

Regulating Bank Portfolio Choice Under Asymmetric Information

Chris Anderson*

September 1, 2023

Abstract

Regulating bank risk-taking is challenging since banks know more than regulators about the risks of their portfolios and can adjust their portfolios to game regulations. To address this problem, I build a tractable model that incorporates this information asymmetry. I derive the optimal calibration of linear risk-sensitive taxes, which should *not* generally be set more conservatively to address asymmetric information. I further show the efficacy of three novel regulatory tools: Not disclosing taxes to banks until after portfolio selection, non-linear taxes that respond to information contained in banks' portfolio choice, and state-dependent taxes on banks' realized profits.

*Federal Reserve Board. Email: christopher.s.anderson@frb.gov. I thank Andrew Hawley for excellent research assistance. For helpful comments, I thank Chris Clayton, Sharjil Haque, David McArthur, Marco Migueis, Ben Ranish, Doriana Ruffino, Nikola Tarashev, Cindy Vojtech, and Ke Wang. The views in this paper are solely the author's and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System.

1 Introduction

Modern banking regulation generally aims to be risk sensitive. For example, under the current Basel III framework, the primary determinants of capital requirements are risk-based measures that aim to assign higher requirements to riskier portfolios. Risk weights are a major input into these measures, which reflect regulators' best guess of the risks of different exposures within banks' portfolios.

However, regulators typically know less about risk than banks, which can expose these frameworks to potential gaming. When regulators underestimate an asset's risk, banks can take advantage by overweighting it in their portfolios. Therefore, a common heuristic among regulators is to be conservative on average, under the theory that such conservatism will counterbalance this gaming.¹ However, it is unclear under which circumstances this heuristic should work or whether there are better ways of addressing this problem. Furthermore, there has been surprisingly little theoretical work on this topic, in spite of its practical importance. In this paper, I build a tractable model to address the question of how to regulate banks' portfolio choice in the presence of asymmetric information, which takes into account that banks will game the regulations ex-post.

In this model, a bank selects its portfolio among a wide range of assets, which I interpret as reflecting investments in different granular sub-sectors (e.g., an asset might represent all loans to small biotechnology companies), subject to regulatory constraints. While these constraints can take many forms, for tractability, I focus on a regulator setting risk-sensitive taxes. These taxes could be interpreted as literal taxes or as shadow costs from quantity-based regulation, such as capital requirements. Each asset has a different level of risk, measured as its beta with respect to a single systematic risk factor, and profitability. Information asymmetry exists because the bank perfectly observes each asset's risk and profitability, while the regulator only receives noisy signals of each. The bank and regulator share the objective of increasing profitability and decreasing risk, but bank risk-taking imposes a social externality that the bank does not internalize. Therefore, the regulator wants to reduce the riskiness of the bank's portfolio, but also recognizes the social costs of the bank forgoing profitable investments.

The major contribution of this paper is to analyze how three novel regulatory tools can address this problem of asymmetric information: Not disclosing taxes to banks until after portfolio selection, non-linear taxes that respond to information contained in banks' portfolio choice, and state-dependent taxes on banks' realized profits. Each or all of these tools could be incorporated into the regulatory framework.

Before analyzing the novel tools, I begin by considering a more familiar tool: An

¹This observation comes from my personal experiences working in banking supervision.

asset-specific linear tax, which is similar to the linear risk weights used in most current capital regulation. Assuming a linear social cost to risk exposure, the regulator’s optimal policy is to set the tax exactly equal to the expected risk given their signals. Contrary to common intuition, the bank’s ex-post gaming of regulations does *not* mean the regulator should set taxes conservatively (i.e., higher than expected risk). The reason is that, in this environment, the costs from setting overly high taxes (i.e., dissuading socially optimal investments) are similar to the costs from setting overly low taxes (i.e., allowing banks to take more than the socially-optimal level of risk). Therefore, bank strategic behavior is not a rationale for conservatism by itself: Non-linearity of social costs or underweighting the benefits of profitable investments are necessary.² I then proceed to consider how the three novel regulatory tools can address this asymmetric information problem.

The first tool is strategic non-disclosure. Specifically, I consider a case in which the regulator sets a linear tax on exposure to each asset, but does not reveal the magnitude of this tax to the bank until after it has selected its portfolio. Thus, the bank must then select its portfolio based on its best guess of these taxes. However, if the bank has no information about the regulator’s signals, the bank’s best guess is that the taxes will be correctly calibrated and therefore the regulator can achieve the first-best outcome. However, even if the bank has imperfect information about the likely size of the tax, non-disclosure can still reduce the bank’s ability to game regulation. This type of approach is in the spirit of stress tests, in which banks do not fully know the regulator’s model, but which could be applied more broadly to other parts of the regulatory framework.

The second tool is a non-linear tax on risk exposures. The bank’s decision to invest more in an asset reveals information about the asset’s likely riskiness, which the regulator can incorporate by using a non-linear tax. Specifically, the regulator’s optimal policy is to set the marginal tax so that it equals the regulator’s estimate of risk, conditioned on level of the bank’s investment in that asset (along with the regulator’s other information). However, counterintuitively, the marginal tax may not necessarily be increasing in the bank’s investment. For example, the bank might be investing more in an asset because it has *lower* risk, which would encourage the regulator to set *decreasing* marginal taxes. In general, whether the optimal marginal tax is increasing or decreasing depends on the correlation between each asset’s profitability and riskiness, which determines whether the bank increases or decreases investment in riskier assets on average.

The third tool is a state-dependent tax on ex-post profits. The regulator sets a tax on profits that has a higher rate during “good times” (i.e., when the stochastic discount factor (SDF) of banks’ investors is lower). This type of tax on profits reduces the expected

²While this paper takes a Bayesian perspective, accounting for model uncertainty does not necessarily justify conservatism. Worst-case scenarios could include a regulator overestimating risk, setting an overly high tax, and preventing socially-efficient investments.

return that banks receive for taking a given amount of risk, which worsens the bank’s after-tax risk-return trade-off and effectively makes the bank more risk averse. If perfectly calibrated, this type of tax can perfectly align the bank’s incentives with the regulator’s. Since banks’ profits are typically higher during booms (which reflect “good times”), a progressive tax on profits could approximate the outcome of a perfectly-calibrated tax. A flat tax on profits is not sufficient: Even though a flat tax reduces banks’ after-tax profits, it also reduces their after-tax portfolio risk in a similar way and thus does not change the risk-return trade-off.

While this paper is not the first to consider asymmetric information in the context of banking regulation (for example, see Giammarino et al. (1993), Chan et al. (1992), Wu and Zhao (2016), and Perotti and Suarez (2018)), it differs in that it (1) proposes a different set of regulatory tools, (2) focuses specifically on the problem of portfolio choice given a wide set of assets, and (3) allows for the regulator to have some information, even if imperfect.

This paper is structured as follows: Section 2 lays out the model. Section 3 covers regulation based on taxing expected risk exposure. Section 4 covers regulation based on non-disclosed taxes, which are not revealed to the bank until after it has selected its portfolio. Section 5 covers regulation based on taxing banks’ profits to reduce their effective risk aversion. Section 6 describes the policy applications. Section 7 concludes.

1.1 Related literature

The most closely related papers explicitly model information asymmetries between banks and regulators. However, they generally do not consider regulators with partial information, implications for banks’ portfolio choice, or the performance of certain regulatory tools, such as non-linear risk-sensitive taxes or taxes on ex-post profits.

Perotti and Suarez (2018) is of the most closely related papers. The paper considers optimal regulation when bank illiquidity imposes an externality, similar to how bank risk-taking imposes an externality in this current paper. They consider the use of both Pigovian taxes as well as quantity-based regulation, taking into account regulators’ uncertainty with respect to banks’ investment opportunities and gambling incentives, in the spirit of Weitzman (1974). They abstract away from considering the liquidity of individual assets or liabilities. There is a clear similarity in that both papers consider Pigovian taxation under some form of uncertainty. However, there are major differences between their paper and the current one. The most obvious is their focus on liquidity risk, whereas this paper more naturally focuses on risks arising from losses in asset value. The current paper focuses on bank’s portfolios, which speaks to questions of optimal asset-specific regulations that their paper abstracts away from. The current paper also

considers different regulatory instruments, such as information nondisclosure and taxes on ex-post profits.

Giammarino et al. (1993) address bank regulation under asymmetric information, but with a different focus than this paper. In their model, the regulator has perfect knowledge of banks' loan quality, but does not know how much is due to investment opportunities versus the banker's effort. A key issue in their model is ensuring that the regulator properly incentivizes the banker to exert effort to improve the quality of the bank's loan portfolio. Their focus is not on portfolio choice. In contrast, in the current paper, regulators do not have perfect knowledge about loan quality or assets' riskiness. The focus is on using regulation to ensure the bank selects a less risky portfolio.

Chan et al. (1992) is another paper addressing asymmetry of information between banks and a regulator that provides deposit insurance. In their environment, the regulator has no direct measure of the riskiness of banks' portfolios. They address the problem of how a regulator without direct knowledge of the riskiness of banks' portfolios can offer incentive-compatible choices of deposit insurance premiums and capital requirements so that banks reveal the riskiness of their loans. As in the case of the previous paper, their focus is not on banks' portfolio choice and does not address the question of how a regulator with limited, but not perfect, information should act.

Wu and Zhao (2016) consider the benefits of adding a leverage ratio requirement on top of risk-based capital requirements in the presence of asymmetric information between banks and regulators. They show that including a leverage ratio reduces banks' incentives to misreport their level of risk. As with the other papers, their focus is not on portfolio choice.

Other papers consider information asymmetries in the context of banking, but with a different focus. For example, Goldstein and Leitner (2018) relates to the optimal disclosure of stress test results, but is more concerned with how regulators can affect the asymmetry of information between banks and the market.

This paper also relates to the theoretical literature on optimal capital regulation for banks, particularly the impacts on portfolio choice. Kim and Santomero (1988) derive the optimal risk weights for a regulator with full information on asset riskiness with the goal of limiting failure probabilities below a certain level. Rochet (1992) similarly derives optimal regulatory policy, both in terms of asset risk weights for capital requirements and pricing of deposit insurance, for a regulator with full information. More recently, Glasserman and Kang (2014) considers the problem of a regulator choosing optimal risk weights, including a case in which the regulator does not know the mean return of each asset. The main difference is that this paper focuses on the problem of a regulator with limited information about the riskiness of individual assets, whereas the regulator in their

paper has full knowledge of each assets' riskiness.

This paper also relates to empirical literature on risk in banking. Meiselman et al. (2018) empirically demonstrate that banks with high profits were more likely to crash during the financial crisis. Morgan and Ashcraft (2003) propose measuring the risk of a loan portfolio using interest rates and show that banks with higher loan spreads have more non-performing loans over the next year and a higher likelihood of a CAMELS rating decline over the next two years.

The results on taxing banks' profits relate to a literature on how tax policy affects bank behavior. John et al. (1991) is the most relevant, in that they propose setting a progressive tax to mitigate banks' incentives to tax excessive risks. Shackelford et al. (2010) discuss various ways in which taxation may be used to address externalities in the financial sector. They note that information asymmetries make Pigouvian taxation to address these externalities difficult and broadly discuss how financial transactions taxes, taxes on bonuses, and levies on banks may partially mitigate those externalities. Empirically, Celerier et al. (2019) empirically demonstrates the impact of Belgium's adoption of an equity subsidy, which allowed banks to deduct an estimate of the cost of equity from their taxes, on the composition of their portfolio. They show that the equity subsidy led banks to shift toward holding more loans rather than government bonds. And while not directly related to banks, this paper relates to a longstanding literature on the impact of taxation on risk-taking in general, such as in Domar and Musgrave (1944) and Stiglitz (1969).

