

A Pigovian Approach to Liquidity Regulation*

Enrico Perotti
U. Amsterdam, DSF and CEPR

Javier Suarez
CEMFI and CEPR

First version: October 2010
This version: April 2011

Abstract

This paper discusses liquidity regulation when short-term funding enables credit growth but generates negative systemic risk externalities. It focuses on the relative merit of price versus quantity rules, showing how they target different incentives for risk creation.

When banks differ in credit opportunities, a Pigovian tax on short-term funding is efficient in containing risk and preserving credit quality, while quantity-based funding ratios are distortionary. Liquidity buffers are either fully ineffective or similar to a Pigovian tax with deadweight costs. Critically, they may be least binding when excess credit incentives are strongest.

When banks differ instead mostly in gambling incentives (due to low charter value or overconfidence), excess credit and liquidity risk are best controlled with net funding ratios. Taxes on short-term funding emerge again as efficient when capital or liquidity ratios keep risk shifting incentives under control. In general, an optimal policy should involve both types of tools.

Keywords: Systemic risk, Liquidity risk, Liquidity requirements, Liquidity risk levies, Macroprudential regulation.

JEL Classification: G21, G28.

*We have greatly benefited from numerous discussions and comments from academics and policymakers on our policy writings on the regulation of liquidity and on this paper. Our special thanks to Rafael Repullo for his insightful suggestions for the revision of the paper. Contact e-mails: E.C.Perotti@uva.nl, suarez@cemfi.es.

1 Introduction

The recent crisis has provided a clear rationale for the regulation of banks' refinancing risk, a critical gap in the Basel II framework. This paper studies the effectiveness of different approaches to liquidity regulation.

The basic trade off of short-term funding is that rapid expansion of credit may only be funded by attracting short-term funding (for instance, because deposit supply can be expanded only slowly, or because short term market lenders do not need to be very informed about new credit choices), but this creates refinancing risk, especially in situations where there are doubts on banks' fundamental solvency or evidence of excessive risk-taking on the asset side as well.¹ Sudden withdrawals may lead to disruptive liquidity runs (Diamond and Dybvig, 1983), and cause fire sales or counterparty risk externalities which affect other intermediaries exposed to short term funding (Brunnermeier, 2009; Allen, Babus and Carletti, 2010). As a result, each bank's funding decision has an impact on the vulnerability of other banks to liquidity risk, causing a negative externality. Even if individual banks' funding decision takes into account its own exposure to refinancing risk, it will not internalize its system-wide effect (Perotti and Suarez, 2009). Because of the wedge between the net private value of short-term funding and its social cost, banks will rely excessively on short-term funding. A prime example is the massive build up in wholesale funding which supported the recent securitization wave, and the overnight credit (repo) growth during 2002-2007, which grew explosively to a volume over ten trillion dollars (Gorton, 2009). Rapid withdrawals forced an unprecedented liquidity support by central banks, undermining their control over the money supply.

In the tradition of externality regulation led by Weitzman (1974), we assess the performance of Pigovian taxes (aimed at equating private and social liquidity costs) and quantity regulations in containing this systemic externality. As in Weitzman, the optimal regulatory tool depends on the response elasticity of banks, recognizing that the regulator is informa-

¹In this paper, we adopt the same strategy as the Basel Committee on Bank Supervision (BCBS) and, for simplicity, discuss liquidity risk as if it were distinct and separable from other sources of risk. Arguably, liquidity risk and solvency risk are hardly separable, both in practice, as it is hard to tell apart insolvent banks from those with pure liquidity problems, and conceptually, since solvency concerns are at the root of most liquidity crises.

tionally constrained in targeting individual bank characteristics.² Our results show how the industry response to regulation depends on the composition of bank characteristics. The model recognizes that banks differ in their credit ability and their incentives to take risk. Banks earn decreasing returns to expand credit to their (monitored) borrowers, so better banks naturally lend more. Shareholders of less capitalized banks gain from investing in poor gambles, since they retain the upside and shift downside risk to the public safety net.³ Depending on the dominant source of heterogeneity, the socially efficient solution may be attained with Pigovian taxes, quantity regulations or a combination of both.

To facilitate the discussion, we first analyze the impact of regulation under either investment quality or risk-taking heterogeneity.⁴ When banks differ only in capacity to lend profitably (reflecting credit assessment capability or access to credit opportunities), a simple flat-rate Pigovian tax on short-term funding (possibly scaled up by the systemic importance of each bank) implements the efficient allocation. The intuition is that liquidity risk levies allow better banks to lend more, without requiring the regulators to be able to identify them. In this context, quantity-based instruments such as the net stable funding ratio or the liquidity coverage ratio introduced by Basel III (see BCBS, 2010) may improve over the unregulated equilibrium but are generally distortionary. An optimal quantity-based approach would require introducing contingency on individual bank characteristics, most of which are imprecisely measured or unobservable.⁵

More precisely, net stable funding ratios which impose an upper limit on short-term debt do reduce overall liquidity risk, but redistribute liquidity risk inefficiently across banks. Banks with better credit opportunities will be constrained, while the reduced systemic risk actually encourages banks with low credit ability (for which the requirement is not binding)

²Our analysis is also related to the classical discussion by Poole (1970) on the optimality of price or quantity monetary policy instruments when the system to regulate is affected by several types of shocks.

³An alternative view of gambling incentives is that it is driven by self-interested and overconfident managers, which view excessive risks as profitable

⁴Each form of heterogeneity leads to a situation akin to each of the polar cases that Weitzman (1974, p. 485) describes in terms of the “curvature” of the social benefit function and the private cost function relevant to his analysis—he finds that price (quantity) regulation dominates when the social benefit (private cost) function is linear.

⁵Quantity requirements may be indexed to measures of the “systemic importance” of each bank such size, interconnectedness, capitalization, etc., but it is much harder to find good proxies of banks’ credit opportunities.

to expand.

Liquidity coverage ratios which require banks to hold fractional reserves of liquid assets against short-term funding work as a de facto tax, but one that can turn out to be ineffective or imply large deadweight losses.⁶ When the yield on liquid assets equals the cost of short-term liabilities (roughly the case in normal times, and certainly prior to the crisis), buffers impose no net cost to stacking liquidity. Banks will simply increase their gross short-term funding to keep their “net” short-term funding (i.e. minus the buffers) as high as in the unregulated equilibrium. The only effect is an artificial demand for liquid assets—traditionally kept in money market mutual funds rather than banks—that might be redirected to banks following the new requirement.

When the spread between bank borrowing costs and liquid asset yields is positive, a liquidity requirement operates as a tax on short-term funding that involves deadweight costs and has its effective tax rate partly determined by market spreads.⁷ In the recent experience, the interbank spread over safe assets was minimal just as aggregate liquidity risk was building up, and jumped up with the start of the crisis. This means that liquidity requirements would have to be increased in good times and reduced in bad times so as to avoid making them a source of further banking system procyclicality.

Studying variation in risk-shifting incentives (correlated with charter value or other determinant of risk-taking tendencies, such as overconfidence) alters the results radically. Low charter value (or more risk loving) banks have strong incentives to gamble to shift risk to the deposit insurance provider (Keeley, 1990). We show that decisions driven by such gambling incentives are not properly deterred by levies, while quantity constraints are more effective. Both short-term funding limits (e.g. a net stable funding ratio) and capital requirements can contain risk shifting by limiting the scale of lending. Levies will not be as effective as in the previous scenario because the most gambling-inclined banks will also be the most inclined to pay the tax and expand their risky lending. In this case, quantity instruments such as net funding or capital ratios are best to contain excess credit expansion.

⁶Liquid assets which can be sold at no fire-sale loss in a crisis are essentially cash, central bank reserves, and treasury bills.

⁷Specifically, the tax rate will equal the product of the buffer requirement per unit of short-term funding times the interest rate spread.

Our analysis identifies the relative merits of price versus quantity instruments, and suggests that combining them may be adequate for the simultaneous control of gambling incentives and systemic risk externalities. However, this presumes that the regulator controls only instruments connected to liquidity risk. If strengthening capital requirements is an effective strategy for the control of gambling incentives (Hellmann, Murdock, and Stiglitz, 2000; Repullo, 2004), the case for levies on short-term funding is considerably reinforced.⁸

Other considerations may qualify the recommendation for the use of price-based or quantity-based instruments. For instance, levies may be less costly to adjust than ratios. First, they might be easier to change for institutional reasons (e.g. if regulatory ratios are embedded in some law or international agreement while the levies are, at least partly, under control of a macroprudential authority). More importantly, they may imply lower adjustment costs at bank level than changing bank funding volumes on short notice. Similarly, changes in levies are less likely to induce procyclicality, since the Pigovian “tax rate” is directly controlled by the regulator rather than implicitly set by the interaction of some (controlled) quantitative requirement and the (freely fluctuating) market price of the required resource (namely, capital, liquid assets or stable funding). For preventive policy, controlling time varying liquidity risk may then be best achieved by a combination of stable ratios and variable levies.

The rest of the paper is organized as follows. Section 2 describes some related literature and some recent evidence on liquidity risk. Section 3 describes the baseline model. Section 4 characterizes the unregulated equilibrium. Section 5 finds the socially optimal allocation. In Section 6, we discuss the possibility of restoring efficiency with a Pigovian tax on short-term funding. Section 7 considers alternative quantity-based regulations. In Section 8 we analyze the implications of introducing gambling incentives as a second dimension of bank heterogeneity. Section 9 discusses further implications and extensions of the analysis. Section 10 concludes the paper.

⁸Hellmann, Murdock, and Stiglitz (2000) point out that “capital requirements also have a perverse effect of harming banks’ franchise values, thus encouraging gambling” but the analysis in Repullo (2004) suggests that the standard deleveraging effect tends to dominate the charter value effect.

