{\partial h}{\partial a_{t,i,s}} (\bar{a}_t - \tilde{a}_t, \bar{d}; \bar{a}_t, \tilde{I}_t) di =: 1 + MC_t \quad (25)$$

$$Y_t = C_t + K_t + h(\bar{a}_t - \tilde{a}_t, \bar{d}; \bar{a}_t, \tilde{I}_t) + \tilde{h}(\bar{a}_t; \bar{a}_t, \tilde{I}_t) =: C_t + K_t (1 + AC_t), \quad (26)$$

where we have defined marginal and average holding costs for the aggregate capital

stock MC_t and AC_t , respectively. The first equation is an intertemporal Euler equation for aggregate capital; the appropriate price of capital is one plus MC_t , the weighted average of individual security marginal holding costs. The second equation is the resource constraint; here average holding costs drive a wedge between capital and savings. The homogeneity properties of h and \tilde{h} imply that MC_t and AC_t depend only on the within-period allocation of capital. In a deterministic frictionless benchmark, both are equal to zero. In an equilibrium without a central bank purchase program, we have $AC = MC$ since h is homogeneous of degree one in portfolio holdings.

Simple special case: log utility and lognormal cost of capital. A particularly simple characterization obtains with log utility $u(C) = \log C$, constant holding costs, and a central bank purchase program with constant ratios $(\delta, \tilde{a}/K)$. The allocation of capital across firms, and hence constant marginal and average holding costs. We obtain a constant savings rate and simple evolution of capital

$$s := \frac{Y_t - C_t}{Y_t} = \bar{\alpha}\beta \frac{1 + AC_t}{1 + MC_t},$$

$$K_t = \frac{\bar{\alpha}\beta}{1 + MC_t} Y_t.$$

In an equilibrium without a central bank purchase program, $AC = MC$ implies the savings rate $s = \bar{\alpha}\beta$ familiar from the standard growth model. Holding costs – whether they capture risk or financial frictions – have income and substitution effects on consumption and savings that cancel out with log utility. The dynamics of the model without central bank are therefore those of a growth model with an extra loss of capital of MC percent every period.

We emphasize that intermediary holding costs can lead to either over- or underinvestment in capital at the firm level. In typical macroeconomic models with financial frictions in firms, the marginal funding cost is always higher than the rate of return on savings, so the return on capital is too high relative to the first best, and there is an underinvestment problem. The same is true here if there is only one firm and no central bank debt. With heterogeneous firms, however, some firms may have a marginal funding cost below the interest rate on savings. As a result, those firms invest more than in the first best. This is because low cost sectors may make it cheaper for the financial system to hold high cost sectors.

3.4 Central bank policy and the allocation of capital

We now use our framework to interpret our empirical findings and compare notions of market neutrality that are important in policy discussions. We note first that our model generates predictions for all data moments featured in Section 2. Formula (21) determines the market portfolio κ_t of firms' capital shares. We can aggregate these shares to the sector level to obtain a model counterpart of the market portfolio in Figure 1. Moreover, if we make bonds one of the securities in S , we can speak to the role of eligibility requirements and sectoral bond issuance. The model explains differences in bond shares by sector with differences in holding cost parameters. In particular, firms in dirty industries have more tangible assets and hence face flatter cost functions that make bonds cheaper relative to other financing.

Policy and the cost of capital. The central message of our model for policy is that any real effects of central bank policy must work through firms' cost of capital (20). To understand policy impact, it is useful to distinguish partial from general equilibrium effects. By changing eligibility requirements \tilde{I} , the central bank can directly affect the cost of capital of individual firms. The size of this effect on the cost of capital depends on (i) the firm's initial bond leverage and (ii) how the firm adjusts its capital structure in response to the policy. This is a partial equilibrium effect: it can be studied by considering only the # S markets for an individual firm's securities.

There are two distinct general equilibrium effects. First, central bank purchases \tilde{a} and debt issuance \tilde{d} alter intermediaries' marginal holding costs which affects premia on securities through the intermediaries' first-order conditions (14). The nature and size of these effects depend crucially on the structure of financial system; the key question is whether changes in risk exposures of intermediaries lead to spillovers across securities markets. The model captures these features through the shape of the holding cost function h . Second, policy generally changes society's overall holding costs in the resource constraint (7) and thereby the real interest rate R_{t+1}^f . The real rate affects all securities and hence all costs of capital equally.

Partial equilibrium: eligibility and bond leverage. While eligibility requirements provide the most direct way to channel funds to individual firms, their impact is limited by both firms' access to the bond market and the amount of their bond leverage. Suppose that including a firm i in the eligible set \tilde{I} lowers private intermediaries' marginal holding cost of bonds issued by that firm. By the intermediaries' first-order condition (14), the required return on the firm's bonds declines. In principle, this is a way for the central bank to guide funding towards a particular firm. However, the central bank

cannot affect the cost of capital of firms that do not issue bonds. Moreover, the effect on the firm's cost of capital (20) multiplies any bond price effect by the share $\bar{a}_{t,j,s}/k_{t,j}$ of bonds in total firm assets, unless eligibility has a direct effect on the firm's other funding sources as well.

The model also predicts that if the central bank focuses purchases on a particular type of securities, such as bonds, then firms will tend to issue more of those securities, which alters overall costs. For example, if a firm's bonds become eligible for central bank purchases, the firm's additional bond issuance increases its risk, and purchase programs make the economy riskier. At the same time, the effect of capital structure choice on the cost of capital and hence capital allocation is small relative to the effect on capital structure itself. Both points follows from the equilibrium condition for capital structure (19). It says that when the marginal cost on a security shifts, say because the firm's bonds become eligible, firms respond by adjusting their securities outstanding $\bar{a}_{t,j}$. However, by the envelope theorem the marginal effect on the cost of capital (20) is negligible. Therefore, we can read off the real effects of the purchase program from direct impact on marginal costs weighted by the level of leverage before the policy change.

General equilibrium: premia and spillovers. The impact of central bank purchases depends not only on the portfolio purchased, but also on the safe debt that funds it. Any program works through two channels, apparent from the marginal holding cost that determines the cost of capital (20). The first channel is absorption of risk: when the central bank purchases \tilde{a} , private intermediaries are left with private security holdings $\bar{a} - \tilde{a}$. If h is strictly convex in holdings of some risk factor or sector, say, private intermediaries now hold fewer securities that are costly to hold and therefore have lower marginal cost. The second channel is issuance of riskfree central bank debt \bar{d} that enters private intermediaries' portfolios. If h is strictly quasi-convex in private and central bank securities, the transformation of risky securities into safe debt by the central bank additionally lowers marginal cost.

We would generally expect spillovers to securities that the central bank does not actually purchase. This is because the effects of central bank purchases work through the marginal holding cost of intermediaries. As long as holding costs are not fully separable in individual securities, changes in private intermediaries' holdings of one security affect their marginal cost of holding other securities. Policy that affects premia in one market therefore spills over to premia in other markets. In our leading example below, this is a key force because intermediaries' marginal cost of risk taking depend on their exposure to a small number of risk factors.

Policy impact is generally stronger the more integrated is the financial system. In our model, integration corresponds to a nonseparable function h : the marginal cost of holding one security depends on the holdings of other securities. In a completely segmented system, where each security type is held by a different intermediary, we would have neither of the two spillover effects above: securities that are not purchased would be unaffected by what happens in markets where the central bank is active, and no other security price would respond to issuance of central bank debt.

Market neutrality and portfolio rules. What is the appropriate definition of market neutrality? Discussion of neutrality often emphasizes that policy should not favor certain firms over others – the goal of the purchasing program is macroeconomic stimulus, not altering the market portfolio κ_t . The determination of market shares (21) suggests that a natural definition of neutrality would be based on *relative costs of capital*. Formally, a *laissez-faire equilibrium* is an equilibrium with no central bank purchasing program, or $\delta_t = 0$. We say that a purchase program is *market neutral* if relative costs of capital R_t^i across firms in the equilibrium that features the purchase program are the same as in the *laissez-faire equilibrium*. Under a market neutral policy, the market portfolio is the same as in the *laissez-faire equilibrium*. Nevertheless, output and total capital may be higher if the purchase program lowers the real interest rate.