2 The general framework

This is a single-period model containing a bank and a regulator. There are many assets whose riskiness (which is reflected by each asset's beta with respect to a common risk factor) and expected return are drawn from known prior distributions. Although the framework is general, I think of a single asset in the model as mapping to an investment in some granular sub-category, such as "loans to small biotechnology companies". The bank has perfect information on the realized betas and expected returns, whereas the regulator only receives noisy signals. The regulator can use these signals to establish regulatory constraints, which in general can take many forms, and then the bank selects its investments subject to those constraints.

I will typically assume that the bank's starting equity is exogenous, although Section B in the appendix considers the case in which the bank endogenously selects its desired level of equity. Without loss of generality, normalize the bank's starting equity to one

unit.³ There is a continuum of risky assets indexed by $i \in [0, 1]$. The bank selects q_i , its quantity of investment in each asset. The overall expected excess return to investing in asset i is $a_i q_i - \frac{c}{2} q_i^2$. I assume that $c > 0$, so that there are diminishing returns to investment. q_i may be negative, which can be interpreted as the bank taking a short position.

Each asset is exposed to a single systematic factor F with an asset-specific β_i loading.⁴ Without loss of generality, I set $E[F] = 0$ and $Var(F) = 1$. Thus, each asset i has a payoff X_i of the form⁵

$$X_i = a_i - \frac{c}{2} q_i + \beta_i F. \quad (1)$$

The bank's objective is to maximize its market value, which is given by

$$E_B[M(\Pi - T)] \quad (2)$$

where $\Pi = \int_0^1 q_i X_i di$ is the bank's pre-tax profit, T represents potential taxes paid to the regulator, M is the investors' SDF, and E_B denotes an expectation with respect to the bank's information set. I assume that the investors' SDF is⁶

$$M = 1 - \gamma F, \quad (3)$$

which implies that the bank maximizes

$$E_B[(\Pi - T)] - \gamma Cov_B((\Pi - T), F) \quad (4)$$

$$= \int_0^1 (a q_i - \frac{c}{2} q_i^2 - \gamma \beta_i q_i) di - E_B[T] - \gamma Cov_B(T, F). \quad (5)$$

Here, risk is measured based on the bank's portfolio beta rather than the variance. A negative beta portfolio is considered less risky than a zero beta portfolio because it

³By normalizing the bank's equity to one, all quantities can be interpreted as relative to the bank's total capital. Large quantities of investments in risky assets correspond to high leverage since banks must issue debt (whether as deposits or in another form) to fund those positions.

⁴The portfolio-invariant risk weights in Basel II and III can be justified by the assumption of a single systematic factor, as discussed in Gordy (2003). However, the assumption of a single factor is not as important in this paper's framework. As will be explained in more detail shortly, risk in the context of this paper is based on covariance with the investors' stochastic discount factor (SDF) rather than the variance of the bank's portfolio. Under the SDF-based measure of risk, each asset's risk contribution does not depend on the rest of the portfolio even if there are multiple systematic factors.

⁵In this setting, there is no idiosyncratic risk. Since there is a continuum of assets, the idiosyncratic risk would be diversified away and would therefore not affect the results of this model.

⁶While the SDF can technically become negative for large values of F , this issue is not a concern for my application. This SDF results in a price of γ per unit of exposure to the systematic risk factor, which leads to a more tractable model.

provides the bank's investors with insurance by paying out during "bad times" when the marginal value of wealth is high. This linear price of risk exposure matches the approach taken by Froot and Stein (1998).

The bank's risk-taking imposes externalities $\eta > 0$ per unit of portfolio beta $\beta_p = \int_0^1 q_i \beta_i di$. The regulator's objective is to maximize social welfare, which is

$$E_R \left[\left(M\Pi - \eta \int_0^1 q_i \beta_i di \right) \left(\right. \right. \quad (6)$$

$$\left. \left. = E_R \left[\int_0^1 \left(a q_i - \frac{c}{2} q_i^2 - (\gamma + \eta) \beta_i q_i \right) di \right] \right] \quad (7)$$

where E_R denotes an expectation with respect to the regulator's information set, which will be specified in greater detail shortly.⁷ The regulator's objective is identical to the bank's, except with higher effective risk aversion. The regulator wants the bank to be profitable, but is willing to accept a lower amount of risk for each unit of profit.

The timing and information structure of the situation is as follows: First, the β_i loadings and a_i measures of investment profitability are drawn from a prior distribution known to both the bank and the regulator. Then, the regulator receives noisy signals $\hat{\beta}_i$ and $\hat{a}_i$ for each asset. These signals reflect the regulator's imperfect information on how risky each asset is. Next, the regulator uses this information to determine regulatory constraints on the bank. There are several cases of regulatory constraints, but one example is the regulator selecting a tax k_i that the bank must pay for each unit of exposure to asset i . This cost could be interpreted as either a direct cost, such as a premium for deposit insurance, or an indirect cost, such as the cost of higher capital requirements for a position. These payments represent a private cost to the bank, but do not affect social surplus. Finally, the bank optimally selects its portfolio knowing the true β_i and a_i parameters, reflecting superior knowledge of assets' riskiness, subject to paying any costs or following any constraints imposed by the regulator.

In determining optimal regulation, the regulator faces a trade off. Since the regulator has a higher effective risk aversion than the bank, the regulator wants to reduce riskiness. But regulations meant to reduce riskiness might lead to investment efficiencies, which reduce social welfare.

⁷In principle, the bank's investments might also impose positive externalities that the regulator takes into account. For simplicity, I'll assume that there are no such externalities, which means that the bank is capturing all social surplus.

The processes for each asset's true profitability and beta are

$$a_i = \bar{a} + u_i^a \quad (8)$$

$$\beta_i = \bar{\beta} + u_i^b, \quad (9)$$

where u_i^a and u_i^b are mean-zero jointly normal random variables with variances σ_{ua}^2 and σ_{ub}^2 . u_i^a and u_j^b have a correlation of ρ if $i = j$ and are independent if $i \neq j$. Therefore, the draws of a_i and β_i are independent across assets.

The regulator receives a pair of signals

$$\hat{a}_i = a_i + e_i^a \quad (10)$$

$$\hat{\beta}_i = \beta_i + e_i^b, \quad (11)$$

where e_i^a and e_i^b are mean-zero jointly normal random variables with variances σ_{ea}^2 and σ_{eb}^2 respectively. e_i^a and e_i^b are independent of each other.

2.1 No regulation

To build intuition for the cases involving regulation that follow, I begin by considering the bank's problem when there is no regulation.

In this case, the bank solves

$$\max_{\{q_i\}} \int_0^1 (aq_i - \frac{c}{2}q_i^2 - \gamma\beta_i q_i) di, \quad (12)$$

which quickly leads to the solution

$$q_i = \frac{1}{c}(a_i - \gamma\beta_i). \quad (13)$$

The bank's investment in an asset increases in a_i and decreases in β_i . While seemingly straightforward, it is important to remember that, all else equal, banks do not like to take risk. They will only take on additional risk if they are sufficiently compensated for it. Common intuition is that banks will take advantage of weaknesses in regulation to increase their risk, but this intuition is only true if that risk is sufficiently compensated.

The $a_i - \gamma\beta_i$ term will be important for future results. I call this term the "initial risk-adjusted return". In this model, the bank's marginal return from investing in asset i is $a_i - cq_i$, so a_i represents the initial marginal return, starting from when q_i is zero. Since the bank dislikes beta exposure with a cost of γ , the $-\gamma\beta_i$ term represents a risk adjustment from the bank's point of view.

Whether banks invest more in high-beta assets depends on the relationship between each asset's initial risk-adjusted return and its beta. Specifically, it depends on whether

$$Cov(q_i, \beta_i) = Cov\left(\frac{1}{c}(a_i - \gamma\beta_i), \beta_i\right) \quad (14)$$

$$= \frac{1}{c} (Cov(a_i, \beta_i) - \gamma Var(\beta_i)) \quad (15)$$

is positive or negative. This relationship will only be positive if the covariance between a_i and β_i is sufficiently positive, so that higher betas are associated with higher initial returns on average.

Since there is typically an equilibrium relationship between risk and return, it may be tempting to think that a_i should necessarily be strongly positively related to β_i . However, this relationship is typically between risk and the marginal return. In this model, a_i only reflects the *initial* marginal return starting from zero investment in an asset, whereas $a_i - cq_i$ reflects the correct marginal return. In equilibrium, investment in a highly profitable asset will drive down that asset's marginal return until the marginal return equals $\gamma\beta_i$, so that it exactly compensates for the asset's risk. There is no similar logic for why the initial marginal return, a_i , should inherently be strongly related to β_i .

Next, I consider what the portfolio beta and expected return look like in this case. The portfolio beta is

$$\beta_p = \int_0^1 q_i \beta_i di. \quad (16)$$

The expected portfolio beta is

$$E[\beta_p] = E[q_i \beta_i] = E[q_i]E[\beta_i] + Cov(q_i, \beta_i). \quad (17)$$

This familiar last term will determine how the portfolio beta compares to what would be expected ignoring the link between risk and portfolio choice. If this term is positive, the bank will adjust its portfolio to scale up risky investments. If this term is negative, it will adjust its portfolio to scale down risky investments.

The portfolio's expected return is similarly

$$\mu_p = \int_0^1 \left(a_i q_i - \frac{c}{2} q_i^2 \right) di. \quad (18)$$

It can be written as

$$\mu_p = E[a_i q_i] - \frac{c}{2} E[q_i^2] \quad (19)$$

$$= E[a_i]E[q_i] + Cov(a_i, q_i) - \frac{c}{2} (Var[q_i] + E[q_i]^2) \quad (20)$$

2.2 The first-best solution

Here I consider the first-best outcome, in which the regulator knows the true a_i and β_i parameters and can select investment quantities q_i to maximize social welfare. In this case, the regulator maximizes

$$\max_{\{q_i\}} \int_0^1 \left(a_i q_i - \frac{c}{2} q_i^2 - (\gamma + \eta) \beta_i q_i \right) di \quad (21)$$

where there are no expectations due to the regulator's perfect knowledge of all parameters. As will be the case in other sections, I will focus on understanding the behavior of the unconditional expectations of the portfolio beta, portfolio return, and social welfare.

Proposition 1. *In the first-best case, equilibrium investment in asset i is*

$$q_i = \frac{1}{c} (a_i - (\gamma + \eta) \beta_i). \quad (22)$$

The equilibrium expected portfolio beta is

$$E[\beta_p] = \frac{1}{c} (E[a_i \beta_i] - (\gamma + \eta) E[\beta_i^2]) \quad (23)$$

The equilibrium expected portfolio return is

$$E[\mu_p] = \frac{1}{2c} (E[a_i^2] - (\gamma + \eta)^2 E[\beta_i^2]) \quad (24)$$

The equilibrium social welfare is

$$E[\mu_p - (\gamma + \eta) \beta_p] = \frac{1}{2c} E[(a_i - (\gamma + \eta) \beta_i)^2]. \quad (25)$$

Proof. See Section C.1. □

The most interesting result is that social welfare depends on the expectation of $a_i - (\gamma + \eta) \beta_i$ squared. This term can be interpreted as the initial risk-adjusted return from the regulator's point of view. Whenever the initial risk-adjusted return differs from zero, the bank can make socially efficient investments.

I also introduce a general result that is useful for calculating expected social welfare under various cases.