2 Evidence from the crisis and related research

The crisis of 2007-2008 has been described as a wholesale bank crisis, or a repo run crisis (Gorton, 2009). The rapid withdrawing of short-term debt was responsible for propagation of shocks across investors and markets (Brunnermeier, 2009). Brunnermeier and Oehmke (2010) show that creditors have an incentive to shorten their loan maturity, so as to pull out in bad times before other creditors can. This, in turn, causes a lender race to shorten maturity, leading to excessively short-term financing. The consequences are formalized in Martin, Skeie and von Thadden (2010), where increased collective reliance on repo funding weakens solvency constraints, and produce repo runs. Acharya and Viswanathan (2011) model the sudden drying up of liquidity when banks need to refinance short-term debt in bad times. As low asset prices increase incentives for risk shifting, investors may rationally refuse refinancing to illiquid banks.

Acharya, Gale, and Yorulmazer (2010) show that high roll-over frequency can reduce the collateral value of risky securities, but they treat debt maturity as exogenous and do not look at the normative implications. Papers emphasizing the possibility of socially inefficient levels of maturity transformation include Huang and Ratnovski (2011), who focus on the deterioration of information production incentives among banks, Farhi and Tirole (2010), where the distortion comes from the expectation of a bail-out, and Segura and Suarez (2011), where the pricing of refinancing during crisis interacts with banks' financial constraints and gives raise to pecuniary externalities linked to banks' funding maturity decisions.

While the role of liquidity risk in the crisis has been evident from the beginning, more precise empirical evidence is now emerging. Acharya and Merrouche (2009) show that UK banks with more wholesale funding and fire sale losses in 2007-08 contributed more to the transmission of shocks to the interbank market. A concrete measure of the role of short-term debt played in the credit boom and its demise comes from the explosive rise of repo (overnight) financing in the last years, and its rapid deflation since the panic (Gorton, 2009). Repo funding evaporated in the crisis, leading to bursts of front running in the sales of repossessed securities. Adrian and Brunnermeier (2009) present evidence of the correlation between banks' use of short-term wholesale funding and their proposed measure of banks' contribution to systemic risk (CoVaR). A similar result emerges in Acharya et al (2010).

The leading causes of external effects from refinancing risk have been identified as losses due to fire sales, and collective fears about counterparty risk amplified by simultaneous refinancing choices. They have motivated proposals on the creation of private or public clearing arrangements to limit the effects of runs, though purely private arrangements are not expected to be sufficient in systemic liquidity runs. Acharya and Öncü (2010) argue for the establishment of a Repo Resolution Authority to take over repo positions in a systemic event, paying out a fraction of their claims and liquidating the collateral in an orderly fashion. This would force investors to bear any residual loss. On the opposite front, Gorton (2009) has proposed stopping fire sales of seized collateral by a blanket state guarantee, while Gorton and Metrick (2010) propose creating special vehicles they call narrow banks to hold such assets, backed by a public guarantee.

Another critical issue is the consequences of ex post liquidity bailouts (Farhi and Tirole, 2010). If in a systemic run there is no choice but to provide liquidity to mismatched intermediaries, this implies a loss of public control over the money supply, which becomes endogenous to the private sector's short-term funding decisions. This highlights the urgency of measures to contain the private creation of liquidity risk. Finally, systemic crises are the source of important fiscal and real losses not fully internalized by those who make the decisions that lead to the accumulation of systemic liquidity risk (Laeven and Valencia, 2010), making a clear case for regulation.

The paper is related to several other strands of the academic literature which would take too long to revise in a systematic manner. These include the corporate finance and banking literatures on the potentially beneficial incentive effects of short-term funding (e.g. Calomiris and Kahn, 1991, Diamond and Rajan, 2001, and Huberman and Repullo, 2010), on the connection between short-term funding and banks' vulnerability to panics and contagion (e.g. Allen and Gale, 2000, Rochet and Vives, 2004, and Allen, Babus, and Carletti, 2010), and on externalities related to other financial decisions, such as diversification decisions (Wagner, 2010) or decisions regarding the supply of credit over the business cycle (Lorenzoni, 2008, and Jeanne and Korinek, 2010). Finally, our analysis is also connected to a vast economic literature about the choice between quantity-based and price-based regulation in specific

setups.⁹

3 The model

Consider a one-period model of a banking economy in which all agents are risk neutral. The banking system is made up of a continuum of heterogeneous banks run by their owners with the objective of maximizing their expected net present value (NPV). To start with, we assume that banks differ in a parameter θ that affects the NPV that they can generate using short-term funding, whose amount will be their only decision variable for the time being.¹⁰ The parameter θ follows a continuous distribution with positive density $f(\theta)$ over the interval $[0, 1]$. Assuming w.l.o.g. that all banks of each class θ behave symmetrically, the short-term funding decision of each bank of class θ is denoted by $x(\theta) \in [0, \infty)$.

We postulate that the expected NPV associated with a decision x by a bank of class θ can be written as

$$v(x, X, \theta) = \pi(x, \theta) - \varepsilon(x, \theta)c(X), \quad (1)$$

where X is a measure of the aggregate systemic risk implied by the individual funding decisions of all banks, $\pi(x, \theta)$ is the NPV generated if no systemic liquidity crisis occurs, and $\varepsilon(x, \theta)c(X)$ is the expected NPV loss due to the possibility of a systemic liquidity crisis. To facilitate the presentation, we assume a multiplicative decomposition of the expected crisis losses in two terms: the term $\varepsilon(x, \theta) \geq 0$, which captures the purely individual contribution of the funding decision x and the individual characteristic θ to the vulnerability of the bank, and the term $c(X) \geq 0$, which captures the influence of other banks' funding decisions on systemic crisis costs.

We assume that $\pi(x, \theta)$ is increasing and differentiable in its two arguments, strictly concave in x , and with a positive cross derivative, $\pi_{x\theta} > 0$, so that a larger θ implies a larger capability to extract value from short-term funding. To facilitate obtaining interior solutions in x and monotone comparative statics with respect to θ , we also assume that

⁹See contributions such as Glaeser and Shleifer (2001) and Kaplow and Shavell (2002) for an overview of the literature.

¹⁰In Section 8, we introduce a second dimension of bank heterogeneity directed to capture differences in banks' gambling incentives.

$\varepsilon(x, \theta)$ is differentiable, increasing and weakly convex in x , non-increasing in θ , and with $\varepsilon_{x\theta} \leq 0$. Finally, we assume $c(X)$ to be increasing, differentiable, and weakly convex in X .

A structural story consistent with this specification might be that $\pi(x, \theta)$ captures the profitability, in the absence of a systemic liquidity crisis, of using short-term funding to expand lending, $\varepsilon(x, \theta)$ captures the probability that the bank faces refinancing problems in a liquidity crisis and has to accommodate them by, say, selling its assets, and $c(X)$ are the net liquidation losses incurred in such an event. Notice that $c(X)$ might be increasing in X due to the impact on liquidation values of concurrent sales from troubled banks (e.g. under some cash-in-the-market pricing logic or simply because the alternative users of the liquidated assets face marginally decreasing returns).¹¹ Here θ can be taken as a measure of a bank's credit ability or any other determinant of the marginal net value of its investments.

The key results below would be robust to essentially any specification of the aggregator $X = g(\{x(\theta)\})$, where $\{x(\theta)\}$ is the schedule of the short-term funding used by the banks in each class $\theta \in [0, 1]$ and we have $\partial g / \partial x(\theta) \geq 0$ for all θ . For concreteness, however, we focus on the case in which aggregate systemic liquidity risk can be measured as the simple sum of all individual decisions:

$$X = g(\{x(\theta)\}) = \int_0^1 x(\theta) f(\theta) d\theta. \quad (2)$$

In Section 9, we will discuss how to adapt our main results to the case in which banks also differ in a “systemic importance” factor that affects the weight of the contribution of their short-term funding to X .

We assume that all investors, except bank owners, have the opportunity to invest their wealth at exogenously given market rates and provide funding at competitive terms, hence obtaining a zero NPV from dealing with the banks. Then, the total NPV generated by banks (and appropriated by their owners) constitutes the natural measure of social welfare W in this economy. Formally,

$$W(\{x(\theta)\}) = \int_0^1 v(x(\theta), X, \theta) f(\theta) d\theta = \int_0^1 [\pi(x(\theta), \theta) - \varepsilon(x(\theta), \theta) c(X)] f(\theta) d\theta. \quad (3)$$

Notice that the short-term funding decision x of any bank of class θ determines, via $\varepsilon(x, \theta)$,

¹¹Of course, an increasing $c(X)$ may also partly reflect that X increases the very probability of a systemic crisis. For example, the more vulnerable banks' funding structures are, the more likely it is that asset-side shocks such as a housing market bust or a stock market crash get transformed into a systemic liquidity shock.

the vulnerability of that very bank to a systemic crisis, and also, via $c(X)$, the likelihood and/or costs of a systemic crisis to all other banks.

4 Equilibrium

In an unregulated competitive equilibrium each bank chooses x so as to maximize its own expected NPV, $v(x, X, \theta)$, taking X as given. So an *unregulated competitive equilibrium* is a pair $(\{x^e(\theta)\}, X^e)$ that satisfies:

1. $x^e(\theta) = \arg \max_x \{\pi(x, \theta) - \varepsilon(x, \theta)c(X^e)\}$ for all $\theta \in [0, 1]$,
2. $X^e = \int_0^1 x^e(\theta) f(\theta) d\theta$.