In practice, the desire to be market neutral leads policy-makers to formulate simple purchasing rules. These rules are usually stated in terms of portfolio shares, for example holding the market portfolio, or holding bonds in proportion to bonds outstanding, the policy pursued by the ECB. Our model allows us to assess when such policies are in fact neutral. We assume that bonds b are one security in the set S , and directly write $b_{t,i}$ for the component of firm securities $a_{t,i}$ that corresponds to bonds. We further define the *bond market portfolio* by a function \bar{b}_t on set of firms I ; it is equal to zero for all firms that do not issue bonds. A purchase program that holds the bond market portfolio satisfies $\tilde{b}_t = \phi \bar{b}_t$ for some scalar ϕ .

Can buying the bond market portfolio be neutral? Our discussion above of how costs of capital are determined already illustrated several sources of non-neutrality. First, not all firms issue bonds: if central bank purchases selectively lower marginal costs of holding bonds this favors firms with more bond leverage $\bar{b}_{t,i}/k_{t,i}$. Second, any nonseparabilities in the holding cost h work against neutrality. For example, if there are spillovers from bond purchases to the marginal cost of loans or equity, then bond purchases favor firms for which such spillovers are stronger. The same concern applies when there are spillovers from issuance of central bank debt to the marginal cost of any security

–purchased by the central bank or not. We conclude that a simple rule that holds the bond market portfolio will in general not be market neutral.

It is possible to write down a holding cost function such all these sources of nonneutrality are absent. However, such a function describes a very special financial system. Suppose that private and central bank bonds are each held by a distinct specialized intermediary that holds no other securities. Formally, define

$$h(a, d; \bar{a}, \bar{I}) = h^b(b; \bar{I}) + h^d(d) + h^{-b,d}(a; \bar{a}, \bar{I}), \quad (27)$$

where the functions h^b and h^d describe holding costs for the private bond and central bank debt intermediaries, respectively, and $h^{-b,d}$, which describes holding costs for all other securities, is independent of b and d .

In a financial system with technology (27), central bank purchases do not affect premia on securities, their only impact is on the real riskfree rate R^f which affects all costs of capital equally. Since h is homogeneous of degree one in a , we must have that h^b and h^d are homogeneous of degree one in b and d , respectively. In equilibrium, a purchase program \bar{b} that buys proportionally to bonds outstanding $\phi \bar{b}$ satisfies $b = \bar{b} - \bar{b} = (1 - \phi) \bar{b}$. Since the marginal cost of holding private bonds is homogeneous zero in b , it follows that marginal cost do not depend on the size ϕ of the purchase program. Therefore, bond market premia (14) also do not depend on ϕ . Since the purchase program does not affect the intermediaries who invest in all other private securities, spreads between the costs of capital and the riskfree rate are also unaffected. Even if purchases leave premia unchanged, they can have real effects if central bank interventions lower overall holding cost. The purchases program then works like a positive injection of resources which lowers the riskfree interest rate.

While the cost function (27) implies that a purchase program that buys the bond market portfolio is indeed market neutral, it is not consistent with evidence on the price effects of central bank purchase programs. Empirical studies have established that the main impact comes from changes in risk premia on bonds that differ across firms, as opposed to common changes in the riskfree interest rate only. Our leading example in the next section is designed to further characterize holding costs that capture the deviations from neutrality that have been found in the data.

Can buying the market portfolio be neutral? We emphasize that the same line of argument also implies that other simple portfolio rules are generally not market neutral. For example, suppose the central bank were to purchase bonds so that its portfolio shares correspond to the market portfolio κ_t . Such a rule does not guarantee market neutrality:

not all firms issue bonds and if bond purchases are funded with safe central bank debt, there will generally be heterogeneous spillovers across firms.

4 Liquidity and risk exposure

We now introduce our leading example of holding costs that are motivated by evidence on liquidity and risk premia earned on securities – in the model, premia reflect marginal balance sheet costs. The holding cost functions capture an integrated financial system, in which intermediaries invest in various firm securities and central bank debt. We introduce a small number of risk factors and describe security payoffs as being exposed to these factors. By investing in these securities, intermediaries take the risk onto their balance sheets which is costly.

4.1 Cost functions

Firms can issue three securities: the set $S = \{b, \ell, e\}$ contains bonds b , loans ℓ , and equity e . We represent a portfolio of private securities by $a_t = (b_t, \ell_t, e_t)$.

Risk exposure. Intermediaries' cost of risk taking depends on their portfolios only through exposure to a small finite number of risk factors $f \in F$. The idea is that intermediaries are well-diversified so that idiosyncratic firm-level risk washes out. We summarize the contribution of security s issued by firm i to the overall cost of risk taking by an $\#F \times 1$ vector $\omega_{i,s}(\bar{a}_{t,i})$, the security's exposure to the risk factors. The exposure of a portfolio of securities (a, d) with positions that sum to total investment A is then the average of the individual exposures

$$\omega(a, d; \bar{a}) = \sum_{s \in S} \int_I \frac{a_{i,s}}{A} \omega_{i,s}(\bar{a}_i) di. \quad (28)$$

Since we treat central bank debt as riskfree, it does not contribute to portfolio risk exposure – in fact, purchasing central bank debt *reduces portfolio risk exposure* since it dilutes the risk in other positions. Intermediaries accumulate exposure by lending via bonds or loans, as well as by purchasing equity.

Costs of risk taking. We now define the cost of risk taking for a private intermediary with portfolio (a, d) and total investment A as

$$\frac{1}{2} \gamma \left(\omega(a, d; \bar{a})^\top \omega(a, d; \bar{a}) \right) A,$$

where the scalar function γ is strictly increasing and convex, and penalizes total risk on

the books of the private intermediary. Convexity can help capture threshold effects in risk sensitivity: for example, γ might be linear over some range, and then sharply curve upward, say as risk reaches a critical point defined by the regulatory regime.

As we will see, the factor structure makes the model consistent with how observable characteristics predict average excess returns in the cross section. We highlight two additional properties of the functional form. First, the quadratic form in ω makes the cost of risk taking quasi-convex. It makes sense for intermediaries to diversify across factors, and to divide the portfolio between riskfree central bank debt and risky private securities – both strategies lower cost. Moreover, since the number of factors is small relative to the variety of securities, there are typically many ways to generate a target vector of factor exposures by selecting individual securities. In particular, central bank purchases affect private banks' cost of risk taking only by reducing their factor exposure, and many different central bank portfolios will have similar effects on the cost of risk taking.

Individual securities' risk exposures and firm leverage. We allow the risk exposure of a security $\omega_{i,s}(\bar{a}_{t,i})$ to depend both on the identity of the issuing firm i and on its capital structure choice $\bar{a}_{t,i}$. On the one hand, we want to capture risk of the firm's technology that would be relevant even if the firm were all equity-financed. For example, a firm in a more cyclical sector has higher exposure to a risk factor that captures the business cycle, which will result in a higher equilibrium risk premium. Similarly, a firm that is more exposed to a separate climate risk factor will have to pay a "pollution premium". On the other hand, we want to capture the effect of leverage through default risk: more levered firms have both more risky debt and more risky equity. Moreover, we want to allow for deadweight costs of default: the risk exposure of a portfolio of all debt and equity issued by a firm should be more risky than holding the capital.

We assume that equity is the riskiest security, followed by bonds and loans. For all firms and regardless of capital structure, the three functions $\omega_{i,s}$ are ranked as

$$\omega_{i,e}(\bar{a}_{t,i}) > \omega_{i,b}(\bar{a}_{t,i}) > \omega_{i,\ell}(\bar{a}_{t,i}). \quad (29)$$

Equity is the riskiest security since shareholders are residual claimants. In a standard model with explicit shocks, debt securities have a safe payoff in at least part of the state space. The ranking of bonds and loans draws on recent work in corporate finance. The idea here is that default is relatively more costly when bondholders are involved, since the presence of many creditors make renegotiation more costly. In contrast, bank loans tend to be more flexible in default.

We further make assumptions on the shape of the functions $\omega_{i,s}$ to describe how the riskiness of securities increases with firm leverage. We first assume that all functions $\omega_{i,s}$ are homogeneous of degree zero in the outstanding securities $\bar{a}_{t,i}$ of firm i – only leverage *ratios* matter for exposure. Moreover, all functions $\omega_{i,s}$ are increasing in loans and bonds, and they are also decreasing in equity whenever loans or bonds are positive. This assumption says that default is more likely if leverage is higher, which makes both equity and debt securities riskier.¹⁰ Finally, the functions $\omega_{i,s}$ also capture firm-level differences in credit market frictions. For example, for a firm with relatively more tangible assets, the $\omega_{i,s}$ rise more slowly with security issuance. As a result, the firm can issue debt more cheaply.