Proposition 2. *For all given random variables q_i representing a bank's investment choice, the expected social welfare can be expressed as*

$$\frac{c}{2} \left(E[(q_i^{fb})^2] - E[(q_i - q_i^{fb})^2] \right) \left(\quad \right) \quad (26)$$

where the first-best investment is

$$q_i^{fb} = \frac{1}{c} (a_i - (\gamma + \eta)\beta_i). \quad (27)$$

Proof. See Section C.2. □

Intuitively, social welfare can be decomposed into one piece that depends on welfare under the first-best case and another piece that depends on deviations away from the first-best. A particularly useful application of this result is comparing social welfare between two cases, for which it only becomes necessary to calculate the expected squared deviations from first-best.

3 Regulation through taxing estimated risk

Here I consider how a regulator might use taxes on banks' estimated risk to regulate their risk-taking. In interpreting the results, literal taxes are not required, but instead something that imposes a direct or indirect cost to the bank. For example, risk-sensitive deposit insurance premiums would be a direct cost while higher steady-state capital requirements would be an indirect cost. Section A discusses the relationship between setting taxes and capital requirements in further detail. Additionally, this section assumes that the bank's starting level of equity is exogenous. Section B discusses an extension in which the bank selects its desired level of equity.

3.1 Linear tax

Here the regulator picks a tax k_i that the bank must pay for each unit of asset i . The proceeds of this tax do not have any social benefit, so they don't enter the regulator's objective function, but the bank optimizes conditional on the tax. The bank optimizes

$$\max_{\{q_i\}} \int_0^1 \left(a_i q_i - \frac{c}{2} q_i^2 - \gamma \beta_i q_i - k_i q_i \right) di \quad (28)$$

taking k_i as given. The regulator optimizes

$$\max_{\{k_i\}} E_R \left[\int_0^k \left(a_i q_i - \frac{c}{2} q_i^2 - (\gamma + \eta) \beta_i q_i \right) di \right] \quad (29)$$

taking into account how the choice of k_i affects the bank's choice of q_i as well as information contained in the signals $\hat{a}_i$ and $\hat{\beta}_i$.

Proposition 3. *The bank optimally selects*

$$q_i = \frac{1}{c} (a_i - \gamma \beta_i - k_i) \quad (30)$$

The regulator optimally selects

$$k_i = \eta E_R[\beta_i] \quad (31)$$

The equilibrium social welfare relative to the first-best is

$$-\frac{1}{2c} \eta^2 \text{Var}(\tilde{\beta}_i) \quad (32)$$

where

$$\tilde{\beta}_i = \beta_i - E_R[\beta_i] \quad (33)$$

reflects the regulator's expectational error.

Proof. See Section C.3. □

The regulator's optimal policy is to set the tax equal to the expected risk given signals. At first glance this result may seem straightforward, but it does run counter to common intuitions. In practice, financial regulation is often calibrated to be conservative. One argument for this conservatism is that banks will take advantage of any weaknesses in the regulatory framework. If the regulator incorrectly believes that an asset is less risky than it actually is, then banks will overinvest in that asset. Calibrating regulation to be more conservative than the regulator's expectations is then argued as necessary to prevent this type of gaming.

This force exists in this model, since banks will overinvest in assets if regulators underestimate their risk. The reason why the optimal solution does not call for conservatism is that such conservatism carries a cost as well. If regulators overestimate the risk of an asset, then banks will underinvest in it and thus not make socially-desirable investments. Due to the specification of the risk-taking externality, $\eta \beta_p$, as a linear function, the costs

of overinvestment and underinvestment are similar and therefore the regulator targets their expectation of optimal investment.

Extensions that allow for a convex risk-taking externality, so that the extra costs of too much risk are substantially more than the benefits of reducing risk, can give rise to conservatism. Section 3.2 covers the topic in more detail.

3.1.1 Is a linear tax better than command?

Here I consider whether a linear tax outperforms the command case, in which a regulator directly selects the banks' investments. Intuitively, it would seem that the linear tax would always outperform since it controls for the expected externality while still allowing the bank to make use of its private information. However, it turns out that, for certain parameters, it can be a better option for a regulator to select investments directly.

First, I'll describe the outcome of the command case. Here, the regulator optimizes

$$\max_{\{q_i\}} E_R \left[\int_0^1 (a_i q_i - \frac{c}{2} q_i^2 - \gamma \beta_i q_i - k_i q_i) di \right], \quad (34)$$

which through straightforward calculus leads to the solution

$$q_i^{com} = \frac{1}{c} (E_R[a_i] - (\gamma + \eta) E_R[\beta_i]). \quad (35)$$

Proposition 4. *Social welfare under the command case relative to the first-best is*

$$-\frac{1}{2c} \text{Var}(\tilde{a}_i - (\gamma + \eta) \tilde{\beta}_i). \quad (36)$$

Social welfare under the command case is higher than in the linear tax case when

$$2(\gamma + \eta) \text{Cov}(\tilde{a}_i, \tilde{\beta}_i) > \text{Var}(\tilde{a}_i) + \gamma(\gamma + 2\eta) \text{Var}(\tilde{\beta}_i), \quad (37)$$

where $\tilde{x} = x - E_R[x]$ is the expectational error of the regulator.

Proof. See Section C.4. □

This expression is not deeply intuitive, but it illustrates that there are some circumstances under which the command case would be preferred to the linear tax.

A necessary, but not sufficient, condition for command to be preferred to a linear tax is for $\text{Cov}(\tilde{a}_i, \tilde{\beta}_i) > 0$. In that case, when the regulator underestimates the beta, the regulator also likely underestimates the expected return. The bank would then be more likely to invest more in assets for which the regulator has underestimated the risk.

3.2 Generalized social cost function with a linear tax

So far I have considered a known linear social cost function of the form $\eta\beta_p$. However, there may be non-linearities in practice. For example, risk may have a small marginal social cost at low levels if banks are very unlikely to default. But at higher levels of risk, and with higher probabilities of defaulting, the marginal social cost may rise dramatically. Additionally, there may be uncertainty as to the magnitude of the social costs, even for a known level of risk.

To address these concerns, I consider a generalized social cost function that allows for both non-linearities and uncertainty over costs.

Proposition 5. *Suppose that the social cost to bank risk-taking is a function $S(\beta_p, \eta)$ that is twice-differentiable, increasing in both arguments, and satisfies $\frac{\partial^2 S}{\partial \beta_p^2}(\beta_p, \eta) \geq 0$. $\beta_p = \int_0^1 q_i \beta_i di$ is the bank's portfolio beta and η is an exogenous random variable. Then the regulator's optimal linear tax is*

$$k_i = E_R \left[\frac{\partial S}{\partial \beta_p}(\beta_p, \eta) \beta_i \right] \left(\right. \quad (38)$$

$$= E_R \left[\frac{\partial S}{\partial \beta_p}(\beta_p, \eta) \right] \left(E_R [\beta_i] + Cov_R \left(\frac{\partial S}{\partial \beta_p}(\beta_p, \eta), \beta_i \right) \right) \left(\right. \quad (39)$$

Proof. See Section C.5. □

One immediate observation is that if the β_i are independent across assets, then the covariance term will be zero. In that case, the regulator's solution is the same as in the case of the known linear social cost, except replacing the marginal social cost of risk with the *expected* marginal social cost, which is $E_R[\frac{\partial S}{\partial \beta_p}(\beta_p, \eta)]$.

If there is a systematic shock to β_i , then the regulator should additionally consider the covariance term. Assets whose riskiness is higher when the social cost of risk is high should receive a higher tax.

3.3 Non-linear tax

Here the problem is similar to before, except that instead of picking a single linear tax governed by k_i , the regulator sets a non-linear tax schedule $k_i(q_i)$. These non-linear taxes will implicitly make use of information contained in the bank's choice of q_i . All else equal, banks prefer to increase q_i when expected returns are high (i.e., a_i is high) or risk is low (i.e., β_i is low), so the bank's choice of q_i is a signal of the combination of these two parameters. Using the non-linear tax is essentially a way for the regulator to set a tax conditioned not only on the signals $\hat{a}_i$ and $\hat{\beta}_i$, but also on q_i .

More specifically, the bank observes the choice of $k_i(q)$ and then selects its portfolio taking it as given. The bank therefore optimizes

$$\max_{\{q_i\}} \int_0^1 (a_i q_i - \frac{c}{2} q_i^2 - \gamma \beta_i q_i - k_i(q_i)) di \quad (40)$$

while the regulator maximizes

$$\max_{\{k_i(q_i)\}} E \left[\int_0^1 (a_i q_i - \frac{c}{2} q_i^2 - (\gamma + \eta) \beta_i q_i) di \right]. \quad (41)$$

Since adding a constant to taxes has no effect on the bank's choice, I select tax schedules that don't charge banks if they hold nothing (i.e., $k_i(0) = 0$).

Proposition 6. *The bank optimally selects*

$$q_i = \frac{1}{c} (a_i - \gamma \beta_i - k'_i(q_i)). \quad (42)$$

The regulator selects $k_i(q_i)$ to satisfy

$$k'_i(q_i) = \eta E[\beta_i | q_i, \hat{a}_i, \hat{\beta}_i], \quad (43)$$

which leads to a choice of optimal tax schedule

$$k_i(q_i) = \left(\underbrace{\frac{\eta \lambda_0}{1 + (\eta/\gamma) \lambda_z}}_{\omega_0} + \underbrace{\frac{\eta \lambda_a}{1 + (\eta/\gamma) \lambda_z}}_{\omega_a} \hat{a}_i + \underbrace{\frac{\eta \lambda_b}{1 + (\eta/\gamma) \lambda_z}}_{\omega_b} \hat{\beta}_i - \underbrace{\frac{(\eta/\gamma) \lambda_z}{1 + (\eta/\gamma) \lambda_z} \frac{c}{2}}_{\omega_q} q_i \right) q_i, \quad (44)$$

where λ_a , λ_b , and λ_z match the coefficients from a multivariable regression of β_i on $\hat{a}_i$, $\hat{\beta}_i$, and $z_i = \beta_i - \frac{1}{\gamma} a_i$ and λ_0 is the constant from this regression.

Proof. See Section C.6. □

In the case of a linear tax, the regulator sets the tax equal to their best guess of the beta. In the case of a non-linear tax, the regulator sets the marginal tax for quantity q_i equal to their best guess of the beta taking into account the information contained in q_i . These results are more general and do not depend strongly on this particular environment.

The specific assumptions in this case lead to a specific expression for $k_i(q_i)$. Essentially, there is a fixed linear component of the tax that depends on $\hat{a}_i$ and $\hat{\beta}_i$, which reflects regulators using information gleaned from their signals about the true β_i . There is also a quadratic term that depends on q_i .

4 Effects of non-disclosure

Here I consider the policy of specifying taxes, similar to Section 3, but with the added twist that the regulator does not reveal these taxes to the bank until after they have made their portfolio choice.

4.1 Fully undisclosed linear tax

In this case, the regulator specifies a linear tax, but the bank does not know the value of this tax until after it has selected its portfolio. One interpretation is that this scenario is similar to a stress test in which the bank does not know the regulator's model. While not implemented in practice, another interpretation is that the bank selects its portfolio and the regulator only reveals the associated risk weights after the fact.

In this situation, the bank solves

$$\max_{\{q_i\}} \int_0^1 (a_i q_i - \frac{c}{2} q_i^2 - \gamma \beta_i q_i - E_B[k_i] q_i) di, \quad (45)$$

where the bank's information set does not include the regulator's signals. Meanwhile, the regulator solves

$$\max_{\{k_i\}} E_R \left[\int_0^1 (a_i q_i - \frac{c}{2} q_i^2 - (\gamma + \eta) \beta_i q_i) di \right] \left(\quad (46)$$

taking into account the bank's choice function. I solve for the set of solutions $\{q_i\}$ and $\{k_i\}$ that jointly solve each optimization problem.