Under our assumptions, banks' privately optimal choice of x under given values of θ and X can be described as the maximum between zero (which corresponds to a corner solution) and the unique solution $y(\theta, X)$ to the first order condition

$$\pi_x(y(\theta, X), \theta) - \varepsilon_x(y(\theta, X), \theta)c(X) = 0, \quad (4)$$

which characterizes an interior solution. Given the assumed properties of the functions involved in the first order condition, the implicit function theorem implies that $y(\theta, X)$ is increasing in θ and decreasing in X .¹²

The equilibrium value of X can be found as the fixed point of the auxiliary function $h(X) = \int_0^1 \max\{0, y(\theta, X)\} f(\theta) d\theta$. This function is decreasing in X insofar as $y(\theta, X) > 0$ for some positive measure set of values of θ . If we assume $y(1, 0) > 0$, so that at least the banks with the largest valuation of short-term funding (those with $\theta = 1$) find it worthy to use some of it if $X = 0$, then it is guaranteed that $h(0) > 0$ and, by standard arguments, there exists a unique fixed point $X^e > 0$.

In what follows we will refer to an equilibrium as *interior* if it satisfies $x^e(\theta) > 0$ for all $\theta > 0$, in which case the first order condition

$$\pi_x(x^e(\theta), \theta) - \varepsilon_x(x^e(\theta), \theta)c(X^e) = 0, \quad (5)$$

¹²Recall that we have assumed $\pi_{x\theta} > 0$ and $\varepsilon_{x\theta} \leq 0$.

with $X^e = \int_0^1 x^e(\theta)f(\theta)d\theta$, is satisfied for all $\theta \in [0, 1]$. Guaranteeing that the equilibrium is interior requires having

$$\pi_x(0, 0) - \varepsilon_x(0, 0)c(X) \geq 0, \quad (6)$$

for a sufficiently large X (e.g., larger than the possibly emerging X^e), so that even the banks with the lowest inclination for short-term funding (those with $\theta = 0$) want to use some of it in equilibrium.¹³ As shown below, the presence of systemic risk externalities will make the conditions stated in (5) incompatible with social efficiency.

5 The social planners' problem

The socially optimal allocation of short-term funding across banks can be found by maximizing social welfare W taking into account the influence of each individual bank funding strategy on X . Formally, a *socially optimal allocation* can be defined as a pair $(\{x^*(\theta)\}, X^*)$ that satisfies:

$$\begin{aligned} (\{x^*(\theta)\}, X^*) = \arg \max_{(\{x(\theta)\}, X)} & \int_0^1 [\pi(x(\theta), \theta) - \varepsilon(x(\theta), \theta)c(X^*)]f(\theta)d\theta \\ \text{s.t.} & \int_0^1 x(\theta)f(\theta)d\theta = X^*. \end{aligned} \quad (7)$$

After substituting the constraint in the objective function, we can find the social optimum by solving the system of equations:

$$\{x^*(\theta)\} = \arg \max_{\{x(\theta)\}} \int_0^1 [\pi(x(\theta), \theta) - \varepsilon(x(\theta), \theta)c(\int_0^1 x(z)f(z)dz)]f(\theta)d\theta \quad (8)$$

for $\theta \in [0, 1]$, and then finding $X^* = \int_0^1 x^*(\theta)f(\theta)d\theta$, recursively.

The assumptions adopted in Section 3 guarantee the existence of a unique socially optimal allocation. For values of θ with $x^*(\theta) > 0$, the first order condition associated with the maximization in (8) establishes

$$\pi_x(x^*(\theta), \theta) - \varepsilon_x(x^*(\theta), \theta)c(X^*) - E_z(\varepsilon(x^*(z), z))c'(X^*) = 0, \quad (9)$$

¹³A sufficient condition for this would be to have $\pi_x(0, 0) - \varepsilon_x(0, 0)c(\bar{X}) \geq 0$ with \bar{X} implicitly defined by the equation $\pi_x(\bar{X}, 1) - \varepsilon_x(\bar{X}, 1)c(\bar{X}) = 0$, which would characterize the equilibrium value of aggregate systemic risk in the hypothetical scenario in which all banks were of the highest type $\theta = 1$. To obtain most of the results below, we need not constrain attention to interior equilibria, but a full discussion of the solutions involving $x(\theta) = 0$ for some θ would make the presentation unnecessarily cumbersome.

where $E_z(\varepsilon(x^*(z), z)) = \int_0^1 \varepsilon(x^*(z), z)f(z)dz$. Intuitively, and similarly to the equilibrium allocation, insofar as $x^*(\theta) > 0$, the socially optimal allocation will assign larger short-term funding to the banks with higher θ . Hence, the socially optimal allocation will be interior if and only if

$$\pi_x(0, 0) - \varepsilon_x(0, 0)c(X^*) - E_\theta(\varepsilon(x^*(\theta), \theta))c'(X^*) > 0. \quad (10)$$

An interior socially optimal allocation can then be guaranteed if, e.g., the profit function satisfies $\pi_x(0, 0) \rightarrow \infty$ and the functions $\varepsilon(x, \theta)$ and $c(X)$ have finite derivatives with respect to x and X , respectively.¹⁴

Relative to the conditions for individual bank optimization in an interior equilibrium, given in (5), the conditions in (9) add a third, negative term that reflects the marginal external cost associated with each $x(\theta)$. The marginal external cost relevant for a bank of class θ is made of two multiplicative factors: the average vulnerability of all the banks in the system to a systemic crisis, $E_z(\varepsilon(x^*(z), z))$, and the marginal effect of aggregate funding risk on systemic crisis costs, $c'(X^*)$.

The presence of the external cost term in (9), implies that any equilibrium allocation $\{x^e(\theta)\}$ with $x^e(\theta) > 0$ for some positive measure set of θ s cannot coincide with the socially optimal allocation $\{x^*(\theta)\}$. To see this, notice that, if $\{x^e(\theta)\}$ were socially optimal, then, by (9), for any θ with $x^e(\theta) > 0$, we should have

$$\pi_x(x^e(\theta), \theta) - \varepsilon_x(x^e(\theta), \theta)c(X^e) - E_z(\varepsilon(x^e(z), z))c'(X^e) = 0. \quad (11)$$

However, the first two terms add up to zero, by (5), while the third term is strictly negative since $E_z(\varepsilon(x^e(z), z)) > 0$ and $c'(X^e) > 0$. So (11) does not hold.

The presence of a negative externality associated with the contribution of banks' short-term funding to aggregate systemic risk suggests that the socially optimal allocation will generally involve a strictly lower level of systemic risk than the equilibrium allocation. Indeed, under the hypothesis that $X^* \geq X^e > 0$, (5) and (9) would imply that $x^e(\theta) > x^*(\theta)$

¹⁴In the same spirit as in the condition given in footnote 4, an alternative sufficient condition would be to have $\pi_x(0, 0) - \varepsilon_x(0, 0)c(\bar{X}) - \varepsilon(\bar{X}, 0)c'(\bar{X}) \geq 0$ with \bar{X} implicitly defined by the equation $\pi_x(\bar{X}, 1) - \varepsilon_x(\bar{X}, 1)c(\bar{X}) - \varepsilon(\bar{X}, 1)c'(\bar{X}) = 0$, which would characterize the socially optimal level of aggregate systemic risk in the hypothetical scenario in which all banks were of the highest type $\theta = 1$. The inequality above uses $\varepsilon(\bar{X}, 0)$ as an upper bound to $E_\theta(\varepsilon(x^*(\theta), \theta))$, since the vulnerability function $\varepsilon(x, \theta)$ is increasing in x and decreasing in θ .

for all θ with $x^e(\theta) > 0$ (and $x^e(\theta) = x^*(\theta) = 0$ otherwise), which would obviously mean $X^e > X^*$ and, hence, would contradict the hypothesis. In practical terms, the results implies that in the unregulated competitive equilibrium sufficiently many banks (if not all) use more short-term funding than they would use under the socially optimal allocation.¹⁵ The key insights from this section are summarized in the following proposition:

Proposition 1 *The presence of systemic externalities associated with banks funding decisions, $c'(X) > 0$, makes the equilibrium allocation socially inefficient and characterized by an excessive aggregate systemic risk $X^e > X^*$.*

Figure 1 depicts the short-term funding decisions that correspond to the unregulated competitive equilibrium and the socially optimal allocation under a parameterization whose details are easy to work out analytically. θ is assumed to be uniformly distributed over the $[0, 1]$ interval, the profit function is assumed to have the simple quadratic form $\pi(x, \theta) = (1 + \theta)x - \frac{1}{2}x^2$ and the product of the terms $\varepsilon(x, \theta)$ and $c(X)$ is jointly specified as $\varepsilon(x, \theta)c(X) = \alpha x X$, where α is a positive constant. Interior solutions can be guaranteed for $\alpha < 1/2$. The figure corresponds to the case with $\alpha = 0.15$. The horizontal axis reflects the values of θ and the vertical axis the unregulated equilibrium and the socially optimal allocations, $x^e(\theta)$ and $x^*(\theta)$, respectively. The unregulated equilibrium features systemic risk $X^e = 1.30$ and social welfare $W^e = 0.892$, while the social optimum features $X^* = 1.15$ and $W^* = 0.907$.

Summing up, systemic externalities associated with banks' short-term funding decisions create a positive wedge between the social and the private marginal costs of using short-term funding. Banks only internalize the implications of the funding choices for their own vulnerability to refinancing risk, disregarding their contribution to all other banks' systemic risk exposure and costs. Standard marginal reasoning when privately optimizing on short-

¹⁵Proving that any $x^e(\theta) > 0$ is socially excessive given X^e is immediate. However, the comparison between $x^e(\theta)$ and $x^*(\theta)$ for any θ is more tricky because (5) involves X^e while (9) involves $X^* < X^e$, and the differences between the second terms in each of these first order conditions creates potential ambiguity. In the numerical examples that we have explored, all banks use excessive short-term funding relative to the social optimum.

term funding make systemic risk in the unregulated equilibrium higher than socially optimal.