Risk exposure of firm value. We assume that the risk exposure of firm value

$$\omega_i(\bar{a}_{t,i}) := \frac{\omega_{i,b}(\bar{a}_{t,i}) \bar{b}_{t,i} + \omega_{i,\ell}(\bar{a}_{t,i}) \bar{\ell}_{t,i} + \omega_{i,e}(\bar{a}_{t,i}) \bar{e}_{t,i}}{\bar{b}_{t,i} + \bar{\ell}_{t,i} + \bar{e}_{t,i}} \quad (30)$$

is also increasing in debt and decreasing in equity for all portfolios $\bar{a}_{t,i} = (\bar{b}_{t,i}, \bar{\ell}_{t,i}, \bar{e}_{t,i})$. This assumption means that leverage generates welfare costs and not simply redistribution among debtholders and equityholders. The reason is that risk exposure from holding all debt and equity of a firm is higher than risk exposure from holding directly the firm's capital, or equivalently a copy of the firm that is 100% funded with equity. In models with explicit default, this effect typically comes about because there are dead-weight cost of bankruptcy, a prominent ingredient in many models of capital structure choice.

Liquidity costs. We use proportional costs to capture market liquidity, including liquidity effects due to the central bank purchase program, as well as firms' access to the bond market. Private intermediaries that hold security s issued by firm i pay a proportional

¹⁰To understand the increase in risk exposures of debt and equity, suppose a firm i has a leverage ratio ϕ_t . A loan or bond issued by the firm is a portfolio with a share $0 < \rho(\phi_t) < 1$ invested in its risky capital and a share $1 - \rho(\phi_t)$ invested in a safe bond. The risk exposure of debt depends on the firm's leverage $\omega_{i,b}(\bar{a}_{t,i}) = \rho(\phi_t) \omega_i$, where ω_i is the risk exposure of firm i . Similarly, equity is a portfolio with a share $\psi(\phi_t)$ invested in risky capital and a share $1 - \psi(\phi_t)$ invested in a safe bond. The exposure of equity also depends on leverage, $\omega_{i,e}(\bar{a}_{t,i}) = \psi(\phi_t) \omega_i$, and is higher than that of debt, $\psi(\phi_t) > \rho(\phi_t)$. Debt plus equity add up to firm capital, $\phi_t \rho(\phi_t) + (1 - \phi_t) \psi(\phi_t) = 1$. In the absence of bankruptcy costs, higher leverage ϕ_t increases the risk exposures of both debt and equity: $\omega_{i,b}(\bar{a}_{t,i})$ and $\omega_{i,e}(\bar{a}_{t,i})$ are higher. However the overall risk exposure of the firm can still be the same, $\phi \omega_{i,b}(\bar{a}_{t,i}) + (1 - \phi) \omega_{i,e}(\bar{a}_{t,i}) = \omega_i$, as long as the overall risk exposure of the firm ω_i is unaffected by leverage. In this case, Modigliani-Miller holds. In a world with bankruptcy costs, the risk exposure of the firm ω_i increases with leverage, because units of capital are lost in bankruptcy, and Modigliani-Miller does not hold.

cost $\lambda_{i,s}$. The proportional cost of equity is always above that of loans

$$\lambda_{i,e} > \lambda_{i,\ell}.$$

The ranking of the proportional cost on bonds depends on the firms. Some firms have access to the bond market. For these firms, we assume $\lambda_{i,\ell} > \lambda_{i,b}$ to capture a liquidity advantage of bonds over loans that must be traded off against the extra risk exposure from bonds. Other firms do not have access to the bond market and have an infinite proportional cost of bonds.

We further allow for an additional liquidity benefit for the subset of firms $\tilde{I} \subset I$ that is eligible for central bank purchases. We assume that the proportional bond holding cost $\lambda_{i,b}$ is lower for all eligible firms – it is cheaper for private intermediaries to hold bonds also held by the central bank. The idea is that such bonds are more liquid when there is a prospective big player in the market. The effect depends only on eligibility – it does not matter how much of the security the central bank has purchased. Central bank action matters on the intensive margin only through absorbing risk.

Private intermediary balance sheet costs. Putting the pieces together, we now define the cost of holding a portfolio (a, d) that sums to total investment A as the sum of proportional and risk taking costs

$$h(a, d; \bar{a}, \tilde{I}) = \lambda_d d + \sum_{s \in S} \int_I \lambda_{i,s} a_{i,s} di + \frac{1}{2} \gamma \left(\omega(a, d; \bar{a})^\top \omega(a, d; \bar{a}) \right) A. \quad (31)$$

Here dependence on eligibility \tilde{I} works through lower liquidity cost on bonds $\lambda_{i,b}$. We note that h has all the properties introduced in Section 3.1. Indeed, both proportional and risk taking costs are homogenous of degree one and marginal holding costs depend on $\bar{a}_{t,i}$ only through $\omega_{i,s}(\bar{a}_{t,i})$ which is homogenous of degree zero. Moreover, derivatives of h with respect to $\bar{a}_{t,i,s}$ are homogenous of degree one in $a_{t,i}$ and depend on $a_{t,k}$ and d_t only through average exposure ω , on which they have a negligible effect.

Central bank balance sheet costs. The central bank holds only bonds and funds those holdings with central bank debt. In general, we allow the central bank costs to face different liquidity costs, sensitivity to risk as well as weighting of risk factors and define

$$\tilde{h}(a; \bar{a}, \tilde{I}) = \sum_{s \in S} \int_{\tilde{I}} \tilde{\lambda}_{i,s} a_{i,s} di + \frac{1}{2} \tilde{\gamma} \left(\omega(a, 0; \bar{a})^\top \omega(a, 0; \bar{a}) \right). \quad (32)$$

The calculation of central bank exposure works just like for private intermediaries, except

that we omit holdings of central bank debt. It makes sense to allow for different $\tilde{\lambda}$ s and $\tilde{\gamma}$ since these features will be shaped in part by political economy considerations. We emphasize, however, that differences in cost parameters are not essential for the central bank to add value. This is because it has the special ability to issue safe debt, which lowers overall cost even when holding costs are the same as for private intermediaries, that is, $\tilde{h}(a; \tilde{a}, \tilde{I}) = h(a, 0; \tilde{a}, \tilde{I})$. We return to this benchmark case below.

Comparison of the holding costs (31) and (32) clarifies the two levers of central bank intervention. First, changing eligibility determines liquidity cost at the firm level, an instrument that allows fine tuning the cost of capital for individual firms with access to the bond market. Second, buying risky securities funded by riskfree debt alters private intermediaries' risk exposure ω and hence their cost of risk taking. Trading securities is a blunt instrument that has spillover to all risky securities because its impact works through ω . With an integrated financial system, both sides of the central bank balance sheet matter: ω is reduced not only by central bank absorption of highly exposed securities, but also by its issuance of safe debt.

4.2 Asset pricing

Substituting for marginal holding cost in the intermediaries' optimality condition (14), we now write the equilibrium interest rate on central bank debt together with premia on security s issued by firm i as

$$\begin{aligned} M_{t+1}R_{t+1}^d &= 1 + \lambda_d - \left(\gamma'_t \omega_t^\top \omega_t - \frac{1}{2} \gamma_t \right), \\ M_{t+1} \left(R_{t+1,i,s} - R_{t+1}^d \right) &= \lambda_{i,s} - \lambda_d + \gamma'_t \omega_{t,i,s}^\top \omega_t, \end{aligned} \quad (33)$$

where we simplify notation by suppressing the argument of the risk cost function γ and its derivatives, for example we write $\gamma_t = \gamma(\omega_t^\top \omega_t)$ and $\gamma'_t = \gamma'(\omega_t^\top \omega_t)$, and also write $\omega_{t,i,s} = \omega_{i,s}(\tilde{a}_{ti})$ for the $\#F \times 1$ dimensional risk exposure of firm i 's security s . In general, premia reflect both liquidity and risk. Liquidity premia on private securities – described by the coefficients λ – are specific to the firm, even holding fixed the firm's risk exposure. In particular, the liquidity premium $\lambda_{i,b}$ on bonds may reflect central bank investment in the firm's bonds.