Proposition 7. *One solution is for the bank to select*

$$q_i = \frac{1}{c} (a_i - (\gamma + \eta) \beta_i) \quad (47)$$

and the regulator to select

$$k_i = \eta \hat{\beta}_i. \quad (48)$$

This solution achieves the first-best outcome for the regulator.

Proof. See Section C.7. □

Recall that $\hat{\beta}_i \neq E_R[\beta_i]$, so the regulator is explicitly *not* setting the tax equal to their best guess of what the true beta is. The rationale is that the regulator wants the tax to respond one-for-one to changes in the true underlying beta. The regulator's best guess

of the true beta, $E_R[\beta_i]$, includes some regularization toward $\bar{\beta}$ that typically results in a response that is less than one-for-one.

For practical application, the lesson here is that when regulators are not sharing ex-ante details about taxes, they should respond very strongly to any information they receive about an asset's riskiness. Otherwise, banks will assume that regulators will not adjust taxes sufficiently in response to information.

4.2 Partially-disclosed linear tax

In practice, banks may have some sense of deficiencies in regulators' models, even if they may not have full knowledge of those models. For example, in the context of stress tests, banks may learn about some of the features of regulators' models. I model this situation by giving banks noisy signals of the regulator's signals, which they can use to guess the likely level of tax that regulators will specify.

The set-up is exactly the same as before, except now banks receive a noisy signal of the regulator's signals. For convenience, I adopt the notation that

$$x_i = [a_i, \beta_i] \tag{49}$$

$$e_i = [e_i^a, e_i^b] \tag{50}$$

$$\hat{x}_i = [\hat{a}_i, \hat{\beta}_i] = x_i + e_i \tag{51}$$

and then suppose that the bank receives a signal

$$s_i = \hat{x}_i + w_i, \tag{52}$$

where

$$w_i \sim \mathcal{N}(0, \Sigma_w) \tag{53}$$

and the w_i are independent across assets.

Proposition 8. *If the bank receives noisy signals of the regulator's signals, then an optimal k_i must satisfy*

$$E_R[E_B[k_i]] = \eta E_R[\beta_i]. \tag{54}$$

One solution satisfying this condition is for the regulator to set

$$k_i = \eta(\bar{\beta} + m'(\hat{x}_i - \bar{x})), \tag{55}$$

where m is a 2×1 vector equal to

$$m = (Var(x_i) + Var(e_i)(Var(e_i) + Var(w_i))^{-1}Var(e_i))^{-1}Cov(x_i, \beta_i) \quad (56)$$

$$= Var(E_B[\hat{x}_i])^{-1}Cov(E_B[\hat{x}_i], \beta_i), \quad (57)$$

which matches the coefficient of a linear regression of β_i on $E_B[\hat{x}_i]$.

Proof. See Section C.8. □

Intuitively, the condition for the optimality of k_i is that the regulator's best guess of the bank's best guess of the regulator's tax equals the regulator's best guess of the risk externality imposed by the portfolio. Put another way, the regulator wants to calibrate the tax to give the bank the correct incentives on average.

Existing results emerge as special cases of this framework. First, if the bank has a perfect signal of the regulator's information, then $Var(w_i) = 0$. This case is equivalent to when the regulator pre-announces the loadings, in which case the result is $k_i = \eta E_R[\beta_i]$. Second, if the bank has no idea of the regulator's information, which is modeled as w_i approaching infinite variance, then $Var(e_i)(Var(e_i) + Var(w_i))^{-1}Var(e_i) \rightarrow 0$ and $m = [0, 1]$ is the solution, which aligns with the earlier result of the regulator setting $k_i = \eta \hat{\beta}_i$.

To interpret m , note that the bank's best guess of the tax in equilibrium is

$$E_B[k_i] = \eta(\bar{\beta} + m'E_B[\hat{x}_i - \bar{x}]). \quad (58)$$

The bank's best guess of the tax depends on the bank's best guess of the regulator's signal, $E_B[\hat{x}_i]$. If the regulator calibrates m such that it matches the coefficient from a regression of β_i on $E_B[\hat{x}_i]$, it most closely aligns the bank's estimate of the tax with the socially-optimal tax.

What happens as the bank's signal becomes noisier (i.e., the variance of w_i rises)? First, note that $E_B[\hat{x}_i] = x_i + E_B[e_i]$, so the bank's expectation of the regulator's guess depends both on their perfect information of the true state along with their imperfect information of the regulator's error. Holding all else constant, raising the noise of the bank's signal makes the bank's information on e_i less useful, so $E_B[\hat{x}_i]$ will be closer to x_i , and the bank's guess will be more responsive to the true state of the world. Intuitively, the less the bank knows about the regulator's errors, the less able the bank is to take advantage of those errors, and the more the bank relies on its knowledge of the asset's true riskiness.

4.3 Non-disclosure with an aversion to idiosyncratic volatility

A potential concern with non-disclosure of taxes is that it exposes banks to additional uncertainty over their future cash flows. In the baseline framework, banks are only averse to systematic risk, so this additional uncertainty does not pose any cost. But, in practice, factors such as capital market imperfections may make banks averse even to idiosyncratic volatility (Froot and Stein, 1998). Therefore, non-disclosure can pose costs to banks that the baseline framework does not capture.

To incorporate these costs, I add a γ_I term reflecting banks' aversion to idiosyncratic volatility resulting from the uncertain taxes. One important question is whether the regulator considers idiosyncratic risk as imposing a social cost or not. To capture a wide range of possibilities, I assume that the regulator recognizes a social cost η_I associated with the bank bearing idiosyncratic volatility. Since the regulator already recognizes indirect effects of volatility affecting banks' portfolio choice, the η_I term reflects only the direct costs of the bank bearing idiosyncratic volatility, even controlling for portfolio choice.

For tractability reasons, I focus on the simpler case in which there is only one asset. In this case, the idiosyncratic risk only comes from the tax on the single asset. In the case with multiple assets, idiosyncratic risk will depend on the volatility of taxes across many assets, in which case the correlation of the errors becomes important. To avoid introducing those complexities for now, I focus on one asset.

Proposition 9. *Suppose that banks are averse to the volatility of tax payments such that they maximize*

$$\max_q (a - \gamma\beta)q - \frac{c}{2}q^2 - E_B[kq] - \gamma_I \sigma_B(kq), \quad (59)$$

where $\sigma_B(kq)$ is the standard deviation of the taxes paid with respect to the bank's information set.

Additionally, suppose that the regulator recognizes a social cost of η_I from the volatility of tax payments, so that the regulator's objective is to select a tax k to maximize

$$E_R[(a - \gamma\beta)q - \frac{c}{2}q^2 - \eta\beta q - \eta_I \sigma_B(kq)]. \quad (60)$$

The regulator can achieve the first-best outcome by setting

$$k = \eta\hat{\beta} + (\eta_I - \gamma_I) \underbrace{\eta\sigma_e}_{\sigma_B(k)} \text{sgn}(q), \quad (61)$$

which results in overall taxes paid of

$$kq = \eta\hat{\beta}q + (\eta_I - \gamma_I)\eta\sigma_e|q|. \quad (62)$$

Proof. See Section C.9. □

The first piece of this expression, $\eta\hat{\beta}$, is the same as in the baseline case without an aversion to idiosyncratic volatility. As before, the bank's best guess of $\hat{\beta}$ is the true β , so non-disclosure by regulators forces the bank to use their information based on the true state of the world.

The second piece of this expression relates to the idiosyncratic volatility. If the private and social costs of idiosyncratic volatility differ (i.e., to the extent that γ_I and η_I differ), the regulator should adjust the taxes to align the bank's incentives. For example, if $\gamma_I = \eta_I$, so that the two are already aligned, there is no need to take action.

But consider the case in which $\gamma_I > 0$ and $\eta_I = 0$, so that there is only a private cost to volatility. For positive q , the regulator should optimally reduce the size of the tax. The reason is that the volatility is already dissuading the bank from investing, so the tax does not need to be as high to achieve the optimal level of investment. In this case, the regulator can compensate for the costs imposed by higher uncertainty in the tax by reducing the average size of the tax. The idea generalizes: The regulator can compensate for the costs imposed by higher regulatory uncertainty by reducing the average tightness of regulations.

5 Taxes on profits

This section explores the potential of using taxes on bank profits to effectively reduce their risk aversion and thus influence their portfolio choice.

5.1 Achieving first-best through a tax on expected profits

In the generally-infeasible case in which the regulator has perfect knowledge of the bank's expected profits, then it's possible for the regulator to achieve a first-best outcome through an appropriately-calibrated tax on them. The intuition is that reducing expected profits also reduces the profit per unit of risk the bank takes, which incentivizes the bank to reduce its risk and has an identical effect as reducing the bank's risk aversion.

Proposition 10. *The regulator can achieve the first-best outcome by taxing a fraction $\frac{\eta}{\gamma+\eta}$ of the bank's expected profits.*

Proof. In this case, the regulator imposes a tax of

$$T = \frac{\eta}{\gamma + \eta} E_B[\Pi] = \int_0^k \frac{\eta}{\gamma + \eta} (a_i q_i - \frac{c}{2} q_i^2) di. \quad (63)$$

The bank's objective is now to maximize

$$\max_{\{q_i\}} E_B[M(\Pi - T)] = \max_{\{q_i\}} \int_0^1 \left(\frac{\gamma}{\eta + \gamma} \left(a_i q_i - \frac{c}{2} q_i^2 \right) - \gamma \beta_i q_i \right) di. \quad (64)$$

Multiplying by $(\gamma + \eta)/\gamma$ does not change the optimal solution and yields the maximization problem

$$\max_{\{q_i\}} \int_0^k \left(a_i q_i - \frac{c}{2} q_i^2 - (\gamma + \eta) \beta_i q_i \right) di, \quad (65)$$

which exactly matches the regulator's objective function and therefore leads to a first-best outcome. $\square$

I again emphasize that this outcome is generally not feasible since it relies on the regulator having perfect knowledge of a_i . However, it suggests that feasible strategies that aim to closely approximate a tax on expected returns might be a fruitful course of action. The following sections examine different feasible approaches in the same spirit.

5.2 The failure of a flat tax on ex-post profits

If a tax on expected profits can achieve the first-best, then could a tax on realized profits do the same? Since the expected value of a tax on realized profits equals a tax on expected profits, then it would at first appear like this approach should work. But it ultimately does not.

To see why, suppose the regulator adds a tax on profits with the rate set to τ . Note that this tax applies to both positive and negative profits, so that the regulator effectively subsidizes the bank's losses. While explicit subsidies for losses are uncommon in practice, carrying forward losses to reduce future taxes or using losses to reduce taxable income for profitable parts of the business have the same effect. In this case, the bank would maximize

$$E_B[M(1 - \tau)\Pi] = E_B[(1 - \tau)\Pi] - \gamma Cov_B((1 - \tau)\Pi, F). \quad (66)$$

Since the tax only scales the bank's objective by a constant, it does not affect the bank's optimization problem and thus has no effect on the bank's portfolio choice.

The intuition for this result is that while the tax reduces profits, it also reduces risk by the same amount. The bank makes smaller profits on the upside, but also smaller losses on the downside. Therefore, the risk-return trade-off remains unchanged along with the bank's portfolio choice.

5.3 Achieving the first-best through a tax on ex-post profits

While a flat tax on ex-post profits does not work, it is possible to achieve the first-best through an appropriately-calibrated tax on ex-post profits.

Proposition 11. *If the regulator sets an ex-post tax of*

$$T = \frac{1}{M} \frac{\eta}{\eta + \gamma} \Pi, \quad (67)$$

then portfolio choice will match the first-best case.