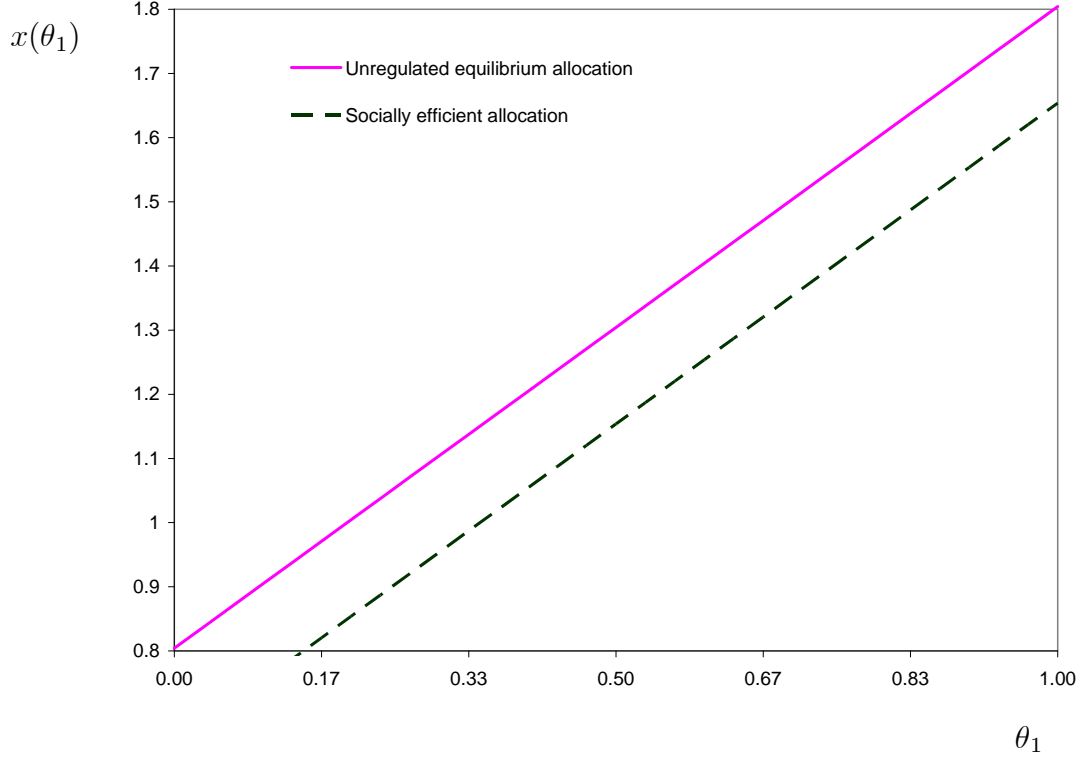


Figure 1. Unregulated equilibrium vs socially optimal allocations

6 The Pigovian tax: an efficient solution

As in the standard textbook discussion on the treatment of negative production externalities, the social efficiency of the competitive equilibrium can be restored by imposing a Pigovian tax, i.e. by taxing the activity causing the externality at a rate equal to the wedge between the social marginal cost and the private marginal cost of the activity (evaluated, if applicable, at the anticipated socially optimal allocation). In our case, this boils down to setting a flat tax per unit of short-term funding equal to

$$\tau^* = E_z(\varepsilon(x^*(z), z))c'(X^*). \quad (12)$$

Obviously, the introduction of a tax on short-term funding alters the first order condition relevant for banks' optimization in the competitive equilibrium with taxes.

Formally, we can define a *competitive equilibrium with taxes* $\{\tau(\theta)\}$ as a pair $(\{x^\tau(\theta)\}, X^\tau)$ satisfying:

1. $x^\tau(\theta) = \arg \max_x \{\pi(x, \theta) - \varepsilon(x, \theta)c(X^\tau) - \tau(\theta)x\}$ for all $\theta \in [0, 1]$,
2. $X^\tau = \int_0^1 x^\tau(\theta)f(\theta)d\theta$.

In an interior equilibrium, the first order conditions for the private optimality of each $x^\tau(\theta)$ imply

$$\pi_x(x^\tau(\theta), \theta) - \varepsilon_x(x^\tau(\theta), \theta)c(X^\tau) - \tau(\theta) = 0 \quad (13)$$

for all $\theta \in [0, 1]$. And it is immediate to see that the flat tax schedule $\tau(\theta) = \tau^*$, with the tax rate defined as in (12), makes (13) equivalent to (9), implying $(\{x^{\tau^*}(\theta)\}, X^{\tau^*}) = (\{x^*(\theta)\}, X^*)$ and, hence, implementing the socially optimal allocation as a competitive equilibrium.

To set the optimal tax rate τ^* properly, it is of course necessary that the regulator knows the functions that characterize the economy (including the density of the parameter θ that captures banks' heterogeneity) and is, hence, able to compute the socially optimal allocation that enters in (12). For the numerical example shown in Figure 1, the flat Pigovian tax that implements the socially optimal allocation as a competitive equilibrium with taxes is $\tau^* = 0.17$.¹⁶

An important practical difficulty when regulating heterogeneous agents is that the particulars of the regulation applicable to each agent may depend on information that its private to the agent. This problem does not affect the efficient Pigovian tax τ^* , which is the same for all values of θ . The following proposition summarizes the key results of this section.

Proposition 2 *When banks differ in the marginal value they can extract from short-term funding, the socially optimal allocation can be reached as a competitive equilibrium by charging banks a flat Pigovian tax τ^* on each unit of short-term funding.*

¹⁶Of course this number depends on the underlying parameterization. Without touching other details of our simple example, which was not intended to provide realistic predictions, reducing the parameter α from 0.15 to, e.g., 0.015 reduces the efficient Pigovian tax to $\tau^* = 0.02$.

7 Quantity-based and ratio-based alternatives

Pigovian taxation is frequently described as a price-based solution to the regulation of externalities. Such description emphasizes the capacity of the tax solution to decentralize the implementation of the desired allocation as a market equilibrium. The polar alternative is to go for a “centralized” quantity-based solution in which each regulated agent (bank) is directly mandated to choose its corresponding quantity (short-term funding) in the optimal allocation ($x^*(\theta)$ in the model).

In the context of our model, pure quantity-based regulation would require detailed knowledge by the regulator of individual marginal value of short-term funding for each bank (i.e., the derivatives $\pi_x(x, \theta)$ and $\varepsilon_x(x, \theta)$, which vary with θ and appear in (9)). Possibly due to the strong informational requirements that this implies, none of the alternatives for liquidity regulation considered in practice these days opts for directly setting individualized quantity prescriptions such as $x^*(\theta)$.

The alternatives to Pigovian taxes actually under discussion are ratio-based regulations, i.e. regulations that consist on forcing banks to have some critical accounting ratios above or below some regulatory minima or maxima. To be sure, some proposals include making the regulatory bounds functions of individual characteristics of each bank, such as size, interconnectedness, capitalization, etc. but none of the considered characteristics (except perhaps those referring to the regional or sectorial specialization of some banks) seem targeted to control for the heterogeneity in banks’ capacity to extract value from short-term funding. These qualifiers can be rather rationalized as an attempt to capture what, in an extension discussed in Section 9, we describe as the *systemic importance* of each bank (the relative importance of the contribution of its short-term funding to the systemic risk measure X).

The most seriously considered ratio-based alternatives for the regulation of liquidity are those put forward by the new Basel III agreement (see BCBS, 2010). The agreement introduces two new regulatory ratios: a *liquidity coverage ratio*, similar in format and spirit to one already introduced by the Financial Services Authority in the UK in October 2009, and a more innovative *net stable funding ratio*. To facilitate the discussion, we analyze each of these instruments as if it were introduced in isolation, starting with the last one, whose potential effectiveness for the regulation of funding maturity is somewhat less ambiguous.

7.1 A stable funding requirement

The net stable funding requirement calls banks to hold some accounting ratio of “stable funding” (i.e. equity, customer deposits, and other long-term or “stable” sources of funding) to “non-liquid assets” above some regulatory minimum. To translate this to our model, where banks’ asset composition and stable sources of funding have been so far treated as exogenously fixed, we can think of this requirement as equivalent to imposing an upper limit \bar{x} to the short-term debt that each bank can issue. In a more general version of our model, the effective upper limit applicable to each bank could be considered affected by prior decisions regarding the maturity and liquidity structure of the bank’s assets, its retail deposits base, its level of capitalization, etc. But here, for simplicity, one can think of changes in these decisions as possible interpretations for the comparative statics of \bar{x} .

The introduction of a minimum stable funding requirement has then the implication of adding an inequality constraint of the type $x \leq \bar{x}$ to the private optimization problem of the banks. Formally, a *competitive equilibrium with a stable funding requirement* parameterized by \bar{x} can be defined as a pair $(\{x^{\bar{x}}(\theta)\}, X^{\bar{x}})$ satisfying:

1. $x^{\bar{x}}(\theta) = \arg \max_{x \leq \bar{x}} \{\pi(x, \theta) - \varepsilon(x, \theta)c(X^{\bar{x}})\}$ for all $\theta \in [0, 1]$,
2. $X^{\bar{x}} = \int_0^1 x^{\bar{x}}(\theta)f(\theta)d\theta$.

Since the preference for short-term funding is strictly increasing in θ , we may have up to three possible configurations of equilibrium. For $\bar{x} \geq x^e(1)$, the stable funding requirement will not be binding for any bank (since $\theta = 1$ identifies the banks with the highest incentives to use short-term funding), and the equilibrium will then coincide with the unregulated competitive equilibrium characterized in Section 4. For $\bar{x} \leq x^e(0)$, the stable funding requirement will be binding for all banks (since $\theta = 1$ identifies the banks with the lowest incentives to use short-term funding), implying $x^{\bar{x}}(\theta) = \bar{x} < x^e(\theta)$ for all θ and, hence, $X^{\bar{x}} = \bar{x} < X^e$. Finally, for $\bar{x} \in (x^e(0), x^e(1))$, the stable funding requirement will be binding for at least the banks with the largest θ s and perhaps for all banks. To see the latter, notice that inducing the limit choice of $x^{\bar{x}}(\theta) = \bar{x} < x^e(\theta)$ to the banks with relatively large θ s will push $X^{\bar{x}}$ below X^e , but this, in turn, will push the banks with relatively low θ s into choices

of $x^{\bar{x}}(\theta) > x^e(\theta)$, possibly (but not necessarily) inducing some or even all of them to also hit the regulatory limit \bar{x} .