Risk premia on private securities depend on the marginal risk appetite of private intermediaries, captured by the first derivative γ'_t of the risk cost function, and their covariance with the risk exposure of the intermediaries $\omega_{t,i,s}^\top \omega_t$. They differ across securities only due to differences in the security risk exposures $\omega_{t,i,s}$: two securities issued

by different firms but with the same risk exposure pay the same risk premium.¹¹ For a given firm i , the ranking of exposures (29) implies that the equity risk premium is larger than risk premia on bonds and loans. At the same time, the interest rate on central bank debt R_{t+1}^d declines with total risk born by private intermediaries $\omega_t^\top \omega_t$, since the function γ is convex. The presence of liquidity and risk effects make it different from the discount rate of households R_{t+1}^f .

The risk premia expressions (33) illustrate how our model captures familiar stylized facts from the empirical asset pricing literature. The typical result is that one can find security characteristics that predict average excess returns in the cross section. In the case of classical risk factors, there exist factor-mimicking portfolios such that (i) the security characteristic is the factor beta, that is, the covariance between the security return and factor-mimicking returns, divided by the variance of the latter and (ii) the average excess return on the security is its factor beta times the average excess return on the factor-mimicking portfolio. While our deterministic framework does not make shocks explicit, it does capture the structure of premia that compensate for systematic factor risk. More generally, it allows for characteristics that are not classical risk factors to also shape premia.

To see the connection to standard asset pricing formulas, we define $\beta_{t,i,s}^f = \omega_{t,i,s}^f / \omega_t^f$ as the security's factor beta for factor f , given by the security's risk exposure divided by the risk ω_t^f in factor f . The exposures to factor f are the f th entry in the security's risk exposure $\omega_{t,i,s}$ and the intermediaries' risk exposure ω_t . We can now rewrite the premium on security s issued by firm i as

$$M_{t+1} \left(R_{t+1,i,s} - R_{t+1}^d \right) = \lambda_{i,s} - \lambda_d + \pi_{t,i,s}, \text{ where } \pi_{t,i,s} = \sum_{f \in F} \beta_{t,i,s}^f \pi_t^f \text{ and } \pi_t^f = \gamma_t' \left(\omega_t^f \right)^2 \quad (34)$$

where the risk premium $\pi_{t,i,s}$ is represented by the sum of security-specific factor betas multiplied by factor risk premia π_t^f that are independent of the security. If f is a classical risk factor, its factor mimicking portfolio works like a security with $\beta_{t,i}^f = 1$ and zero betas for the other factors, so that it exactly earns the factor premium π_t^f .

In our model, as in the data, the same set of factors is relevant for stock and bond returns. Investors – including the central bank – can make use of the factor structure

¹¹A key property of the cost function (31) is that it penalizes only systematic risk represented by factors, not exposure to individual firms. This assumption is sensible when intermediaries hold diversified portfolios. To illustrate, consider two firms i and j with the same exposure functions $\omega_{i,s}$ and the same leverage. Consider further a portfolio strategy that holds bonds issued by firm i and equity issued by firm j . The total cost of this portfolio is the same as that of a portfolio that holds both bonds and equity from firm i only – both strategies exhibit the same factor exposure and idiosyncratic firm specific risk is irrelevant.

to assemble portfolios that load more or less on the individual factors. For example, a greener portfolio would have negative weight on a climate risk factor. By holding bonds rather than stocks, investors can take on factor exposure but keep that exposure lower than that of the market portfolio. In contrast, holding equity of levered firms generally implies exposure higher than that of the market portfolio. Indeed, we have assumed in our exposure ranking (29) that equity is more exposed than debt, but also that total firm exposure $\omega_{t,i}$ in equation (30) – which determines firm i 's contribution to market portfolio – is a weighted average of individual security exposures.

Time variation in aggregate factor exposures ω_t create common variation in premia, that is, joint predictable variation in expected excess returns. For example, if the risk on an overall market factor or the exposure of intermediaries to that factor increases in recessions, then we should see higher premia on both stocks and bonds, as well as a lower riskfree interest rate. We note that the basic effect is present even if the function γ is linear and the marginal risk appetite γ' is constant, but is amplified by convexity of γ . That way, our model captures that additional risk can have very large effects when risk is already elevated, such as during crisis episodes.

Unconventional monetary policy matters for premia for two reasons. First, it can lower bond liquidity premia $\lambda_{i,b}$ for eligible firms. Second, it can lower private sector exposure and hence risk premia. To see this, we assume the central bank invests an amount $\delta_t = \bar{d}_t/K_t$ measured as a fraction of total capital and selects bond portfolio shares $\tilde{\kappa}_{i,t} = \bar{a}_{t,i,b}/K_t$ associated with risk exposures $\omega_{i,b}$ for the various firms. Now we substitute (30) into (28) to obtain

$$\omega_t = \int_I (\kappa_{t,i} \omega_i (\bar{a}_{t,i}) - \delta_t \tilde{\kappa}_{t,i} \omega_{i,b} (\bar{a}_{t,i})) di = \omega_t^* - \delta_t \tilde{\omega}_t. \quad (35)$$

Without a central bank purchase program ($\delta_t = 0$), private exposure is equal to the exposure of the market portfolio ω_t^* . More generally, the central bank risk taking lowers private exposure by substituting risky private securities with riskfree central bank debt. At the same time, policy affects exposure of the market portfolio itself by encouraging firm leverage. We discuss these effects in more detail below.

4.3 Leverage and the cost of capital

Consider next firm i 's choice of leverage. The firm knows its liquidity cost parameters as well as private intermediaries' exposure (35) and can therefore compute its cost of

capital as a function of the securities $\bar{a}_{t,i}$ it issues

$$M_{t+1}R_{t+1}^i = \sum_{s \in S} \frac{\bar{a}_{t,i,s}}{k_{t,i}} \lambda_{i,s} + \gamma'_t (\omega_{t,i} - \omega_t)^\top \omega_t + \frac{1}{2} \gamma_t.$$

Here the contribution of risk depends on the identity of the firm through the exposure $\omega_{t,i}$ of firm value as defined in (30). If there are no bankruptcy costs, that is, $\omega_{t,i}$ is independent of $\bar{a}_{t,i}$, then the risk premium required of the firm is also independent of capital structure. The firm would then only issue the securities with the lowest liquidity cost. If moreover there are no differences in liquidity costs, the Modigliani-Miller theorem holds.¹²

More generally, firm i chooses optimal leverage by trading off the liquidity benefit of bonds – which makes bonds a cheaper source of funding – against the increase in risk and hence costs of risk taking that leverage brings. At an interior solution, optimal loan and bond leverage ratios follows from

$$\begin{aligned} \lambda_{i,e} - \lambda_{i,b} &= \gamma'_t \left(\frac{\partial \omega_{t,i}}{\partial \bar{b}_{t,i}} (\bar{a}_{t,i}) - \frac{\partial \omega_{t,i}}{\partial \bar{e}_{t,i}} (\bar{a}_{t,i}) \right)^\top \omega_t k_{t,i}, \\ \lambda_{i,e} - \lambda_{i,l} &= \gamma'_t \left(\frac{\partial \omega_{t,i}}{\partial \bar{l}_{t,i}} (\bar{a}_{t,i}) - \frac{\partial \omega_{t,i}}{\partial \bar{e}_{t,i}} (\bar{a}_{t,i}) \right)^\top \omega_t k_{t,i}. \end{aligned} \quad (36)$$

Firms balance the liquidity benefit of both types of debt against the larger contribution debt makes to marginal cost. In particular, the right hand side in the first equation is positive because issuing equity *lowers* exposure.

The model predicts that leverage should decline with risk and lenders' cost of risk taking, as described by risk exposure ω and the curvature of the function γ . These ingredients increase the marginal cost of more leverage on the right hand side of (36). This feature is consistent with a large body of evidence on capital structure choice. For example, leverage is procyclical: in recessions when risk premia are large, leverage declines. Moreover, studies of financial crises emphasize "freezing" of loan markets.

4.4 Quantitative and quantitative policy effects

In this section, we characterize how policy affects asset prices and real activity through liquidity and risk channels. We focus on "small" purchase programs with δ close to zero,

¹²Even if the exposure of firm value ω_i does not depend on capital structure, we can still have security exposures $\omega_{i,s}$ depend on leverage. The model thus allows for differences in risk premia even if the Modigliani-Miller theorem holds.

the relevant case in light of recent central bank programs. For small programs, there is a close connection between welfare effects and observed security premia. Formally, we work out this connection below by taking linear approximations of risk premia around $\delta_t = 0$. We use Δ to indicate the change in a variable between an equilibrium without a purchase program and one with a program $(\delta_t, \tilde{\omega}_t)$. For example, we denote by $\Delta\lambda_{i,b}$ the difference in liquidity premia for bonds issued by firm i : it is negative for eligible firms and zero for ineligible firms.