Proof. From Proposition 10, a tax that is set to be a fraction $\eta/(\eta + \gamma)$ of the bank's expectation of profits achieves the first best. The regulator can set an ex-post tax T with a present value equal to this tax by selecting T to satisfy

$$E_B[MT] = \frac{\eta}{\eta + \gamma} E_B[\Pi]. \quad (68)$$

Setting T as proposed leads to

$$E_B \left[M \underbrace{\frac{\eta}{\eta + \gamma} \frac{\Pi}{M}}_T \right] = \frac{\eta}{\eta + \gamma} E_B[\Pi], \quad (69)$$

so that the present value of the tax to the bank equals a tax on expected profits. $\square$

From before, the regulator would ideally want to tax the bank's expected returns. Unfortunately, a direct tax is infeasible because the regulator does not know asset-specific expected returns. However, the regulator can set a stochastic tax whose present value, from the bank's perspective, equals some fraction of the portfolio's expected return. State-dependent tax rates are required to achieve this result. Since M is high in "bad" times and low in "good" times, the taxes should be higher in "bad" times than in "good."

In practice, the taxes may not necessarily have to explicitly be state-dependent. For example, since bank profits are likely higher during good times than bad, setting a progressive tax rate on profits might achieve a similar outcome. Even a binary approach, such as imposing an extra tax on profits above a particular threshold, may achieve an approximately similar outcome.

6 Policy applications

While the model is stylized, it yields several general insights for specifying financial regulation in a way that addresses the information asymmetry between the regulator and the regulated, particularly for capital regulation.

First, regulators can benefit by not disclosing information to banks. Intuitively, if a bank does not know the regulator’s model ahead of time, they will need to use their best guess. If the bank thinks that the regulator’s model will be correct on average, then the bank will act as though the regulator had used correct estimates of risk. Even if the bank has some information on the regulator’s model, there is still benefit in disclosing as little information as possible to reduce possibilities to game the model.

A potential concern from this approach is that hiding information may cause regulatory uncertainty that would dissuade socially-valuable investment. However, if uncertainty over an asset’s risk weight is dissuading banks from investing in it, reducing the average level of that risk weight could counteract the effect of the uncertainty.

This broad approach of not disclosing information to banks is reflected in the spirit of current stress tests, in which banks provide information on their portfolios and do not have full knowledge of regulators’ model. One of the touted benefits is that it is harder for banks to game the model. While non-disclosure is most strongly associated with stress tests, it could apply to other parts of the regulatory framework. For example, in the Basel III capital framework, risk weights are currently pre-specified by regulators. A possibility is for regulators to update the risk weights and only reveal either the individual risk weights or overall capital requirements to banks after they have specified their portfolios.

Second, regulators can apply non-linear formulas that take into account information contained in the bank’s portfolio choice. The most direct application would be risk weights that automatically change as banks concentrate investment in particular sectors. While the model is too stylized for the formulas to directly apply to policy settings, regulators can ask the question: If a bank is investing more in a particular type of asset, what information does it reveal? In some cases, concentration in an asset may indicate that the regulator has underestimated risk, and using a tool similar to a non-linear risk weight may be a way of leaning back against that underestimation.⁸ But higher concentration might not indicate risk, but instead indicate an abundance of investment opportunities. Distinguishing these cases is important.

Practical application of such non-linear formulas should also account for bank specialization. Some banks may specialize in particular lines of business and would therefore

⁸Sector-specific countercyclical capital buffers can provide a similar outcome if adjusted appropriately. But if there are any difficulties in activating such buffers, an automatic non-linear rule may be able to adjust more quickly.

naturally have higher concentrations. Penalizing specialized banks is likely socially undesirable. For practical implementation, it may be better to use measures of concentration for the aggregate banking system rather than individual banks.

Third, regulators could implement taxes on ex-post profits. The intuition is that banks take risk to generate higher returns. Reducing expected returns through a tax thus reduces the bank’s incentives to take risk, which can effectively make them more risk averse while still using their private information. However, it is important to take into account that taxes on ex-post profits affect both the after-tax risk and expected return.

In theory, one optimal solution is a state-dependent tax that is higher during “good times” (i.e., when investors’ SDF is lower). The level of this tax would rise with the magnitude of the externality that the bank’s risk-taking imposes. An explicitly state-dependent tax could be calibrated based on measures of economic performance, such as unemployment or GDP growth. However, since bank profitability is likely to be highly correlated with economic performance, a progressive tax might also approximate the ideal outcome. For example, a higher rate on profits above some predetermined threshold might discourage banks from excessive risk-taking by reducing the after-tax payoff of doing so.

Empirically, Meiselman et al. (2018) show that high profitability for banks predicts higher tail risk for both the 2007-2008 financial crisis as well as the 1980s savings and loan crisis. Given the empirical link between profits and risk, a well-calibrated tax on profits can effectively be viewed as a tax on risk, without requiring any knowledge on the regulator’s part of the riskiness of the bank’s portfolio.

Finally, regulators could supplement their calibrations of risk weights or similar regulatory tools with information on assets’ profitability as well. The model frequently indicates that regulators should set taxes proportionally to the regulator’s estimate of the asset’s risk. Implicit in the model is that the regulator’s estimate uses all available signals, including signals of profitability. Since risk and profitability are likely closely related in practice, information on profitability can help to improve estimates of risk, particularly if such information is easier to observe. For example, it’s easier to observe a loan’s credit spread than its probability of defaulting. In that case, a regulator might adjust risk weights (or other tools) based on observed credit spreads.

7 Conclusion

In this paper, I address the problem of how to regulate bank portfolio choice taking into account the asymmetry of information between banks and regulators. I construct a tractable model that explicitly accounts for this asymmetry. The model is flexible

enough to allow for the consideration of a wide range of regulatory tools. I then consider the efficacy of several common regulatory tools along with proposing three less-common tools: Not disclosing taxes to banks until after portfolio selection, non-linear taxes that respond to information contained in banks' portfolio choice, and state-dependent taxes on banks' realized profits.

While the model is intentionally stylized to aid in communicating intuitions, the broad takeaways could be applied in practice (whether separately or concurrently). Regulators could consciously not disclose information from banks to prevent them from gaming regulation, similar to how they already do so for stress tests. Non-linear taxes (or risk weights) could be used to automatically respond to banks concentrating their investments in a particular sector. And taxes on ex-post profits can be used to incentivize banks to act in a more risk averse manner, even if regulators know almost nothing about banks' risks.

Future work can extend the model along several dimensions. One extension is to focus on dynamic interactions, particularly in the case of information non-disclosure. Since banks can learn information about regulators' models over time by observing outcomes, regulators may need to adjust their models or intentionally introduce noise to keep banks from learning too much. Another extension is to ease the single-factor assumption. While a single factor is an implicit assumption underlying much regulation (for example, Gordy (2003) discusses the importance of a single-factor assumption within capital regulation of the banking book), market risk must account for many correlated factors and hedges. Addressing this problem would further aid in regulating bank portfolio choice in the important, but more complicated, setting of banks' trading books.

References

- Celerier, C., Kick, T. K., and Ongena, S. (2019). Taxing Bank Leverage: The Effects on Bank Portfolio Allocation. SSRN Scholarly Paper ID 2829326, Social Science Research Network, Rochester, NY.
- Chan, Y.-S., Greenbaum, S. I., and Thakor, A. V. (1992). Is Fairly Priced Deposit Insurance Possible? *The Journal of Finance*, 47(1):227–245.
- Domar, E. D. and Musgrave, R. A. (1944). Proportional Income Taxation and Risk-Taking. *The Quarterly Journal of Economics*, 58(3):388–422.
- Froot, K. A. and Stein, J. C. (1998). Risk management, capital budgeting, and capital structure policy for financial institutions: An integrated approach. *Journal of Financial Economics*, 47(1):55–82.

- Giammarino, R. M., Lewis, T. R., and Sappington, D. E. M. (1993). An Incentive Approach to Banking Regulation. *The Journal of Finance*, 48(4):1523–1542.
- Glasserman, P. and Kang, W. (2014). Design of Risk Weights. *Operations Research*, 62(6):1204–1220.
- Goldstein, I. and Leitner, Y. (2018). Stress tests and information disclosure. *Journal of Economic Theory*, 177:34–69.
- Gordy, M. B. (2003). A risk-factor model foundation for ratings-based bank capital rules. *Journal of Financial Intermediation*, 12(3):199–232.
- John, K., John, T. A., and Senbet, L. W. (1991). Risk-shifting incentives of depository institutions: A new perspective on federal deposit insurance reform. *Journal of Banking & Finance*, 15(4):895–915.
- Kim, D. and Santomero, A. M. (1988). Risk in Banking and Capital Regulation. *The Journal of Finance*, 43(5):1219–1233.
- Meiselman, B. S., Nagel, S., and Purnanandam, A. (2018). Judging Banks’ Risk by the Profits They Report. SSRN Scholarly Paper ID 3169730, Social Science Research Network, Rochester, NY.
- Morgan, D. P. and Ashcraft, A. B. (2003). Using Loan Rates to Measure and Regulate Bank Risk: Findings and an Immodest Proposal. *Journal of Financial Services Research*, 24(2):181–200.
- Perotti, E. and Suarez, J. (2018). A Pigovian Approach to Liquidity Regulation. *27th issue (November 2011) of the International Journal of Central Banking*.
- Rochet, J.-C. (1992). Capital requirements and the behaviour of commercial banks. *European Economic Review*, 36(5):1137–1170.
- Shackelford, D. A., Shaviro, D. N., and Slemrod, J. (2010). Taxation and the financial sector. *National Tax Journal*, 63(4.1):781–806.
- Stiglitz, J. E. (1969). The Effects of Income, Wealth, and Capital Gains Taxation on Risk-Taking. *The Quarterly Journal of Economics*, 83(2):263–283.
- Weitzman, M. L. (1974). Prices vs. Quantities. *The Review of Economic Studies*, 41(4):477.
- Wu, H.-M. and Zhao, Y. (2016). Optimal Leverage Ratio and Capital Requirements with Limited Regulatory Power. *Review of Finance*, 20(6):2125–2150.

A Relationship between taxes and capital requirements

The aim of this section is to demonstrate how approaches based on taxes and capital requirements compare with each other. The exact relationship depends on the assumptions about capital.

In several cases, there is an exact equivalence. In the most straightforward case, the supply of capital is elastic with constant cost relative to other funding sources of r_c . In that case, a capital requirement of m_i per quantity held of asset i has the same effect on the bank's portfolio choice as an asset-specific tax of $k_i = m_i \cdot r_c$.

There is also an exact equivalence between the two approaches when the quantity of capital is exogenous and shocks to a_i and β_i are idiosyncratic (i.e., independent across assets). In that case, setting a capital requirement on a specific asset has the same effect on the bank's portfolio choice as setting a tax equal to the shadow cost of the capital requirement.

Proposition 12. *Suppose that capital is exogenously set to 1 and the regulator sets a capital requirement of the form*

$$\int_0^1 m_i q_i di \leq 1, \quad (70)$$

where m_i indicates the capital requirement for asset i . If there are no aggregate shocks to a_i and β_i (i.e., a_i and β_i are drawn independently across assets), then the Lagrange multiplier on the capital requirement λ is deterministic. The regulator can induce identical portfolio choice by the bank through setting a linear tax k_i according to

$$k_i = \lambda m_i. \quad (71)$$

Proof. See Section C.10. □

However, the two approaches are not equivalent with systematic shocks. In that case, the Lagrange multiplier λ will be stochastic. The difference between capital requirements and taxes will essentially be a choice of price-based versus quantity-based regulation, the considerations of which are discussed most notably by Weitzman (1974). In the baseline case in which there is a known linear social cost to bank risk-taking, a tax is likely the better solution. But for particular forms of non-linear social cost functions, quantity-based regulation may be preferable.