It is then obvious that, in general, a sufficiently tight stable funding requirement $\bar{x} < x^e(1)$ can reduce the equilibrium measure of aggregate systemic risk $X^{\bar{x}}$ relative to the unregulated equilibrium X^e , thus moving it closer to its value in the socially optimal allocation X^* . The induced allocation will, however, be necessarily inefficient. The reason for this is that the reduction in the activities that generate negative externalities comes at the cost of distorting the allocation of short-term funding across bank classes: (i) constraining the banks with relatively higher valuation for short-term funding to the common upper limit \bar{x} , and (ii) encouraging the banks with relatively low valuation for short-term funding to use more of it than it would be socially optimal (since they will choose $x^{\bar{x}}(\theta) > x^e(\theta) > x^*(\theta)$). In fact, there is no guarantee that introducing a \bar{x} that simply brings $X^{\bar{x}}$ closer to X^* improves, in welfare terms, over the unregulated equilibrium.

Proposition 3 *A binding net stable funding requirement will affect the measure of aggregate systemic risk X in the same direction as the efficient arrangement (i.e. will reduce X) but it will also redistribute short-term funding inefficiently from banks that value it more to banks that value it less, so that the socially optimal allocation cannot be reached and the improvement in social welfare is not guaranteed.*

Assuming a solution in which \bar{x} is binding for some but not all banks, the socially optimal choice of \bar{x} can be defined as follows:

$$\begin{aligned} \bar{x}^{SB} = \arg \max_{(\bar{x}, X^{\bar{x}})} & \int_0^{\bar{\theta}} [\pi(y(\theta, X^{\bar{x}}), \theta) - \varepsilon(y(\theta, X^{\bar{x}}), \theta) c(X^{\bar{x}})] f(\theta) d\theta + \int_{\bar{\theta}}^1 [\pi(\bar{x}, \theta) - \varepsilon(\bar{x}, \theta) c(X^{\bar{x}})] f(\theta) d\theta \\ \text{s.t.} & \int_0^{\bar{\theta}} y(\theta, X^{\bar{x}}) f(\theta) d\theta + \bar{x}[1 - F(\bar{\theta})] = X^{\bar{x}}, \end{aligned} \tag{14}$$

where the function $y(\theta, X)$ is defined as in (4) and $F(\theta)$ is the cumulative distribution function associated with $f(\theta)$. This problem identifies the “second best” allocation attainable if the only available instrument for liquidity regulation is the stable funding requirement \bar{x} , which is assumed to be binding only for the banks with $\theta \in (\bar{\theta}, 1]$, where $\bar{\theta}$ is implicitly defined by $y(\bar{\theta}, X^{\bar{x}}) = \bar{x}$.

To solve the problem in (14), the simplest approach is to treat the constraint as a definition of the aggregate systemic risk $X^{\bar{x}}$ induced by the choice of \bar{x} , so that \bar{x} is effectively the only decision variable. The first order condition for an optimal interior solution in this variable will then require that the full differential of the objective function above with respect to \bar{x} is made equal to zero at the optimum \bar{x}^{SB} . Such full differential of the objective function will generally have terms related to (i) the direct effects of a marginal variation in \bar{x} on the integrands, (ii) the effects of a marginal variation in \bar{x} on $\bar{\theta}$ (and through it on social welfare) and (iii) the effects of a marginal variation in \bar{x} on $X^{\bar{x}}$ (and through it on social welfare). However, it turns out that the effects of type (i) are only relevant in the second term of the objective function (i.e. for the net value generated by banks for which the stable funding requirement is binding), (ii) is zero, and the effects of type (iii) channelled via $y(\theta, X^{\bar{x}})$ for the banks for which the stable funding requirement is not binding are also zero by virtue of (4).¹⁷

Eventually, the only non-zero terms resulting from the full differentiation of social welfare with respect to \bar{x} give rise to the following first order condition for an interior solution to the maximization problem in (14):

$$\int_{\bar{\theta}}^1 [\pi_x(\bar{x}, \theta) - \varepsilon_x(\bar{x}, \theta)c(X^{\bar{x}})]f(\theta)d\theta - E_{\theta}(\varepsilon(x^{\bar{x}}(\theta), \theta))c'(X^{\bar{x}})\frac{dX^{\bar{x}}}{d\bar{x}} = 0, \quad (15)$$

where

$$\frac{dX^{\bar{x}}}{d\bar{x}} = \frac{1 - F(\bar{\theta})}{1 - \int_0^{\bar{\theta}} y_X(\theta, X^{\bar{x}})f(\theta)d\theta} \in [0, 1]. \quad (16)$$

The first term in the LHS of (15) collects the direct effects of a marginal change in \bar{x} on the expected NPV generated by the banks for which the stable funding requirement is binding. This term is positive because, for given $X^{\bar{x}}$, relaxing the constraint associated with \bar{x} would produce more NPV at each constrained bank. The second term captures the impact of the requirement \bar{x} on aggregate systemic risk $X^{\bar{x}}$ and, through it, on the negative externalities suffered by all banks. This term is clearly negative and the only reason why introducing a binding \bar{x}^{SB} may be possibly optimal in a second-best sense. As reflected in (16), $dX^{\bar{x}}/d\bar{x}$ will typically be lower than one because moving \bar{x} will directly reduce the systemic risk generated

¹⁷This last effect is simply the particularization of the envelope theorem to the individual optimization decisions of unconstrained banks.

by the constrained banks only, while indirectly encouraging (due to the very reduction in $X^{\bar{x}}$) the use of short-term funding by the unconstrained banks.¹⁸

To gain further intuition on the second-best nature of the trade-offs behind the choice of \bar{x}^{SB} , it is convenient to compare (15) with the condition for first-best efficiency in (9). First, (9) applies point-wise, defining an efficient $x^*(\theta)$ for each θ ; in contrast, (15) determines a common upper limit \bar{x} to the short-term funding decisions of all banks. Hence, in (15) the costs and benefits of marginally moving \bar{x} are “averaged” over all the θ s. The terms in the integral that appears in (15) resemble the first two terms in the left hand side of (9), but the ones “averaged” here correspond to the set of high θ s only, for which the requirement \bar{x} is binding—the remaining banks are not directly affected by \bar{x} . The second term in (15) and the third in (9) reflect the marginal externality caused by \bar{x} and each $x^*(\theta)$, respectively. The relevant difference regarding these terms is due to the presence of $dX^{\bar{x}}/d\bar{x}$ (typically lower than one) in (15).

Intuitively, the stable funding requirement emerges as a just too rough way to deal with the underlying systemic externalities. It hinges on constraining in a fully indiscriminate manner the decisions made by the banks that, under the source of heterogeneity considered here, are precisely those that can generate more value from short-term funding. And it has the undesirable (though possibly quantitatively small) side effect of encouraging those that generate less value from short-term funding to use more of it. In the numerical example depicted in Figure 2, built on the same assumptions as Figure 1, we have that $\bar{x}^{SB} = 1.47$, which means the best stable funding requirement is binding for the banks in the top 34% quantile of the distribution of θ s. For lower values of θ , the resulting allocation involves some $x^{\bar{x}}(\theta)$ slightly above those in the unregulated equilibrium, $x^e(\theta)$, and the whole second-best allocation attainable with \bar{x}^{SB} is quite distant from the socially efficient one, $x^*(\theta)$. The best stable funding requirement leads to a social welfare of $W^{\bar{x}} = 0.896$ (only 0.4% larger than in the unregulated equilibrium), while the optimal flat Pigovian tax reaches the first-best level $W^* = 0.907$ (1.6% more than in the unregulated equilibrium). The best stable funding requirement reduces aggregate systemic risk by 3.9%, while the optimal Pigovian tax reduces

¹⁸The fact that $y(\theta, X^{\bar{x}})$ is decreasing in $X^{\bar{x}}$ follows directly from applying the implicit function theorem on (4).

it by 11.5%.