Measuring policy impact. The key piece of evidence that guides our inference about magnitudes is the response of yield spreads to the introduction of a purchase program. In the model, yield spreads are equivalent to premia on bonds. When a purchase program $(\delta_t, \tilde{\omega}_t)$ is introduced, the model predicts that spreads change by

$$\Delta \left(M_{t+1} \left(R_{t+1,i,b} - R_{t+1}^d \right) \right) = \Delta\lambda_{i,b} - \delta_t \sum_{f \in F} \beta_{t,i,b}^f \pi_t^f \left(\tilde{\beta}_t^f + 2 \frac{\gamma_t'' \omega_t^\top \omega_t}{\gamma_t'} \frac{\tilde{\pi}_t}{\pi_t} - \frac{\partial \omega_{t,i,s}^f}{\partial \delta_t} \omega_{t,i,s}^{-1} \right), \quad (37)$$

where $\tilde{\beta}_t^f := \tilde{\omega}_t^f / \omega_t^f$ is the factor f beta of the central bank's portfolio and $\tilde{\pi}_t$ and π_t^* are risk premia on the central bank and market portfolios, respectively. In addition to the liquidity component $\Delta\lambda_{i,b}$, there is a risk premium component that scales with the size of the central bank purchase program δ_t . Since we are approximating around $\delta_t = 0$, both the market portfolio exposure $\omega_t = \omega_t^*$ and the premia π_t^f and π_t^* can be taken to be pre-intervention baselines.

Policy affects risk premia in three ways, described by the three terms in the bracket. First, central bank purchases of a factor directly lower factor risk premia introduced in (34), and more so if the factor beta $\tilde{\beta}_t^f$ of the central bank portfolio is larger. By loading more on a particular factor – that is, by purchasing more securities that are exposed to that factor – the central bank can thus assume factor risk, leave the private sector less exposed and thereby make bond issuance for exposed firms cheaper. With a linear cost function γ and fixed leverage, this would be the only effect. More generally, the purchase program also lowers overall risk, as described by $\omega_t^\top \omega_t$ and hence the risk appetite of the private sector. This effect is not factor-specific but lowers the entire risk premium; accordingly its strength is given by the ratio of central bank to market risk premia $\tilde{\pi}_t / \pi_t$, and the risk cost elasticity $\gamma_t'' \omega_t^\top \omega_t / \gamma_t'$ of the private sector. Finally, risk premia increase if firms increase leverage.

The empirical literature has documented two facts that help us measure key elements of (37). First, Mota and Papoutsis (2020) find that the decline in yields for eligible firms relative to ineligible firms in response to the introduction of the ECB purchase program

was about 15bp, controlling for other firm characteristics such as risk. Todorov (2020) finds that yields for eligible A-rated bond yields drop by 22 bp compared to ineligible bonds with the same rating, while lower rated BBB bond yields drop by 34bp. In terms of our model, this number corresponds to the average change in the purely firm-specific liquidity premium $\Delta\lambda_{i,b}$. Second, both Todorov (2020) and De Santis and Zaghini (2021) provide evidence that the policy also changed spreads on non-eligible firms, and more so if the firm is more risky. In particular, the estimated effect are about 20bp for investment grade firms and 130bp for low grade firms. Through the lens of the model, these numbers imply that some of the effect of policy must have come through risk premia.

To further identify the important components of changes in risk premia, we first note that the direct effect on factor risk premia is small when the program itself is small. The order of magnitude of betas on the central bank portfolio can be gauged from the fact that the central bank holds bonds, and that debt earns premia below the market portfolio which also includes stocks and loans. As an example, if the central bank holds bonds proportional to outstanding bonds, and the premium on a typical bond is, say, 2%, which is, say, one quarter of the premium on the market, then a program of size $\delta_t = 10\%$ generates a direct effect of $\delta_t 0.02/4$ or half a basis point. We further have that the effect of leverage on yields should be positive, whereas the measured effect is in fact negative.

We conclude that the bulk of the negative effect of policy on premia must come from the response of private sector risk appetite to total risk in the second term. If the elasticity of cost to risk $\gamma_t'' \omega_t^\top \omega_t / \gamma_t'$ is sufficiently large, even a purchase program that is relatively small (low δ) and relatively safe (low $\tilde{\beta}_t^f$ and hence low $\tilde{\pi}_t / \pi_t$) can have a sizeable effect on spreads. Since we do not have numbers for the effect through leverage, we cannot exactly identify the elasticity. However, we can use the spread response to provide a lower bound for the effect on the cost of capital. Moreover, since the effect is not factor specific, we can do so without precise measures of central bank factor betas.

Cross sectional effects of unconventional monetary policy. We are now ready to characterize how policy shapes the cross section of costs of capital. The difference between two individual firms' cost of capital is approximately

$$\Delta (M_{t+1} (R_{t,i} - R_{t,k})) \approx \frac{\bar{b}_{t,i}}{k_{t,i}} \Delta\lambda_i^b - \frac{\bar{b}_{t,k}}{k_{t,k}} \Delta\lambda_k^b - \delta_t \sum_f \left(\beta_{t,i}^f \pi_t^f - \beta_{t,k}^f \pi_t^f \right) \left(\tilde{\beta}_t^f + 2 \frac{\gamma_t'' \omega_t^\top \omega_t}{\gamma_t'} \frac{\tilde{\pi}_t}{\pi_t} \right). \quad (38)$$

where $\beta_{t,i}^f$ is the factor f beta for the firm value of firm i . Here the first two terms reflect differences in liquidity premia, and the last term differences in risk premia. The differences in liquidity premia depend on individual firms' bond leverage such as $\bar{b}_{t,i}/k_{t,i}$

but not on the size δ_t of the purchase program because eligibility matters on the extensive margin only. The last term is the change in relative risk premia which is closely related to the change in bond risk premia (37). A crucial difference is that changes in leverage do not have a first order effect on the cost of capital so derivatives with respect to capital structure are absent.

The central bank has three distinct policy levers that operate at different layers of aggregation, from securities to factors to aggregates. First, the central bank can target individual firms by making them eligible. It thereby improves bond market liquidity and lowers the liquidity premium of the firm. This effect works only through the bond market of the individual firm targeted, it affects neither other firms nor other securities of the same firm. As a result, the effect on the cost of capital is limited by the importance of bond financing for the firm – any yield effect must be multiplied by bond leverage (that is, the ratio of bonds to assets) in order to arrive at the cost of capital. At the typical bond leverage of 10%, the effect on the cost of capital is reduced by a factor of 10. At the same time, the liquidity effect does not scale with the size of the purchase program, and can therefore matter even for very small programs.

The second lever is *factor trading*. By assembling a portfolio exposed to factor f , the central bank makes it cheaper for firms exposed to f – as measured by firms' factor premium – to raise funds. In a multifactor world, factor trading thus also allows some targeting through central bank portfolio choice. In contrast to the firm-specific liquidity effect, however, factor trading is a blunter instrument: buying bonds exposed to a factor affects all firms exposed to that factor. Moreover, those firms are affected not only through the bond market, but also via loans and equity, which by (29) are even more exposed to the factor f and held by the same intermediaries. In terms of magnitudes, this means that the policy impact on the cost of capital is larger than the bond market impact – indeed, the risk premia terms in (38) scale with the firm premium $\beta_{t,i}^f \pi_t^f$ which is larger than the bond premium $\beta_{t,i,b}^f \pi_t^f$. Concretely, if the typical firm value premium is, say, three times larger than a bond premium, the policy impact is also multiplied by a factor of three.

The final lever is *total risk absorption*, described by the second risk premium term in (38). As we have seen, this is an important force in the data, since a high risk cost elasticity can generate a large policy impact from small purchase programs. It is also the bluntest instrument, as it affects all firms, regardless of the factor composition of their risk premia. Indeed, the cross sectional effect is proportional to overall firm risk premia. Importantly, this means that the effect is not neutral: absorption of risk lowers costs of capital of riskier firms relatively more. This is because replacing any risk on

private intermediaries' balance sheet with safe central bank debt makes it cheaper for those intermediaries to take on more risk.