B Linear taxes with endogenous capital structure

I now consider the effects endogenous capital structure choice into the model for the case of linear taxes. The bank now selects the amount of equity with which it will fund itself. More equity reduces the social cost of bank risk taking for a given level of risk exposure. However, equity may potentially have both social and private costs relative to other sources of funding. Without delving into exactly why these costs exist, I suppose that the social cost of equity is $r_s > 0$ and the private cost of equity is $r_p \geq r_s > 0$.⁹ These reflect the change in the overall cost of capital from funding with an additional unit of equity, not the per-unit cost of equity.

For example, if the Modigliani-Miller capital structure irrelevance theorem were to hold, then $r_s = r_p = 0$. I focus on the cases in which both r_s and r_p are positive since these reflect the interesting cases; if equity were socially costless, then the solution would be to hold sufficient equity such that there is no longer any social cost.

The social cost of risk exposure is

$$\eta \frac{\int_0^1 \beta_i q_i di}{e}. \quad (72)$$

This social cost is exactly the same as before, except now the portfolio risk is scaled by the amount of equity e .

I consider a case in which the regulator uses two tools. First, the regulator sets a tax that takes the form

$$\frac{\int_0^1 k_i q_i di}{e}, \quad (73)$$

where k_i sets the magnitude of the asset-specific tax on asset i , but now it is also scaled by the level of equity e since the social cost of risk also depends on the bank's leverage. For this piece, the regulator solves for the optimal k_i .

Second, the regulator could also potentially subsidize equity at a linear rate of s . So the regulator solves for the optimal s as well.

Proposition 13. *If the bank solves*

$$\max_{\{q_i\}, e} \int_0^1 \left((a_i - \gamma \beta_i) q_i - \frac{c}{2} q_i^2 - \frac{k_i q_i}{e} \right) di - (r_p - s)e \quad (74)$$

⁹If $r_s = 0$, so that there is no social cost of equity, then the solution is for the bank to fund itself with as much equity as possible. For reasons of tractability, this model allows for potentially unlimited losses on the portfolio, including losses that are larger than the initial investment. Therefore to cover any potential loss would require equity levels approaching infinity. In a model that incorporated limited liability on investments, costless equity would imply that the bank should fund itself with 100% equity.

and the regulator solves

$$\max_{\{k_i\}, s} E_R \left[\int_0^k \left((a_i - \gamma\beta_i)q_i - \frac{c}{2}q_i^2 - \eta \frac{\beta_i q_i}{e} \right) di - r_s e \right] \left(\quad \right) \quad (75)$$

taking into account the bank's optimal choices of q_i and e , then the regulator optimally sets

$$k_i = \eta E_R[\beta_i] \quad (76)$$

and

$$s = r_p - r_s. \quad (77)$$

Proof. See Section C.11. □

The first part of this result, that the regulator sets the asset-specific tax k_i based on the asset's expected risk, is intuitively the same as the case with exogenous equity. The second part, that the regulator should set an equity subsidy, is new. Setting the subsidy like this reduces the bank's marginal cost of equity from r_p to r_s , which aligns it with the social cost.

The key difference when introducing endogenous capital structure is that tax policy also affects the bank's incentives to fund itself with equity. When the regulator can align the private and social costs of equity, then the regulator only needs to worry about aligning asset-specific taxes with expected risks.

However, the situation becomes more complicated when the regulator cannot introduce an equity subsidy and there are differences in the social and private costs of equity. In that case, asset-specific taxes can be used to influence the bank's overall selected level of equity. For example, in typical specifications, setting k_i uniformly higher than $\eta E_R[\beta_i]$ can serve as a crude tax on leverage that encourages banks to fund themselves with more equity.

C Additional Details on Proofs

C.1 The first-best solution

The first-order condition for equilibrium q_i is

$$\begin{aligned} a_i - cq_i - (\gamma + \eta)\beta_i &= 0 \\ \implies q_i &= \frac{1}{c}(a_i - (\gamma + \eta)\beta_i) \end{aligned} \tag{78}$$

as desired. The second derivative, $-c$, is negative, indicating that this solution is a maximum.

The expected portfolio beta is

$$\begin{aligned} E[\beta_p] &= E \left[\int_0^1 q_i \beta_i di \right] \\ &= E[q_i \beta_i] \\ &= \frac{1}{c} (E[a_i \beta_i] - (\gamma + \eta)E[\beta_i^2]) \left(\right. \end{aligned} \tag{79}$$

The expected portfolio return is

$$\begin{aligned} E[\mu_p] &= E \left[\int_0^1 (a_i q_i - \frac{c}{2} q_i^2) di \right] \left(\right. \\ &= E[a_i q_i] - \frac{c}{2} E[q_i^2] \\ &= \frac{1}{c} (E[a_i^2] - (\gamma + \eta)E[a_i \beta_i]) - \frac{c}{2} \frac{1}{c^2} (E[a_i^2] - 2(\gamma + \eta)E[a_i \beta_i] + (\gamma + \eta)^2 E[\beta_i^2]) \\ &= \frac{1}{2c} (E[a_i^2] - (\gamma + \eta)^2 E[\beta_i^2]) \left(\right. \end{aligned} \tag{80}$$

Social welfare, represented by the regulator's objective function, is

$$\begin{aligned} E[\mu_p - (\gamma + \eta)\beta_p] &= \frac{1}{2c} (E[a_i^2] - (\gamma + \eta)^2 E[\beta_i^2]) \left(\frac{1}{c} (\gamma + \eta) (E[a_i \beta_i] - (\gamma + \eta)E[\beta_i^2]) \right) \left(\right. \\ &= \frac{1}{c} \left(\frac{1}{2} E[a_i^2] + \frac{1}{2} (\gamma + \eta)^2 E[\beta_i^2] - (\gamma + \eta)E[a_i \beta_i] \right) \left(\right. \\ &= \frac{1}{2c} E[(a_i - (\gamma + \eta)\beta_i)^2]. \end{aligned} \tag{81}$$

C.2 Expression of social welfare

By definition, the regulator's expected social welfare conditional on their information set is

$$\begin{aligned} & E_R[a_i q_i - \frac{c}{2} q_i^2 - (\gamma + \eta) \beta_i q_i] \\ &= E_R[(a_i - (\gamma + \eta) \beta_i) q_i] - \frac{c}{2} E_R[q_i^2]. \end{aligned} \quad (82)$$

Using the fact that $q_i^{fb} = \frac{1}{c}(a_i - (\gamma + \eta) \beta_i)$ and expanding $q_i = q_i^{fb} + (q_i - q_i^{fb})$ leads to

$$\begin{aligned} & c E_R[q_i^{fb} q_i] - \frac{c}{2} E_R[(q_i^{fb} + (q_i - q_i^{fb}))^2] \\ &= c E_R[q_i^{fb} q_i] - \frac{c}{2} E_R[(q_i^{fb})^2 + 2q_i^{fb}(q_i - q_i^{fb}) + (q_i - q_i^{fb})^2] \\ &= c E_R[q_i^{fb} q_i] - \frac{c}{2} E_R[(q_i^{fb})^2] - c E_R[q_i^{fb} q_i] + c E_R[(q_i^{fb})^2] - \frac{c}{2} E_R[(q_i - q_i^{fb})^2] \\ &= \frac{c}{2} (E_R[(q_i^{fb})^2] - E_R[(q_i - q_i^{fb})^2]). \end{aligned} \quad (83)$$

Taking unconditional expectations leads to the desired result.

C.3 Linear tax

The first-order condition for the bank with respect to q_i is

$$a_i - c q_i - \gamma \beta_i - k_i = 0 \quad (84)$$

$$\implies q_i = \frac{1}{c} (a_i - \gamma \beta_i - k_i) \quad (85)$$

and this q_i is a maximum since the second derivative, $-c$, is negative.

The first-order condition for the regulator with respect to k_i , taking as given the bank's choice of q_i , is

$$E_R \left[(a_i - c q_i - (\gamma + \eta) \beta_i) \frac{\partial q_i}{\partial k_i} \right] \stackrel{!}{=} 0 \quad (86)$$

$$E_R \left[(k_i - \eta \beta_i) \left(-\frac{1}{c} \right) \right] \stackrel{!}{=} 0 \quad (87)$$

$$k_i = \eta E_R[\beta_i]. \quad (88)$$

Using the result from Proposition 2, the difference in social welfare compared to the

first-best is

$$\frac{c}{2}E[(q_i - q_i^{fb})^2] \quad (89)$$

$$= \frac{c}{2} \left(\frac{\eta}{c} \right)^2 E[(-E_R[\beta_i] + \beta_i)^2]. \quad (90)$$

I define $\tilde{\beta}_i = \beta_i - E_R[\beta_i]$, which reflects the regulator's expectational error. I can then rewrite the expression as

$$= \frac{1}{2c} \eta^2 Var(\tilde{\beta}_i). \quad (91)$$

C.4 Command vs. linear tax

Using the result from Proposition 2, the social welfare relative to the first-best is

$$- \frac{c}{2} E[(q_i^{com} - q_i^{fb})^2] \quad (92)$$

$$= - \frac{c}{2} \frac{1}{c^2} E[(\tilde{a}_i - (\gamma + \eta)\tilde{\beta}_i)^2] \quad (93)$$

$$= - \frac{1}{2c} Var(\tilde{a}_i - (\gamma + \eta)\tilde{\beta}_i) \quad (94)$$

Using results on the social welfare in the linear tax case from Proposition 3, the difference in social welfare between the command and the linear tax case is

$$- \frac{1}{2c} \left(Var(\tilde{a}_i - (\gamma + \eta)\tilde{\beta}_i) - \eta^2 Var(\tilde{\beta}_i) \right) \quad (95)$$

This quantity is positive (indicating that the command solution provides higher social welfare) when

$$\eta^2 Var(\tilde{\beta}_i) > Var(\tilde{a}_i - (\gamma + \eta)\tilde{\beta}_i) \quad (96)$$

$$\implies \eta^2 Var(\tilde{\beta}_i) > Var(\tilde{a}_i) + (\gamma + \eta)^2 Var(\tilde{\beta}_i) - 2(\gamma + \eta) Cov(\tilde{a}_i, \tilde{\beta}_i) \quad (97)$$

$$\implies 2(\gamma + \eta) Cov(\tilde{a}_i, \tilde{\beta}_i) > Var(\tilde{a}_i) + \gamma(\gamma + 2\eta) Var(\tilde{\beta}_i) \quad (98)$$

C.5 Generalized social cost function and a linear tax

The regulator's goal is to maximize social welfare, which is given by

$$E_R \left[\int_0^1 \left(a q_i - \frac{c}{2} q_i^2 - \gamma \beta_i q_i \right) \left(d_i - S(\beta_p, \eta) \right) \right], \quad (99)$$

where $\beta_p = \int_0^1 q_i \beta_i di$ is the portfolio beta. The regulator sets constraints subject to the bank's choice of q_i .