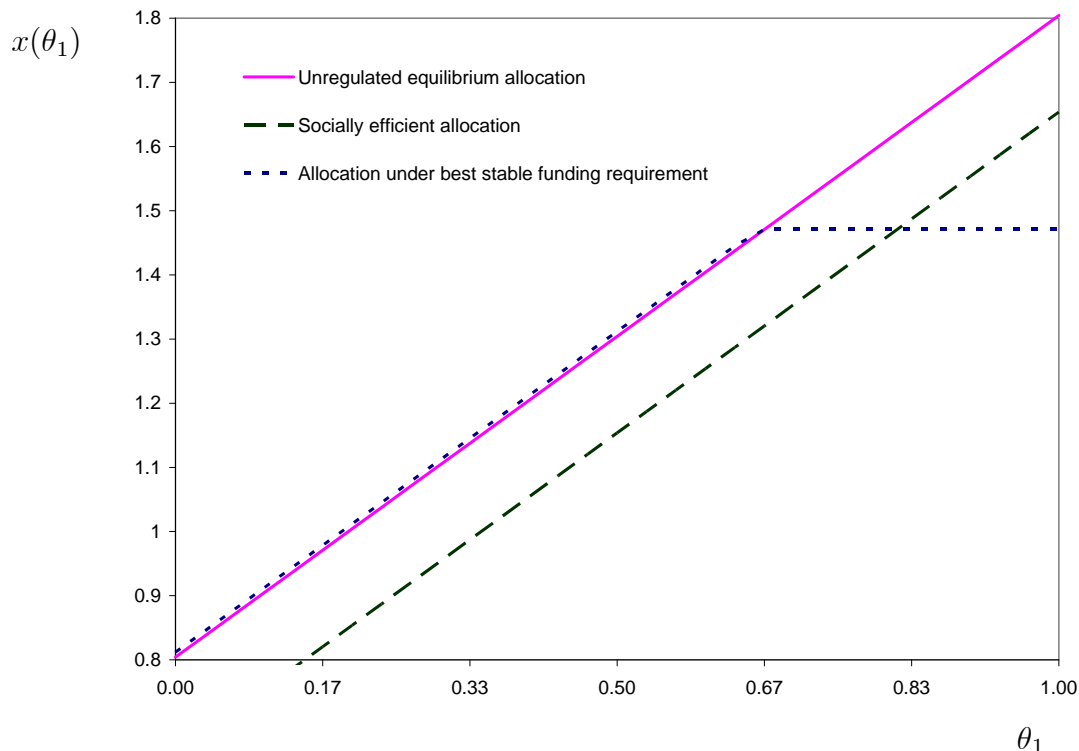


Figure 2. Equilibrium allocation under the best stable funding requirement

7.2 A liquidity requirement

The liquidity coverage ratio introduced by Basel III requires banks to back their use of short-term funding with the holding of high-quality liquid assets, i.e. assets that could be easily sold, presumably at no fire-sale loss, in case of a crisis. In its original description this requirement responds to the motivation of providing each bank with its own liquidity buffer, which, presumably might also expand the liquidity available in the system in case of a crisis (on top of that possibly provided by the lender of last resort).

Specifically, it is proposed that banks estimate the refinancing needs that they would accumulate if the functioning of money markets or other conventional borrowing sources were disrupted for some specified period (one month) and keep enough high-quality liquid

assets so as to be able to confront the situation with their sale.¹⁹ Qualifying assets would essentially be cash, central bank reserves and treasury bonds.

How can we capture this requirement in the context of our model? Leaving details aside, the liquidity requirement can be seen as a requirement to back some minimal fraction $\phi \leq 1$ of each bank's short-term funding x with the holding of qualifying liquid assets m , thereby introducing the constraint $m \geq \phi x$. Additionally, the impact of m on the bank's objective function could be taken into account by considering the following extended value function:

$$v(x, m, X, \theta) = \pi(x - m, \theta) - \varepsilon(x - m, \theta)c(\widehat{X}) - \delta m, \quad (17)$$

where

$$\widehat{X} = \int_0^1 [x(\theta) - m(\theta)]f(\theta)d\theta, \quad (18)$$

and $\delta = r_b - r_m \geq 0$ is the difference between the bank's short-term borrowing rate r_b and the yield r_m of the qualifying liquid assets. This formulation credits for both the individual and the systemic "buffering" role of the liquid assets by making each bank's individual vulnerability factor $\varepsilon(x - m, \theta)$ a function of its *net* short-term funding and by redefining the systemic risk measure \widehat{X} as the result of aggregating banks' net short-term funding positions.

The other terms in (17) capture the NPV generated in the absence of a systemic crisis. Our formulation is based on assuming that the former function $\pi(x, \theta)$ only captured the NPV generated by the bank's core lending or investment activity, which does not include investing in the qualifying liquid assets. The new first argument of $\pi(x - m, \theta)$ is justified by the fact that if a part m of the resources obtained as short-term funding x is invested in liquid assets, the net amount available for core banking activities becomes $x - m$. We assume that the funds m invested in liquid assets yield a risk-free rate r_m but have a cost equal to the bank's short-term borrowing rate $r_b \geq r_m$, which for simplicity, we will treat r_b as exogenous (and equal for all banks).²⁰ Hence the spread $\delta = r_b - r_m \geq 0$ acts as a direct net cost of holding liquid assets.

¹⁹Or by posting them as collateral at the central bank's discount window.

²⁰In principle, r_b should depend on each bank's risk of insolvency as assessed by market participants. In the logic of the model one might think of the spread as directly connected to $\varepsilon(x - m, \theta)$ and perhaps also to $c(X)$. However, the exogeneity of r_b can be justified by alluding either to informational asymmetries or safety-net guarantees that effectively make investors not to discriminate across banks. Additionally, one could also argue that the main driver of banks' insolvency risk is asset risk, from which we are abstracting.

In this extended framework, social welfare can be written as:

$$W(\{x(\theta), m(\theta)\}) = \int_0^1 [\pi(x(\theta) - m(\theta), \theta) - \varepsilon(x(\theta) - m(\theta), \theta)c(\widehat{X}) - \delta m(\theta)]f(\theta)d\theta, \quad (19)$$

where the presence of $-\delta m(\theta)$ implies considering banks' direct costs of holding liquidity as a deadweight loss.²¹

A *competitive equilibrium with a liquidity requirement* parameterized by ϕ can be defined as a pair $(\{(x^\phi(\theta), m^\phi(\theta))\}, \widehat{X}^\phi)$ satisfying:

1. $(x^\phi(\theta), m^\phi(\theta)) = \arg \max_{m \geq \phi x} \{\pi(x - m, \theta) - \varepsilon(x - m, \theta)c(\widehat{X}^\phi) - \delta m\}$ for all $\theta \in [0, 1]$,
2. $\widehat{X}^\phi = \int_0^1 (x^\phi(\theta) - m^\phi(\theta))f(\theta)d\theta$.

Before continuing it may be worth commenting on how equilibrium looks in the special limit case with $\phi = 1$. Obliging banks to match each unit of short term funding with one unit of liquid assets implies that banks are left with no funds to expand their value creating activities. In other words, banks will have $\widehat{x} = x - m = 0$ for any possible choice of x . However any $x > 0$ will imply $m = x > 0$, and hence paying the corresponding direct cost $\delta m > 0$ of holding liquidity, so banks will optimally fix $x = m = 0$. Of course, aggregate systemic risk will be zero, i.e. $\widehat{X}^{\phi=1} = 0$, but the welfare generated by the banks, as measured by (19), will also be zero. In what follows we will focus on the case with $\phi < 1$.

In general the liquidity requirement can be taken as binding (necessarily so if $\delta > 0$ and binding without loss of generality if $\delta = 0$). For $\phi < 1$ it is possible and convenient to reformulate banks' optimization problem in terms of their choice of *net* short-term funding $\widehat{x}(\theta) = x(\theta) - m(\theta)$ only, since the binding liquidity constraint allows us to write $m(\theta)$ as $\frac{\phi}{1-\phi}\widehat{x}(\theta)$ and eliminate it as an independent decision variable. Hence, equilibrium can be redefined as a pair $(\{\widehat{x}^\phi(\theta)\}, \widehat{X}^\phi)$ satisfying:

1. $\widehat{x}^\phi(\theta) = \arg \max_{\widehat{x}} \{\pi(\widehat{x}, \theta) - \varepsilon(\widehat{x}, \theta)c(\widehat{X}^\phi) - \frac{\delta\phi}{1-\phi}\widehat{x}\}$ for all $\theta \in [0, 1]$,
2. $\widehat{X}^\phi = \int_0^1 \widehat{x}^\phi(\theta)f(\theta)d\theta$.

²¹This view is consistent with having assumed that investors provide (short-term) funding to the banks at competitive market rates and thus make zero NPV when doing so. In this context, $\delta > 0$ is a premium that compensates for (unmodeled) utility losses derived from either the risk or the lower liquidity of an investment in bank liabilities.

We will proceed with the analysis by looking first at the case in which the net cost of holding liquid assets is zero ($\delta = 0$) and then at the case in which it is positive ($\delta > 0$).

7.2.1 The case in which holding liquidity is costless ($\delta = 0$)

The following proposition establishes a somewhat shocking result for the relevant case in which the spread δ is zero (roughly the case in “normal times”, when banks are perceived as essentially risk-free borrowers):

Proposition 4 *With $\delta = 0$, the competitive equilibrium with a liquidity requirement $\phi < 1$ involves the same amount of net short-term funding and, hence, the same level of systemic risk as the unregulated equilibrium. That is, it involves $x^\phi(\theta) - m^\phi(\theta) = x^e(\theta)$ and $\widehat{X}^\phi = X^e$.*

The proof of this proposition follows immediately from the equivalence, when $\delta = 0$, between the equilibrium conditions for $(\{\widehat{x}^\phi(\theta)\}, \widehat{X}^\phi)$ and those for $(\{x^e(\theta)\}, X^e)$ (see Section 4). Hence, the only effect of the liquidity requirement relative to the unregulated equilibrium is to induce an artificial demand $M^\phi = \frac{\phi}{1-\phi} E_\theta(x^e(\theta))$ for the qualifying liquid assets and a spurious increase in banks’ gross short-term funding, which becomes $E_\theta(x^\phi(\theta)) = E_\theta(x^e(\theta)) + M = \frac{1}{1-\phi} E_\theta(x^e(\theta))$.

Therefore, when the direct net cost δ of each unit of liquidity that the requirement forces banks to hold is zero (not implausible in “normal times”), the liquidity coverage ratio totally fails to bring the equilibrium allocation any closer to the socially optimum than in the unregulated scenario. Banks respond to regulation by increasing their short-term funding and their liquidity holding so as to make their net short-term funding as high as in the unregulated equilibrium. The artificial demand for high-quality liquid assets may imply that liquid assets, kept somewhere else in the financial system (e.g. money market mutual funds) prior to imposing the ratio, end up kept by banks after imposing the ratio. However the (net) systemic risk generated by the banks will not change.