Quantitatively, we can conclude that a central bank purchase program that has a significant effect on the cross section of yields also has a significant effect on the cross section of costs of capital. This point follows from comparing the risk premium terms in (37) and (38). In particular, the evidence shows that the difference in premia responses between high yield and investment grade bonds is about 20bp.

Aggregate effects in the log-normal model. We now turn to aggregate effects. Here we make use of the log-normal example introduced in Section 3.3, where we can work out changes in closed form. Unconventional monetary policy works through two channels. First, it alters the average cost of holding capital AC_t defined in (26) – with log utility, this gain is used by households to increase consumption. Moreover, it increases aggregate TFP by reducing the effect of financial frictions on the allocation of capital across firms. We note that a small purchase program does not alter the aggregate marginal cost of capital MC_t . Intuitively, a purchase program matters through its effect on private intermediary portfolio composition. The first dollar of purchases thus lowers average cost, but it does not make it cheaper to attract additional capital.

The change in consumption is

$$\Delta C_t = \frac{\bar{\alpha}\beta}{1 - \bar{\alpha}\beta} \delta_t \left(-\lambda_d - \left(\int_I \kappa_{t,i} (\tilde{\lambda}_{i,b} - \lambda_{i,b}) di \right) + \tilde{\pi}_t - \frac{1}{2} \tilde{\gamma}_t \right). \quad (39)$$

Again both liquidity and risk effects are relevant; because the effect on consumption works through average cost, both scale with the size of the purchase program δ_t . On the one hand, holding a dollar of capital in the central bank saves private holding cost for corporate bonds $\lambda_{i,b}$ but also adds central bank holding cost $\tilde{\lambda}_{i,b}$ as well as private holding cost of central bank debt λ_d . On the other hand, absorption of risk by the central bank saves private holding cost, as measured by the risk premium on the central bank portfolio, but requires central bank risk cost.

The change in consumption is positive provided that aggregate risk premium is sufficiently large. In general, a purchase program can thus be too small or too large – we return to this point in Section 5 when we characterize the optimal program. Here we are interested in the magnitude of the effects. Assuming that liquidity effects make a small contribution to overall premia on bonds, the message from (39) is that the effect on consumption is bounded above by $\delta_t \tilde{\pi}_t$, where $\tilde{\pi}_t$ is the risk premium on the central bank portfolio. A sizeable consumption effect from a small purchase program therefore

requires very high risk premia, as one might observe for example in a financial crisis. In normal times, with a premium on high grade bonds held by the central bank at, say, 2%, the consumption effect of a small program is limited to a few basis points.

Market neutrality and the tilt of the ECB portfolio. Our definition of a market neutral purchase program requires that the difference in costs of capital (38) is zero. The formula clarifies the two key forces of non-neutrality present in our model. First, liquidity effects imply that central bank purchases favor relative bond-levered firms. In particular, those effects do not matter for firms that do not issue bonds. They further favor firms with more tangible assets for which bond leverage is cheaper. Second, risk effects imply that central bank purchases favor relatively risky firms. Factor trading does allow some targeting of groups of firms by selectively reducing risk premia on some factors. At the same time, the quantitatively important effect of total risk absorption favors firms with higher overall risk premia.

The formula thus implies that market neutrality is generally impossible to achieve. The result follows from simple counting of equations and unknowns. Neutrality requires that (38) hold for *any* two firms i and k . Observed risk premia and leverage differ widely across firms, so in the typical data set we have many differences in risk premia, as well in factor risk premia. The central bank, however, can choose only eligibility as well as a finite number of exposures $\tilde{\omega}_t$ and the program scale δ_t . As a result, there generally does not exist a portfolio that makes the equation hold for all i and k , unless there is a "divine coincidence" whereby risk premia and liquidity premia cancel each other. We note that it does not matter for the argument that the central bank only hold bonds – it relies on the fact that the central bank affects a small number of factors and many factor exposures. Even a central bank that holds equity, say, cannot typically be neutral.

5 Optimal policy

5.1 Social planner problem

The social planner chooses an allocation $(C_t, k_t, \bar{a}_t, \bar{a}_t - \tilde{a}_t, \bar{d}_t, y_t, l_t, \tilde{I}_t)$ to maximize utility (1) subject to technology and resource constraints (2)-(8). Let μ_t denote the multiplier on the evolution of the environment (3), that is, the shadow cost of the environment, and let $\rho_{t,i}$ denote the multiplier on the financing constraint (10) for firm i , both normalized

by marginal utility $u'(C_t)$. The first order condition for capital of firm i in sector n is

$$u'(C_t)(1 + \rho_{t,i}) = \beta u'(C_{t+1}) \left(\theta_n \frac{Y_{t+1}}{Y_{t+1,n}} \left(\frac{Y_{t+1,n}}{y_{t+1,i}} \right)^{\frac{1}{\sigma}} - \varepsilon_{t+1,i} \mu_{t+1} \right) \alpha_n \frac{y_{t+1,i}}{k_{t,i}}. \quad (40)$$

Without climate externalities ($\varepsilon_{t+1,i} = 0$) and capital holding costs ($\rho_{t,i} = 0$), this is a standard intertemporal Euler equation: the marginal cost of investment today—the marginal utility of consumption—is equated to the discounted future benefit of consumption times the marginal product of capital.

Both holding costs and climate externalities alter this tradeoff. First, holding costs increase the marginal cost of investment on the left-hand side by a factor $1 + \rho_{t,i}$ which discourages capital accumulation. For example, for two firms with identical production functions, it is optimal to invest more in the firm that is less costly to hold. Second, the marginal benefit of investment on the right-hand side is reduced by the marginal cost of emissions. For two firms with the same technology and holding costs, it is optimal to invest more in the firm that generates less emissions.

The shadow cost of temperature μ_t reflects all future damages to production. From the planner's first order condition for environmental quality η_t , it evolves according to

$$\mu_t = \frac{\beta u'(C_{t+1})}{u'(C_t)} \mu_{t+1} + \int_I \frac{1 + \rho_{t,i}}{\alpha_n} \frac{z'_{t+1,i}(\eta_{t+1})}{z_{t+1,i}(\eta_{t+1})} k_{t,i} di. \quad (41)$$

Allowing the environment to degrade at date t has persistent effects by changing TFP that the planner discounts using the representative agent's marginal rate of substitution. The overall effect consists of two parts: lower TFP lowers the marginal product of capital and also lowers future emissions, a counteracting force.

Optimal choice of labor equates marginal products across firms. Let v_t denote the multiplier on the resource constraint for labor, or the shadow wage. With Cobb-Douglas technology, the planner would like ensure that factor expenditures satisfy

$$\frac{(1 + \rho_{t,i}) k_{t,i}}{v_t l_{t+1,i}} = \frac{1 - \alpha_n}{\alpha_n}. \quad (42)$$

Again marginal products take into account not only marginal utility from consumption, but also marginal environmental cost. For example, for two sectors that are identical in TFP, the optimal labor to capital ratio should be higher in the sector that produces fewer emissions. An important difference between labor and capital choice is that labor choice does not depend on holding costs, because labor choice in our model is not an

intertemporal decision.

What is the extra cost to the planner – in units of marginal utility – of raising capital $k_{t,i}$ in a frictional market? Since it is free to make securities eligible, the central bank should extend eligibility for all firms that are able to issue bonds. As for actual purchases, the first order condition with respect to total issuance $\bar{a}_{t,i,s}$ of security s is

$$\rho_{t,i} = \frac{\partial h}{\partial a_{t,i,s}} (\bar{a}_t - \tilde{a}_t, \bar{d}_t; \bar{a}_t, \tilde{I}_t) + \frac{\partial h}{\partial \bar{a}_{t,i,s}} (\bar{a}_t - \tilde{a}_t, \bar{d}_t; \bar{a}_t, \tilde{I}_t) + \frac{\partial \tilde{h}}{\partial \bar{a}_{t,i,s}} (\tilde{a}_t; \bar{a}_t, \tilde{I}_t) \quad (43)$$

Here the planner takes into account not only the marginal holding cost of private intermediaries – the first term – but also the effect of total issuance on private and central bank costs. For example, raising more debt to fund capital increases bankruptcy costs born by both private intermediaries and the central bank.