First, consider the case in which a regulator sets an asset-specific linear tax k_i . Given a linear tax, the bank maximizes

$$\max_{\{q_i\}} a q_i - \frac{c}{2} q_i^2 - \gamma \beta_i q_i - k_i q_i, \quad (100)$$

which leads to the familiar optimal solution

$$q_i = \frac{1}{c} (a_i - \gamma \beta_i - k_i). \quad (101)$$

Next, turn to the regulator's problem of selecting optimal k_i taking the bank's behavior as given. The regulator solves

$$\max_{\{k_i\}} E_R \left[\int_0^1 \left(a q_i - \frac{c}{2} q_i^2 - \gamma \beta_i q_i \right) \left(di - S(\beta_p, \eta) \right) \right], \quad (102)$$

taking into account the effect on the bank's choice of q_i . The regulator's first-order condition for k_i is

$$E_R \left[\left(\left(a_i - \gamma \beta_i - c q_i - \frac{\partial S}{\partial \beta_p} \beta_i \right) \frac{\partial q_i}{\partial k_i} \right) \right] = 0, \quad (103)$$

which leads to an optimal choice of

$$k_i = E_R \left[\frac{\partial S}{\partial \beta_p} \beta_i \right] \left(\quad (104)$$

The tax should be set to the expected product of the asset beta multiplied by the marginal social cost. This term can be expanded to include a covariance as

$$k_i = E_R \left[\frac{\partial S}{\partial \beta_p} \right] \left(E_R [\beta_i] + Cov \left(\frac{\partial S}{\partial \beta_p}, \beta_i \right) \right) \left(\quad (105)$$

I then verify that the second-order condition holds with respect to k_i . The first derivative with respect to k_i can be written as

$$E_R \left[\left(\left(k_i - \frac{\partial S}{\partial \beta_p} (\beta_p, \eta) \beta_i \right) \left(\frac{1}{c} \right) \right) \right] \left(\quad (106)$$

which leads to an expression for the second derivative as

$$E_R \left[\left(1 + \frac{1}{c} \frac{\partial^2 S}{\partial \beta_p^2}(\beta_p, \eta) \beta_i^2 \right) \left(\frac{1}{c} \right) \right] \left(\right. \quad (107)$$

which is negative since $\frac{\partial^2 S}{\partial \beta_p^2}(\beta_p, \eta) \geq 0$.

C.6 Non-linear tax

I begin by assuming that the regulator has selected equilibrium $k_i(q_i)$ functions that are twice-differentiable, then solve for the bank's optimal portfolio choice. I then show that given the bank's choice, the regulator will optimally select twice-differentiable $k_i(q_i)$ functions.

Taking a twice-differentiable $k_i(q_i)$ as given, the bank's first-order condition for q_i is

$$q_i = \frac{1}{c}(a_i - \gamma \beta_i - k'_i(q_i)). \quad (108)$$

The second-order condition requires that $-c - k''_i(q_i) < 0$, which is not necessarily guaranteed since the $k_i(q)$ functions can be arbitrary. However, I will later verify that the equilibrium $k_i(q_i)$ functions for the regulator lead to the second-order conditions being satisfied.

Taking a first-order condition with respect to $k'_i(q_i)$ leads to

$$E[(a_i - cq_i - (\gamma + \eta)\beta_i) \underbrace{\frac{\partial q_i}{\partial k'_i(q_i)}}_{=-\frac{1}{c}} | q_i, \hat{a}_i, \hat{\beta}_i] = 0 \quad (109)$$

$$\implies E \left[a_i - b \left(\frac{1}{c}(a_i - \gamma \beta_i - k'_i(q_i)) \right) \right] \left((\gamma + \eta)\beta_i | q_i, \hat{a}_i, \hat{\beta}_i \right) \left(\right. \quad (110)$$

$$\implies k'_i(q_i) = \eta E[\beta_i | q_i, \hat{a}_i, \hat{\beta}_i]. \quad (111)$$

Next, I focus on constructing an explicit expression for the expectation of beta. Note that rewriting the bank's optimal quantity reveals a noisy signal of β_i and a_i as

$$z_i = -\frac{1}{\gamma}(cq_i + k'_i(q_i)) = \beta_i - \frac{1}{\gamma}a_i. \quad (112)$$

Based on the assumption of joint normality and the linearity of the signals, the regulator's optimal estimate of β_i will be some linear combination of $\hat{a}_i$, $\hat{\beta}_i$, and z_i . The coefficients on the linear combination will match the coefficients of a regression of β_i on $\hat{a}_i$, $\hat{\beta}_i$, and

z_i , which leads to

$$E[\beta_i|q_i, \hat{a}_i, \hat{\beta}_i] = \lambda_0 + \lambda_a \hat{a}_i + \lambda_b \hat{\beta}_i + \lambda_z z_i. \quad (113)$$

The λ coefficients other than the constant come from $Var(\hat{x})^{-1}Cov(\beta_i, \hat{x})$, where $\hat{x} = [\hat{a}_i, \hat{\beta}_i, z_i]$ and $\lambda_0 = \bar{\beta} - \lambda_a \bar{a} - \lambda_b \bar{\beta} - \lambda_z (\bar{\beta} - \frac{1}{\gamma} \bar{a})$. Next, apply the relationship

$$k'_i(q_i) = \eta E[\beta_i|q_i, \hat{a}_i, \hat{\beta}_i] \quad (114)$$

$$= \eta(\lambda_0 + \lambda_a \hat{a}_i + \lambda_b \hat{\beta}_i + \lambda_z (-\frac{1}{\gamma}(cq_i + k'_i(q_i)))). \quad (115)$$

Solving for $k'_i(q_i)$ yields

$$k'_i(q_i) = \frac{\eta\lambda_0}{1 + (\eta/\gamma)\lambda_z} + \frac{\eta\lambda_a}{1 + (\eta/\gamma)\lambda_z} \hat{a}_i + \frac{\eta\lambda_b}{1 + (\eta/\gamma)\lambda_z} \hat{\beta}_i - \frac{(\eta/\gamma)\lambda_z}{1 + (\eta/\gamma)\lambda_z} cq_i \quad (116)$$

Integrating and setting $k_i(0) = 0$ yields

$$k_i(q_i) = \left(\frac{\eta\lambda_0}{1 + (\eta/\gamma)\lambda_z} + \frac{\eta\lambda_a}{1 + (\eta/\gamma)\lambda_z} \hat{a}_i + \frac{\eta\lambda_b}{1 + (\eta/\gamma)\lambda_z} \hat{\beta}_i \right) \left(q_i - \frac{(\eta/\gamma)\lambda_z}{1 + (\eta/\gamma)\lambda_z} \frac{c}{2} q_i^2 \right) \quad (117)$$

Returning to the second-order condition from earlier, I plug in the result

$$k''_i(q_i) = -c \frac{(\eta/\gamma)\lambda_z}{1 + (\eta/\gamma)\lambda_z} \quad (118)$$

to obtain that the second derivative of the bank's objective with respect to q_i is

$$-c \left(1 - \frac{(\eta/\gamma)\lambda_z}{1 + (\eta/\gamma)\lambda_z} \right) \left(\quad (119)$$

$$= -c \left(\frac{1}{1 + (\eta/\gamma)\lambda_z} \right) \left(\quad (120)$$

which is negative if and only if

$$\lambda_z > -\frac{\gamma}{\eta}. \quad (121)$$

[TO DO: Investigate whether this condition holds. I have not found any numerical counterexamples so far, but I have seen that λ_z is negative when $\gamma = 2$, $\eta = 2$, $\sigma_{ub} = 4$, $\sigma_{ua} = 7$, $\sigma_{ea} = 3$, $\sigma_{eb} = 1$, and $\rho = -0.5$.]

C.7 Fully undisclosed linear tax

To verify that these pair of choices are a solution, I first suppose that the regulator sets $k_i = \eta \hat{\beta}_i$. The bank's first-order condition for q_i leads to

$$q_i = \frac{1}{c}(a_i - \gamma\beta_i - \eta E_B[\hat{\beta}_i]) \quad (122)$$

$$= \frac{1}{c}(a_i - \gamma\beta_i - \eta E[\bar{\beta} + u_i^b + e_i^b | u_i^b]) \quad (123)$$

$$= \frac{1}{c}(a_i - \gamma\beta_i - \eta \underbrace{(\bar{\beta} + u_i^b)}_{\beta_i}) \quad (124)$$

which is the desired solution. Since this quantity matches the first-best quantity from Proposition 1, it also maximizes the regulator's objective function.

Note that this solution is not unique. For all random variables x_i satisfying $E_B[x_i] = \beta_i$, the regulator will achieve first-best by setting $k_i = \eta x_i$. While $\hat{\beta}_i$ is an obvious candidate, $\hat{\beta}_i$ with added noise or the fitted value from regressing β_i on $\hat{a}_i$ are other possibilities.

C.8 Partially-disclosed linear tax

First, I'll establish the optimality condition for k_i . Starting with the bank's problem, the bank maximizes given knowledge over the distribution of k_i . The bank solves

$$\max_{\{q_i\}} \int_0^1 (a_i q_i - \frac{c}{2} q_i^2 - \gamma \beta_i q_i - E_B[k_i] q_i) di, \quad (125)$$

which leads to the first-order condition

$$q_i = \frac{1}{c}(a_i - \gamma\beta_i - E_B[k_i]). \quad (126)$$

The second-order condition is satisfied since $-c < 0$, which ensures that the solution is a maximum.

Given this choice of q_i for the bank, the regulator solves

$$\max_{\{k_i\}} E_R \left[\int_0^1 (a_i q_i - \frac{c}{2} q_i^2 - (\gamma + \eta) \beta_i q_i) di \right] \quad (127)$$

which leads to a first-order condition of

$$E_R \left[\begin{pmatrix} q_i - cq_i - (\gamma + \eta)\beta_i & \frac{\partial q_i}{\partial k_i} \\ \underbrace{\phantom{\frac{\partial q_i}{\partial k_i}}}_{=-(1/c)} \end{pmatrix} \right] \begin{pmatrix} \\ \end{pmatrix} \neq 0 \quad (128)$$

$$\implies E_R[E_B[k_i]] = \eta E_R[\beta_i] \quad (129)$$

as desired.

Next, I show that the proposed solution satisfies this condition. I conjecture that the regulator sets a linear tax according to

$$k_i = \eta(\bar{\beta} + m'(\hat{x}_i - \bar{x})), \quad (130)$$

where $\hat{x}_i = [\hat{a}_i, \hat{\beta}_i]$ and m is a 2×1 vector. I then solve for the m that satisfies the previous first-order condition.