7.2.2 The case in which holding liquidity is costly ($\delta > 0$)

When the direct net unit cost of holding liquidity, δ , is positive, the implications are quite different. The conditions that characterize $(\{\widehat{x}^\phi(\theta)\}, \widehat{X}^\phi)$ become analogous to those that

characterize a competitive equilibrium with taxes in which $\tau(\theta) = \frac{\delta\phi}{1-\phi}$ (see Section 6). This allows us to directly state the following proposition:

Proposition 5 *With $\delta > 0$, the competitive equilibrium with a liquidity requirement $\phi < 1$ involves the same individual net short-term funding decisions and aggregate systemic risk as a competitive equilibrium with a tax on short-term funding with rate $\tau(\theta) = \frac{\delta\phi}{1-\phi}$ for all θ .*

For a given $\delta > 0$, the implicit tax rate $\delta\phi/(1 - \phi)$ described above moves from zero to infinity as the liquidity requirement ϕ moves from zero to one. Thus the regulator can *seemingly* replicate the effects of any flat tax (including the efficient Pigovian tax τ^* of Section 6) by setting $\phi^* = \frac{\tau^*}{\delta + \tau^*}$. However, banks' demand for the qualifying liquid assets would be $m^\phi(\theta) = \frac{\phi}{1-\phi}\widehat{x}^\phi(\theta) = \frac{\tau}{\delta}x^\tau(\theta)$ (implying an aggregate demand $M^\phi = \frac{\tau}{\delta}X^\tau$) and their *gross* short-term funding would be $x^\phi(\theta) = x^\tau(\theta) + m^\phi(\theta) = \frac{\delta + \tau}{\delta}x^\tau(\theta) > x^\tau(\theta)$ (implying $X^\phi = E_\theta(x^\phi(\theta)) = X^\tau + M^\phi = \frac{\delta + \tau}{\delta}X^\tau > X^\tau$ at the aggregate level). Importantly, the total direct net costs of holding liquidity would cause a deadweight loss of $\delta m^\phi(\theta) = \tau x^\tau(\theta)$ to each bank. Not surprisingly, the aggregate deadweight loss $\delta M^\phi = \tau X^\tau$ equals the tax revenue that the replicated tax on short-term funding could have raised.

The presence of the deadweight loss τ^*X^* implies that the liquidity requirement that *seemingly* replicates the Pigovian solution ($\phi^* = \frac{\tau^*}{\delta + \tau^*}$) is not socially efficient.

Proposition 6 *With $\delta > 0$, replicating the net short-term funding allocation and aggregate systemic risk of the efficient allocation using a liquidity requirement $\phi^* = \frac{\tau^*}{\delta + \tau^*}$ is feasible, but entails a deadweight loss $\tau^*X^* > 0$.*

Actually, ϕ^* will not generally be optimal even from a second best perspective. Except in the non-generic situation in which the efficient Pigovian tax τ^* happens to be at a critical point of the Laffer curve τX^τ . This is because moving the liquidity requirement marginally away from ϕ^* (in one direction) will reduce the deadweight loss δM^ϕ , while other components of social welfare will not change (since they are maximized precisely with $\phi = \phi^*$).

For a given spread $\delta > 0$, the socially optimal liquidity requirement will be some $\phi^{SB} = \frac{\tau^{SB}}{\delta + \tau^{SB}}$ whose associated implicit tax rate τ^{SB} satisfies:

$$\begin{aligned} \tau^{SB} = \arg \max_{\tau \geq 0} & \int_0^1 [\pi(x^\tau(\theta), \theta) - \varepsilon(x^\tau(\theta), \theta)c(X^\tau) - \tau x^\tau(\theta)] f(\theta) d\theta \\ \text{s.t.} & x^\tau(\theta) = \arg \max_x \pi(x, \theta) - \varepsilon(x, \theta)c(X^\tau) - \tau x \text{ for all } \theta \\ & \int_0^1 x^\tau(\theta) f(\theta) d\theta = X^\tau. \end{aligned} \tag{20}$$

This formulation of the optimization problem exploits the analogy explained above, which conveniently allows us to write the deadweight loss suffered by each bank as $\tau x^\tau(\theta)$, which is actually independent of δ and will end up making the solution in terms of τ^{SB} also independent of δ . The constraints in the optimization problem are simply the conditions that define an equilibrium with a tax τ on short-term funding (see Section 6).

Typically, the optimal liquidity requirement ϕ^{SB} will be inferior to ϕ^* , implying more aggregate systemic risk than in the first best allocation. The intuition for this is that moving away from the unregulated equilibrium allocation by increasing ϕ will typically monotonically increase the aggregate deadweight loss δM^ϕ , while the remaining marginal benefits of moving towards the first best allocation decline towards zero as ϕ approaches ϕ^* . In fact, for the parameterization behind Figures 1 and 2, the optimal liquidity requirement ϕ^{SB} is just zero, meaning that the deadweight losses associated with forcing banks to hold liquidity do not compensate, for any level of ϕ , the gains from the reduction in systemic risk.²²

Interestingly, the writing of the problem as in (20) makes clear that τ^{SB} does not depend on δ , implying that the total variation of $\phi^{SB} = \frac{\tau^{SB}}{\delta + \tau^{SB}}$ with respect to δ is just given by the partial derivative

$$\frac{\partial \phi^{SB}}{\partial \delta} = \frac{-\tau^{SB}}{(\delta + \tau^{SB})^2} < 0. \quad (21)$$

Hence, if the regulator wants to implement an interior second best allocation as defined above (or to seemingly replicate the efficient Pigovian tax), it should be ready to move the imposed liquidity requirement ϕ^{SB} (or ϕ^*) in response to the fluctuations in the spread δ . In practice, moving ϕ and the implied adjustments in quantities may be a source of trouble. On the one hand, authorities will have to be effective in changing ϕ in due course. On the other hand, frequent and sudden changes ϕ might produce changes in M^ϕ that, for reasons left outside the model (such as monetary stability) might not be admissible. This might be especially so if δ approaches zero, in which case the prescriptions for ϕ^{SB} (or ϕ^*) imply that M^ϕ would tend to infinity.

²²In our example, the liquidity requirement that replicates the effects of the optimal Pigovian tax ($\tau^* = 0.17$) has deadweight costs equivalent to 22% of the value generated by banks in the unregulated equilibrium. These costs are more than ten times the increase in welfare that the pure Pigovian solution might produce relative to the unregulated equilibrium.

With potentially large variations in the demand for liquid assets, the reference to possible general equilibrium implications is inexcusable. Our results are definitely valid if the supply of liquid assets is sufficiently elastic at the rate r_m used in the definition of the spread δ . However, if changes in liquidity requirements have an impact on the equilibrium value of r_m due to a more general interaction between the demand and supply for liquid assets, then the required analysis may well exceed the scope of this paper. This is especially so if r_m is also the reference risk-free rate in the economy, because then the net present values captured by our reduced-form value functions might be expected to change with r_m and the whole analysis in (at least) this section should be redone using a somewhat more structural formulation.

8 Risk-shifting and the case for quantity regulation

In this section we extend the model to address formally one of the main criticisms to the proposal of a Pigovian approach to liquidity risk regulation. Such criticism is based on the “robustness” of the price-based approach to modeling mistakes and, specifically, to the possibility of having some “crazy” or just particularly risk-inclined banks that, for the sake of expanding their risky lending, are willing to pay large amounts of the established tax so as to use large amounts of short-term funding.

In our baseline formulation, banks that like to take more short-term funding are those that can extract more expected NPV from it. In such formulation, the considered dimension of heterogeneity makes banks with larger θ essentially more valuable, privately and, if properly regulated, also socially. We will now denote that dimension of heterogeneity by θ_1 and introduce a second dimension of heterogeneity, $\theta_2 \in [0, 1]$, intended to capture differences in banks’ inclination towards risk-taking.²³ The joint distribution of (θ_1, θ_2) will be described by the density function $f(\theta_1, \theta_2)$.

²³The literature has identified several sources of such differences. Corporate governance arrangements may affect the severity of the conflicts of interest between shareholders and debtholders, making the former more or less capable to ex post expropriate the former by shifting risk (Jensen and Meckling, 1976). In the case of banks, risk-shifting problems are exacerbated by the existence of safety net guarantees (e.g. deposit insurance) provided at risk-insensitive rates. In such a setup, banks’ charter values reduce excessive risk-taking (Keeley, 1990). Capital requirements (especially if risk-based) generally improve the alignment of incentives between the bankers and other stakeholders (Holmstrom and Tirole, 1997) and can specifically attenuate the risk-shifting problem (Hellmann, Murdock, and Stiglitz, 2000, and Repullo, 2004).

To capture heterogeneity in banks' risk-shifting inclinations formally, we are going to treat θ_2 as a parameter that determines the fraction of the losses incurred by a bank during a crisis which are *not* internalized by its owners but passed (without compensation) to other stakeholders (e.g. the deposit insurer). We then assume each bank, when privately deciding on x , only considers the fraction $1 - \theta_2$ of $\varepsilon(x, \theta)c(X)$ as an expected value loss, leaving the remaining fraction θ_2 to other stakeholders. Hence, the social welfare measure $W(\{x(\theta)\})$ must now explicitly consider, in addition to the NPV appropriated by the bank owners, the losses $-\theta_2\varepsilon(x, \theta)c(X)$ passed on to other bank stakeholders.

So the new objective function for banks is:

$$v(x, X, \theta_1, \theta_2) = \pi(x, \theta_1) - (1 - \theta_2)\varepsilon(x, \theta_1)c(X), \quad (22)$$

while social welfare is given by:

$$W(\{x(\theta_1, \theta_2)\}) = \int_0^1 \int_0^1 [v(x(\theta_1, \theta_2), X, \theta_1, \theta_2) - \theta_2\varepsilon(x(\theta_1, \theta_2), \theta_1)c(X)]f(\theta_1, \theta_2)d\theta_1d\theta_2, \quad (23)$$

where

$$X = g(\{x(\theta_1, \theta_2)\}) = \int_0^1 \int_0^1 x(\theta_1, \theta_2)f(\theta_1, \theta_2)d\theta_1d\theta_2. \quad (24)$$

Plugging (22) into (23), social welfare can be written as

$$W(\{x(\theta_1, \theta_2)\}) = \int_0^1 \int_0^1 [\pi(x(\theta_1, \theta_2), \theta_1) - \varepsilon(x(\theta_1, \theta_2), \theta_1)c(X)]f(\theta_1, \theta_2)d\theta_1d\theta_2, \quad (25)$$

which is conceptually identical to (3).