Optimal central bank holdings of securities, that is, the size and the composition of the purchase program now follow a static first order condition for central bank purchases $\tilde{a}_{t,i,s}$ that minimizes total holding cost $h + \tilde{h}$:

$$\frac{\partial h}{\partial a_{t,i,s}} (\bar{a}_t - \tilde{a}_t, \bar{d}_t; \bar{a}_t, \tilde{I}_t) = \frac{\partial \tilde{h}}{\partial \tilde{a}_{t,i,s}} (\tilde{a}_t; \bar{a}_t, \tilde{I}_t) + \frac{\partial h}{\partial d_t} (\bar{a}_t - \tilde{a}_t, \bar{d}_t; \bar{a}_t, \tilde{I}_t) \quad (44)$$

A purchase program can transform private holdings of any security into public holdings plus central bank debt: at the optimum, the marginal costs of both ways of holding the securities must be equal. Put differently, the planner wants to equate marginal private and central bank holding costs across securities, and takes into account the marginal cost of central bank debt.

We also note that optimality conditions (43)–(44) do not depend on the presence of climate externalities. The planner is concerned with raising capital $k_{t,i}$ per firm at minimal total cost. In this sense the purchase program targets only the financial frictions encoded in h and \tilde{h} , and does not target the externality. Importantly, the result does not mean that the size and composition of optimal purchases does not change with the climate once a carbon tax is in place. For example, if higher taxation of dirty firms reduces tangible assets and hence increases marginal costs of raising funds for the typical firm, then provision of safe central bank debt by the central bank becomes more valuable. The key is that this value is determined only through the effect of climate policy on financial frictions.

Implementing optimal policies with carbon taxes. Can we implement the social planner solution as an equilibrium? Consider an allocation that solves the planner problem

and a purchase program $(\tilde{a}_t, \tilde{I}_t)$ and implies a shadow cost of the environment μ_t . Define the carbon tax to reflect the social cost of carbon $\tau_t = \mu_t$. An equilibrium with this tax path and the optimal purchase program (44) in place generates the same allocation as the planner problem. Indeed, consider an equilibrium allocation given this policy, that is, a feasible allocation that together with security returns $R_{t,i,s}$, firm costs of capital R_t^i and wages w_t solves (19)-(21). It is equivalent to a solution of the planner problem that solves (40)-(43) with shadow prices $\mu_t = \tau_t$, $\rho_{t,i} = M_{t+1}R_{t+1}^i - 1$ and $v_t = w_t$.

Substituting out goods prices from firms' first order conditions (11), we have that equilibrium factor allocations are optimal as (40) and (42) hold for $\rho_{t,i} = M_{t+1}R_{t+1}^i - 1$. Moreover, if (44) holds, then (19) is equivalent to (43), so that equilibrium choices of capital structure are also optimal. Indeed, the last two terms in (43) can be rewritten as

$$\begin{aligned} & \frac{\partial h}{\partial \bar{a}_{t,i,s}} (\bar{a}_t - \tilde{a}_t, \bar{d}_t, \bar{a}_t, \tilde{I}_t) + \frac{\partial \tilde{h}}{\partial \bar{a}_{t,i,s}} (\tilde{a}_t, \bar{a}_t) \\ &= \sum_{\hat{s} \in S} \left\{ (\bar{a}_{t,i,s} - \tilde{a}_{t,i,s}) \frac{\partial^2 h}{\partial a_{t,i,\hat{s}} \partial \bar{a}_{t,i,s}} (\bar{a}_t - \tilde{a}_t, \bar{d}_t; \bar{a}_t, \tilde{I}_t) + \tilde{a}_{t,i,\hat{s}} \frac{\partial^2 \tilde{h}}{\partial a_{t,i,\hat{s}} \partial \bar{a}_{t,i,s}} (\tilde{a}_t; \bar{a}_t) \right\} \\ &= \sum_{\hat{s} \in S} \bar{a}_{t,i,s} \frac{\partial^2 h}{\partial a_{t,i,\hat{s}} \partial \bar{a}_{t,i,s}} (\bar{a}_t - \tilde{a}_t, \bar{d}_t; \bar{a}_t, \tilde{I}_t) - k_{t,i} \frac{\partial^2 h}{\partial d_t \partial \bar{a}_{s,\hat{s}}} (\bar{a}_t - \tilde{a}_t, \bar{d}_t; \bar{a}_t, \tilde{I}_t) \end{aligned} \quad (45)$$

which is the same as the second term in (19) since we have assumed that $\partial^2 h / \partial d \partial \bar{a}_{t,i,s}$ is negligible. Here the first equality uses that $\partial h / \partial \bar{a}_{t,i,s}$ is homogeneous of degree one in $\bar{a}_{t,i}$ and the second follows from (44).

Securities markets thus induce firms to issue optimal amounts of debt and equity – capital structure decisions are constrained efficient. This result holds not only with an optimal central bank purchase program in place, but also when there is no central bank. Without a central bank, only the first terms in each line of (45) are relevant. Homogeneity of $\partial h / \partial \bar{a}_{t,i,s}$ ensures that the marginal cost of issuance is linear in prices, so markets guide firms to the optimal capital structure. With a purchase program, the marginal cost to society also involves the contribution by the central bank. One might imagine that prices are not sufficient to convey the benefits of the program to firms. However, as long as the planner tunes the program by equating costs and benefits security by security according to (44), prices reflect the benefits. It is important here that individual firms are small so their capital structure does not affect intermediaries' marginal cost of central bank debt.

Optimal central bank portfolio. Using the cost function from our leading example (31), we can sharpen the characterization of optimal policy. We focus on equilibria with a rich cross section of bond risk premia, that is, $\omega_{t,i,s}$ differs across firms. To ensure existence

of such an equilibrium, we assume that the difference in bond holding costs between private intermediaries and the central bank does not depend on the firm i itself, and can thus be written as $\lambda_b - \tilde{\lambda}_b$. For all firms $i \in I_b$ with access to the bond market, the condition (44) for an interior solution becomes

$$\lambda_b - \lambda_d + \omega_{t,i,b}^\top \gamma'_t \omega_t = \tilde{\lambda}_b + \omega_{t,i,b}^\top \tilde{\gamma}'_t \tilde{\omega}_t - \tilde{\gamma}'_t \tilde{\omega}_t^\top \tilde{\omega}_t + \frac{1}{2} \tilde{\gamma}_t, \quad (46)$$

The left hand side is the premium on firm i bonds over the central bank debt, that is, the marginal cost to private intermediaries of replacing firm i bonds on their portfolio with central bank debt. The planner chooses the central bank portfolio so as to equate this premium to the central bank holding cost.

The exact composition of the optimal central bank portfolio is indeterminate – the central bank affects risk premia by changing exposures, not individual security holdings. Formally, policy enters (46) only through the $\#F$ -dimensional exposure vector $\tilde{\omega}_t$. In order for the condition to hold for all firms i that access the bond market in an equilibrium with heterogeneous exposure $\omega_{t,i,b}$, we must have equal private and central bank exposures $\omega_t = \tilde{\omega}_t$. Using (35), the central bank should therefore choose exposure proportional to exposure of the market portfolio ω_t^*

$$\tilde{\omega}_t = \frac{\gamma'_t}{\tilde{\gamma}'_t + \delta_t} \omega_t^*. \quad (47)$$

While this is not a closed form solution for central bank exposure since the fraction depends on the endogenous objects $\tilde{\omega}_t$ and δ_t , it nevertheless has content because the same fraction multiplies all exposures.

We draw four lessons from the optimal portfolio condition (47). First, the central bank should hold *factors* – not securities – in the same proportions as the market portfolio. For example, if there are two factors, say a factor that moves with the business cycle, but also a climate risk factor, then optimal policy simply mimics relative exposure of the overall economy to those factors – the central bank should not overweigh either one. In a world with a small number of risk factors, this is how the central bank respects the general principle derived in (44) that public and private marginal costs should be equated.

Second, the central bank portfolio should be safer – in the sense of lower exposure – the higher its relative cost of risk taking, and the larger its footprint in the economy captured by δ_t . In particular, if the central bank has no advantage in lending to private firms, so its cost of risk taking is the same as that of the private sector ($\tilde{\gamma} = \gamma$), then the optimality condition $\tilde{\omega}_t = \omega_t$ implies equal marginal costs of extra risk $\tilde{\gamma}' = \gamma'$ at the

optimum, and we have that the central bank portfolio should be *strictly safer* than the market portfolio for any $\delta_t > 0$. Intuitively, in this case the central bank's only advantage is safe borrowing. Since its debt makes private intermediaries safer, holding the market portfolio would make the central bank more risky than private intermediaries, which is not efficient as it does not equate marginal cost. Moreover, as the central bank portfolio grows, its risk grows linearly, whereas private sector risk, which is diluted by safe central bank debt, grows more slowly. Equating marginal costs calls for making the central bank portfolio even safer.