First, I begin by computing the left-hand side of the first-order condition, $E_R[E_B[k_i]]$. Start by taking the expectation of the tax with respect to the bank's information set as

$$E_B[k_i] = \eta(\bar{\beta} + m'E_B[\hat{x}_i - \bar{x}]) \quad (131)$$

$$= \eta(\bar{\beta} + m'[x_i + E_B[e_i] - \bar{x}]), \quad (132)$$

where $e_i = [e_i^a, e_i^b]$ is the vector of the error in the regulator's signals. Recall that the bank receives a noisy signal $s_i = \hat{x}_i + w_i$. Since the bank perfectly observes x_i , then $s_i - x_i = e_i + w_i$ is a noisy signal of e_i . $E_B[e_i]$ emerges from the result of regressing e_i on $s_i - x_i$, which is

$$E_B[e_i] = [Var(s_i - x_i)^{-1}Cov(s_i - x_i, e_i)]'(s_i - x_i) \quad (133)$$

$$= \underbrace{[Var(e_i + w_i)^{-1}Cov(e_i + w_i, e_i)]'}_{\Lambda}(e_i + w_i). \quad (134)$$

Thus, the bank's expectation of the tax is

$$E_B[k_i] = \eta(\bar{\beta} + m'[x_i + \Lambda'(e_i + w_i) - \bar{x}]). \quad (135)$$

Now take the regulator's expectation of this quantity. The key is $E_R[x_i + \Lambda'(e_i + w_i) - \bar{x}]$, which can be computed based on a linear regression on $\hat{x}_i$ given the assumptions of

normality. The expectation is then

$$E_R[x_i + \Lambda'(e_i + w_i) - \bar{x}] = [Var(\hat{x}_i)^{-1}Cov(\hat{x}_i, x_i + \Lambda'(e_i + w_i))]'(\hat{x}_i - \bar{x}) \quad (136)$$

$$= [Var(\hat{x}_i)^{-1}(Var(x_i) + Var(e_i)\Lambda)]'(\hat{x}_i - \bar{x}), \quad (137)$$

which leads to an estimated tax of

$$E_R[E_B[k_i]] = \eta(\bar{\beta} + m'[Var(\hat{x}_i)^{-1}(Var(x_i) + Var(e_i)\Lambda)]'(\hat{x}_i - \bar{x})). \quad (138)$$

The right-hand side of the first-order condition is

$$\eta E_R[\beta_i] = \eta(\bar{\beta} + [Var(\hat{x}_i)^{-1}Cov(\hat{x}_i, \beta_i)]'(\hat{x}_i - \bar{x})) \quad (139)$$

$$= \eta(\bar{\beta} + [Var(\hat{x}_i)^{-1}Cov(x_i, \beta_i)]'(\hat{x}_i - \bar{x})). \quad (140)$$

Substituting in for the left- and right-hand sides of the first-order condition yields

$$\underbrace{\eta(\bar{\beta} + m'[Var(\hat{x}_i)^{-1}(Var(x_i) + Var(e_i)\Lambda)]'(\hat{x}_i - \bar{x}))}_{E_R[E_B[k_i]]} = \underbrace{\eta(\bar{\beta} + [Var(\hat{x}_i)^{-1}Cov(x_i, \beta_i)]'(\hat{x}_i - \bar{x}))}_{\eta E_R[\beta_i]} \quad (141)$$

which can be written as

$$m'[Var(\hat{x}_i)^{-1}(Var(x_i) + Var(e_i)\Lambda)]' = [Var(\hat{x}_i)^{-1}Cov(x_i, \beta_i)]'. \quad (142)$$

Solving for m yields

$$m = (Var(x_i) + Var(e_i)\Lambda)^{-1}Cov(x_i, \beta_i) \quad (143)$$

$$= (Var(x_i) + Var(e_i)Var(e_i + w_i)^{-1}Var(e_i))^{-1}Cov(x_i, \beta_i) \quad (144)$$

Finally, I show that m is the coefficient from regressing β_i on $E_B[\hat{x}_i]$. To see this, first note that

$$E_B[\hat{x}_i] = E_B[x_i + e_i] \quad (145)$$

$$= x_i + E_B[e_i] \quad (146)$$

$$= x_i + \Lambda'(e_i + w_i). \quad (147)$$

Since x_i is independent of e_i and w_i , it follows that

$$\text{Var}(E_B[\hat{x}_i]) = \text{Var}(x_i + \Lambda'(e_i + w_i)) \quad (148)$$

$$= \text{Var}(x_i) + \Lambda' \text{Var}(e_i + w_i) \Lambda. \quad (149)$$

Substituting in

$$\Lambda = \text{Var}(e_i + w_i)^{-1} \text{Cov}(e_i + w_i, e_i) \quad (150)$$

$$= \text{Var}(e_i + w_i)^{-1} \text{Var}(e_i) \quad (151)$$

and canceling out the resulting $\text{Var}(e_i + w_i) \text{Var}(e_i + w_i)^{-1}$ term yields

$$\text{Var}(E_B[\hat{x}_i]) = \text{Var}(x_i) + \text{Var}(e_i) \text{Var}(e_i + w_i)^{-1} \text{Var}(e_i). \quad (152)$$

Based on this result, it follows that m can be expressed as

$$m = \text{Var}(E_B[\hat{x}_i])^{-1} \text{Cov}(x_i, \beta_i) \quad (153)$$

$$= \text{Var}(E_B[\hat{x}_i])^{-1} \text{Cov}(x_i + E_B[e_i], \beta_i) \quad (154)$$

$$= \text{Var}(E_B[\hat{x}_i])^{-1} \text{Cov}(E_B[\hat{x}_i], \beta_i), \quad (155)$$

which is the coefficient from a linear regression of β_i on $E_B[\hat{x}_i]$.

C.9 Non-disclosure with an aversion to idiosyncratic volatility

If the regulator sets the tax in this manner, then the bank's objective function transforms into the regulator's, which immediately leads to the first best. Setting the two objectives equal, the bank and regulator have perfectly-aligned incentives when

$$E_B[k]q + \gamma_I \sigma_B(k)|q| = \eta \beta q + \eta_I \sigma_B(k)|q| \quad (156)$$

$$\iff E_B[k] + \gamma_I \sigma_B(k) \text{sgn}(q) = \eta \beta + \eta_I \sigma_B(k) \text{sgn}(q). \quad (157)$$

Recall that $\hat{\beta} = \beta + e$ and that the bank knows β , but has no information about e . Therefore $E_B[\hat{\beta}] = \beta$ and $\sigma_B(\hat{\beta}) = \sigma_e$. Using this fact yields

$$E_B[k] = (\eta_I - \gamma_I) \eta \sigma_e \text{sgn}(q) + \eta \beta \quad (158)$$

$$\sigma_B(k) = \eta \sigma_e. \quad (159)$$

Therefore it follows that the equation

$$E_B[k] + \gamma_I \sigma_B(k) \text{sgn}(q) = \eta\beta + \eta_I \sigma_B(k) \text{sgn}(q) \quad (160)$$

$$(\eta_I - \gamma_I) \eta \sigma_e \text{sgn}(q) + \eta\beta + \gamma_I \eta \sigma_e \text{sgn}(q) = \eta\beta + \eta_I \eta \sigma_e \text{sgn}(q) \quad (161)$$

$$\eta_I \eta \sigma_e \text{sgn}(q) + \eta\beta = \eta_I \eta \sigma_e \text{sgn}(q) + \eta\beta \quad (162)$$

is satisfied and thus the regulator's and bank's incentives are aligned.

C.10 Equivalence between taxes and capital requirements with exogenous capital and idiosyncratic shocks

If the regulator has set a capital requirement, then the bank solves

$$\max_{\{q_i\}} \int_0^1 (a_i q_i - \frac{c}{2} q_i^2 - \gamma \beta_i q_i) di \quad (163)$$

subject to the constraint that

$$\int_0^1 m_i q_i di \leq 1. \quad (164)$$

Setting up the Lagrangian, the problem becomes

$$\max_{\{q_i\}} \int_0^1 (a_i q_i - \frac{c}{2} q_i^2 - \gamma \beta_i q_i) di - \lambda \left(\int_0^1 m_i q_i di - 1 \right) \quad (165)$$

which leads to the first-order condition for q_i of

$$a_i - c q_i - \gamma \beta_i - \lambda m_i = 0 \quad (166)$$

$$\implies q_i = \frac{1}{c} (a_i - \gamma \beta_i - \lambda m_i). \quad (167)$$

Compare this first-order condition to the case in which the regulator sets an asset-specific tax k_i , in which case the bank selects

$$q_i = \frac{1}{c} (a_i - \gamma \beta_i - k_i). \quad (168)$$

The regulator can replicate the outcome from the capital requirement by setting a tax $k_i = \lambda m_i$. However, this replication is possible in this case because λ is deterministic, which I will show next.

If the capital requirement is not binding, then $\lambda = 0$. Otherwise, multiplying the

first-order condition for q_i by m_i and then integrating over i yields

$$\int_0^1 m_i q_i di = \frac{1}{c} \int_0^1 (m_i(a_i - \gamma\beta_i) - \lambda m_i^2) di = 1 \quad (169)$$

In the case in which there is only idiosyncratic uncertainty in a_i and β_i , the integral almost certainly equals the expectation, so that

$$\frac{1}{c}(E[m_i(a_i - \gamma\beta_i)] - \lambda E[m_i^2]) = 1 \quad (170)$$

$$\implies \lambda = \frac{E[m_i(a_i - \gamma\beta_i)] - c}{E[m_i^2]}, \quad (171)$$

which is deterministic. Intuitively, the idiosyncratic shocks to a_i and β_i diversify away in the aggregate so that the regulator understands how tightly the capital requirements will bind overall.

C.11 Linear taxes with endogenous capital structure

I will show that the proposed solution of $s = r_p - r_s$ and $k_i = \eta E_R[\beta_i]$ satisfies the conditions for optimal s and k_i .

I start with the bank's problem, taking s and k_i as given. The bank solves

$$\max_{\{q_i\}, e} \int_0^1 \left((a_i - \gamma\beta_i)q_i - \frac{c}{2}q_i^2 - \frac{k_i q_i}{e} \right) di - (r_p - s)e. \quad (172)$$

The first-order conditions yield

$$q_i = \frac{1}{c} \left(a_i - \gamma\beta_i - \frac{k_i}{e} \right) \quad (173)$$

$$e = \sqrt{\frac{k_p}{r_p - s}}. \quad (174)$$

The second derivatives are

$$\text{Second of } q_i : -c \quad (175)$$

$$\text{Second of } e : -2 \frac{k_p}{e^3}. \quad (176)$$

At $e = \sqrt{\frac{k_p}{r_p - s}}$ and $s = r_p - r_s$, the second derivative with respect to e is

$$e = -2 \frac{k_p}{e^2} \frac{1}{e} = -2 \frac{r_s}{e}, \quad (177)$$

which is negative since both $r_s > 0$ and $e > 0$. Thus the second-order conditions for a maximum are satisfied.

Next, switch to the regulator's problem. The regulator solves

$$\max_{\{k_i\}, s} E_R \left[\int_0^1 \left((a_i - \gamma\beta_i)q_i - \frac{c}{2}q_i^2 - \eta \frac{\beta_i q_i}{e} \right) di - r_s e \right] \left(\quad \right) \quad (178)$$

where q_i and e follow the bank's strategy from before. The first-order conditions with respect to k_i and s yield

$$E_R \left[\int_0^1 \left(a_i - \gamma\beta_i - cq_i - \eta \frac{\beta_i}{e} \right) \left(\frac{\partial q_i}{\partial k_i} di + \eta \frac{\int_0^1 \beta_i q_i di}{e^2} - r_s \right) \frac{\partial e}{\partial k_i} \right] = 0. \quad (179)$$

Substitute the expressions for cq_i and for e^2 to obtain

$$E_R \left[\int_0^1 \left(\frac{k_i - \eta\beta_i}{e} \right) \left(\frac{\partial q_i}{\partial k_i} di + \eta \frac{\int_0^1 \beta_i q_i di}{\int_0^1 k_i q_i di} (r_p - s) - r_s \right) \frac{\partial e}{\partial k_i} \right] = 0. \quad (180)$$

To verify that this first-order condition is satisfied for $s = r_p - r_s$ and $k_i = \eta E_R[\beta_i]$, substitute in these values and apply the law of iterated expectations to obtain

$$E_R \left[\left(E \left[\int_0^1 \left(\frac{\eta E_R[\beta_i] - \eta\beta_i}{e} \right) \left(\frac{\partial q_i}{\partial k_i} di + \frac{r_s}{\int_0^1 k_i q_i di} \int_0^1 (\eta\beta_i - \eta E_R[\beta_i]) q_i \frac{\partial e}{\partial k_i} di \right) e, \{q_i\}, \{\hat{\beta}_i\}, \{\hat{a}_i\} \right] \right) \right] = 0, \quad (181)$$

which follows since $E_R[\beta_i - E_R[\beta]] = 0$. Showing that the first-order condition with respect to s is satisfied is identical, except that $\frac{\partial q_i}{\partial k_i}$ and $\frac{\partial e}{\partial k_i}$ are replaced with $\frac{\partial q_i}{\partial s}$ and $\frac{\partial e}{\partial s}$.