The result thus provides a justification for the widespread practice of central banks purchasing bonds rather than equity – from the ranking of security level exposures (29), holding bonds is a convenient way to assemble a diversified portfolio that is less exposed to risk than holding the market outright. We note that this advantage of holding bonds is present even when the financial system is fully integrated, that is, the same intermediaries hold both bonds and equity. Our model further says that the optimal portfolio should become relative more risky when the relative marginal cost of risk taking for the private sector increases. It thus makes sense to tilt towards a riskier portfolio in times of financial sector distress (for example, by purchasing lower grade bonds. At the same time, any large expansion of the central bank balance sheet calls for a more conservative investment approach.

The third takeaway from (47) is that the optimal purchase program is generally not market neutral, but instead favors riskier firms by lowering their cost of capital relatively more. Indeed, for positive δ_t , an optimal purchase program lowers overall risk in private intermediaries $\omega_t^\top \omega_t$ – it is not worthwhile otherwise since it increases central bank risk costs and central bank debt holding cost. Holding the market portfolio of factor exposure thus means that exposure to all risk factors is lower. But then firms with higher exposures to any risk factor experience a stronger reduction in the cost of capital.

Finally, we note that by lowering private intermediary exposure and hence risk premia, the central bank incentivizes risky leverage by firms. This is part of an optimal financial structure: because the central bank has the power to borrow without increasing leverage cost, it dilutes risk on bank balance sheets and hence allows for more risk created via private leverage. Pricing of risky claims ensures that this increase in risk does not overwhelm the overall decline in exposure – if this were true, then risk premia would adjust so as to drive firms to lower leverage.

What is the optimal scale of the central bank balance sheet? Increasing the purchase program is beneficial since it creates riskfree debt, but it is costly since that debt needs to be held by the private sector – overall balance sheets in the economy lengthen. Sub-

stituting (47) into (49) we determine total risk on the central bank's books $\tilde{\omega}'\tilde{\omega}_t$ as

$$\lambda_d + \tilde{\lambda}_b - \lambda_b = \gamma'_t - \frac{1}{2}\gamma_t.$$

Holding costs of central bank debt are linear, whereas there is a diminishing marginal benefit of creating safe debt. As a result, there is an interior solution for government risk taking.

We then find δ_t , the size of the purchase program as a share of the market portfolio by substituting back into (47):

$$\tilde{\omega}'\tilde{\omega} = \frac{\gamma'_t}{\tilde{\gamma}'_t + \delta_t} \omega_t^{*\top} \omega_t^*$$

The central bank should have a larger footprint if total risk in the economy $\omega_t^{*\top} \omega_t^*$ is larger, and when it is relatively better at lending to risky firm ($\tilde{\gamma}'$ small relative to $\tilde{\gamma}$). We obtain the intuitive result that the central bank should expand in size in time of private intermediary distress, even if it is not better than the central bank at lending to the private sector.

References

- Baker, Malcolm, Daniel Bergstresser, George Serafeim, and Jeffrey Wurgler**, “Financing the Response to Climate Change: The Pricing and Ownership of U.S. Green Bonds,” Working Paper 25194, National Bureau of Economic Research October 2018.
- Baldauf, Markus, Lorenzo Garlappi, and Constantine Yannelis**, “Does Climate Change Affect Real Estate Prices? Only If You Believe In It,” *Review of Financial Studies*, 2020, 33 (3), 1256–1295.
- Bolton, Patrick and Marcin T. Kacperczyk**, “Do investors care about carbon risk?,” *Forthcoming Journal of Financial Economics*, 2021a.
- and —, “Global Pricing of Carbon-Transition Risk,” 2021b. Working Paper, Imperial College.
- Boyarchenko, Nina, Anna Kovner, and Or Shachar**, “It’s What You Say and What You Buy: A Holistic Evaluation of the Corporate Credit Facilities,” 2020. CEPR Discussion Paper No. DP15432.
- Choi, Darwin, Zhenyu Gao, and Wenxi Jiang**, “Attention to Global Warming,” *Review of Financial Studies*, 2020, 33 (3), 1112–1143.
- Cúrdia, Vasco and Michael Woodford**, “The Central-Bank Balance Sheet as an Instrument of Monetary Policy,” *Journal of Monetary Economics*, 2011, 58 (1), 54–79.
- D’Amico, Stefania and Thomas King**, “Flow and Stock Effects of Large-Scale Treasury Purchases: Evidence on the Importance of Local Supply,” *Journal of Financial Economics*, 2013, 108 (2), 425–448.
- David, Joel and David Zeke**, “Risk-Taking, Capital Allocation and Optimal Monetary Policy,” 2021. Working Paper, USC.
- De Santis, Roberto A. and Andrea Zaghini**, “Unconventional monetary policy and corporate bond issuance,” *European Economic Review*, 2021, 135, 103727.
- Engle, Robert F., Stefano Giglio, Bryan Kelly, Heebum Lee, and Johannes Stroebel**, “Hedging Climate Change News,” *Review of Financial Studies*, 2020, 33 (3), 1184–1216.
- Ferrari, Alessandro and Valerio Nispi Landi**, “Whatever it Takes to Save the Planet? Central Banks and Unconventional Green Policy,” 2021. Working Paper, Banca d’Italia.
- Gagnon, Joseph, Matthew Raskin, Julie Remache, and Brian Sack**, “The financial market effects of the Federal Reserve’s large-scale asset purchases,” *International Journal of Central Banking*, 2011, 7 (1), 3–43.

- Gertler, Mark and Peter Karadi**, “A Model of Unconventional Monetary Policy,” *Journal of Monetary Economics*, 2011, 58 (1), 17–34.
- Giglio, Stefano, Bryan Kelly, and Johannes Stroebel**, “Climate Finance,” *Forthcoming Annual Review of Financial Economics*, 2021.
- Hambel, Christoph, Holger Kraft, and Rick van der Ploeg**, “Asset Diversification versus Climate Action,” 2021. Working Paper, University of Oxford.
- Hassler, John, Per Krusell, and Tony Smith**, “Environmental Macroeconomics,” *Handbook of Macroeconomics*, 2016, 2, 1893–2008.
- Hong, Harrison and Marcin T. Kacperczyk**, “The price of sin: The effects of social norms on markets,” *Journal of Financial Economics*, 2009, 93 (1), 15–36.
- Koijen, Ralph, François Koulischer, Benoît Nguyen, and Motohiro Yogo**, “Inspecting the Mechanism of Quantitative Easing in the Euro Area,” *Journal of Financial Economics*, 2021, 140 (1).
- Krishnamurthy, Arvind and Annette Vissing-Jorgensen**, “The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy,” *Brookings Papers on Economic Activity*, 2011, p. 215–265.
- , **Stefan Nagel, and Annette Vissing-Jorgensen**, “ECB Policies Involving Government Bond Purchases: Impact and Channels,” *Review of Finance*, 2018, 22 (1), 1–44.
- Kurtzman, Robert and David Zeke**, “Misallocation costs of digging deeper into the central bank toolkit,” *Review of Economic Dynamics*, 2020, 38, 94–126.
- Matikainen, Sini, Emanuele Campiglio, and Dimitri Zenghelis**, “The climate impact of quantitative easing,” 2017. Policy Paper, Grantham Research Institute on Climate Change and the Environment.
- Mota, Lira and Melina Papoutsis**, “Understanding the Effects of Unconventional Monetary Policy on Corporate Bond Market in the Euro Area,” 2020. Working Paper, ECB.
- Ottonello, Pablo and Thomas Winberry**, “Financial Heterogeneity and the Investment Channel of Monetary Policy,” *Econometrica*, 2020, 88 (6), 2473–2502.
- Piazzesi, Monika, Ciaran Rogers, and Martin Schneider**, “Money and Banking in a New Keynesian Model,” 2020. Working Paper, Stanford.
- Schoenmaker, Dirk**, “Greening monetary policy,” *Climate Policy*, 2021, 21, 581–592.
- Todorov, Karamfil**, “Quantify the quantitative easing: Impact on bonds and corporate bond issuance,” *Journal of Financial Economics*, 2020, 135, 340–358.

Zerbib, Olivier David, "The effect of pro-environmental preferences on bond prices: Evidence from green bonds," *Journal of Banking and Finance*, 2019, 98, 39–60.