8.1 Gambling as the sole source of heterogeneity

To highlight our key argument, suppose that the variation due to θ_1 , whose implications we have already discussed in prior sections, is shut down by fixing $\theta_1 = \bar{\theta}_1$ for all banks. So residual bank heterogeneity is due to θ_2 only. How is the unregulated equilibrium determined? And the socially optimal allocation? How do they differ? How should $x(\theta_2)$ be regulated?

Without restating all the relevant definitions (which will follow mechanically from the adaptation of those already presented for the baseline model), the answers to these questions can be found by comparing the first order conditions satisfied by bank decisions, $x^{ee}(\theta_2)$, and

the systemic risk measure, X^{ee} , in an interior unregulated equilibrium, with the conditions satisfied by their counterparts, $x^{**}(\theta_2)$ and X^{**} , in an interior socially optimal allocation. Similarly to (5), the unregulated equilibrium objects satisfy:

$$\pi_x(x^{ee}(\theta_2), \bar{\theta}_1) - (1 - \theta_2)\varepsilon_x(x^{ee}(\theta_2), \bar{\theta}_1)c(X^{ee}) = 0, \quad (26)$$

while, similarly to (9), in the socially optimal allocation we must have:

$$\pi_x(x^{**}(\theta_2), \bar{\theta}_1) - \varepsilon_x(x^{**}(\theta_2), \bar{\theta}_1)c(X^{**}) - E_z(\varepsilon(x^{**}(z), \bar{\theta}_1))c'(X^{**}) = 0, \quad (27)$$

in both cases for all θ_2 . From these conditions, it is immediate to conclude that $x^{ee}(\theta_2)$ is increasing in θ_2 (that is, banks with greater risk-shifting inclinations tend to use more short-term funding) while $x^{**}(\theta_2)$ is independent of θ_2 and, hence, equal to a constant \bar{x}^{**} (since, for any given x , θ_2 determines the distribution of value across bank stakeholders but not the total marginal value of short-term funding).

By simple comparison of the two sets of conditions, it is now obvious that the efficient Pigovian tax schedule is

$$\tau^{**}(\theta_2) = \theta_2\varepsilon_x(x^{**}(\theta_2), \bar{\theta}_1)c(X^{**}) + E_z(\varepsilon(x^{**}(z), \bar{\theta}_1))c'(X^{**}), \quad (28)$$

where the first term is new relative to (12) and reflects that risk shifting incentives produce additional discrepancies between the private and social costs of expanding banks' short-term funding. In contrast to the pure systemic externality term (identical to what we had in the baseline model), the first term depends on θ_2 . Hence, the efficient Pigovian tax schedule is not flat and cannot be enforced without detailed knowledge of each bank's risk-shifting inclination. A flat tax on short-term funding will not implement the first best allocation.

Now, however, proper quantity regulation can do a great job. Specifically, a net stable funding requirement that effectively imposes the first best quantity \bar{x}^{**} as a limit to each bank's use of short-term funding would implement the first best. It is easy to see that the regulatory constraint will be binding for all θ_2 . As for liquidity requirements, the rather negative conclusions obtained in the baseline analysis would still apply: with $\delta = 0$, a liquidity requirement is as ineffective as it was there, while with $\delta > 0$ its effect is very similar to (but with worse welfare properties than) a flat tax on short-term funding. And a flat tax on short-term funding is not a good solution in this environment!

Our conclusions can then be summarized as follows:

Proposition 7 *If gambling incentives constitute the only source of heterogeneity across banks, a stable funding requirement $\bar{x} = \bar{x}^{**}$ implements the socially efficient allocation while no flat-rate tax on short-term funding can do it. A liquidity requirement has the same shortcomings as in the baseline model and is, then, either ineffective (if $\delta = 0$) or very similar (but with larger deadweight costs) than the flat-rate tax solution (if $\delta > 0$).*

As a graphical illustration of these results, Figure 3 develops an example based on extending and modifying the parameterization used in previous figures. We fix $\theta_1 = 1$ and introduce θ_2 as in (22) under the assumption that it is uniformly distributed over the interval $[0, 1]$. The new figure depicts, as a function of θ_2 , the unregulated equilibrium allocation $x^{ee}(\theta_2)$, the socially optimal allocation $x^{**}(\theta_2)$ (which can be attained by limiting banks short-term funding to $\bar{x}^{**} = 1.154$) and the allocation that would arise under the best flat tax on short-term funding.²⁴ In this example, social welfare in the unregulated equilibrium is 95.4% of its first-best level, while under the flat Pigovian tax it reaches 99.8% of the first-best level. This good second-best performance is due to the fact that, as illustrated in the figure, the downward shift in the use of short-term funding by all banks induced by the flat tax (which corrects for the “average” externalities of the unregulated equilibrium) is of relatively larger importance than the “differences” in externalities produced by the heterogeneity in θ_2 .

8.2 Generalizing the analysis

The analysis of the general case in which both θ_1 and θ_2 exhibit significant variation across banks is complicated and unlikely to yield clear-cut analytical results, if anything because first best efficiency will not be generally attainable using instruments that are not explicitly contingent on θ_1 or θ_2 . The analysis of simple instruments will necessarily be based on their second-best performance, that will have to be checked numerically for the general case.

²⁴In our example, the trade-offs involved in determining the best flat tax rate $\bar{\tau}^{SB}$ are linear, so $\bar{\tau}^{SB}$ (which equals 0.26) coincides with the average value of the efficient Pigovian schedule $\tau^{**}(\theta_2)$ characterized in equation (28).

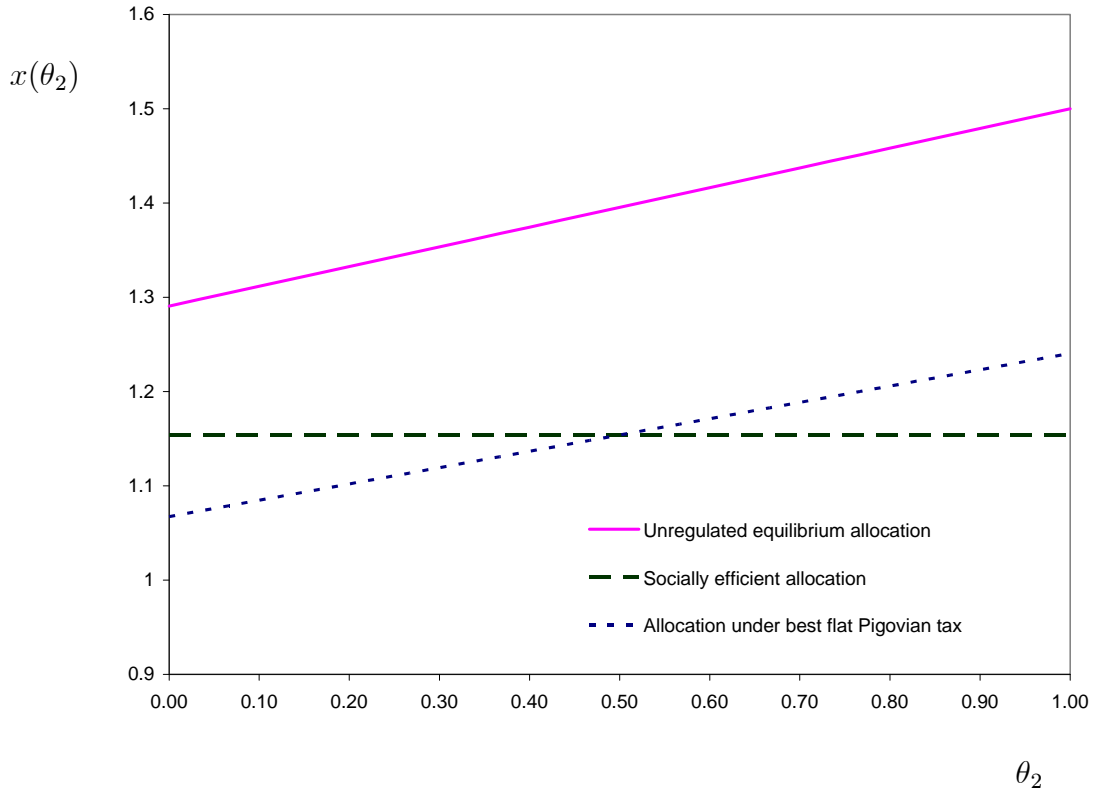


Figure 3. Allocations when banks differ in gambling incentives

Using a continuity argument and building on the polar cases already analyzed above, we can say that a flat tax on short-term funding will tend to perform better than a stable funding requirement if θ_1 is the dominant source of heterogeneity, i.e. if it has ample variation and, specifically, sufficient density at its upper tail, producing sufficiently many banks with value-generating motives to use short-term funding at a larger scale. The opposite will be true if θ_2 is the dominant source of variation, in this case producing sufficiently many banks whose main reason for wanting to use short-term funding at large scale is risk-shifting. For instance, if the banking system had a small group of gambling banks and an ample majority of non-gambling banks, a stable funding requirement might be helpful to control the otherwise excessive short-term funding that the former would like to use.

But continuing with these easy to visualize examples, one can also anticipate possible advantages from combining the instruments. Suppose, in particular, that there is a group of

banks which are essentially diverse in θ_1 and with no specially severe incentives to gamble, like in our main case, but there are also banks, at the top of the overall unconditional distribution of θ_1 , which are also heterogeneous in their incentives to gamble (perhaps due to “too big to fail” problems and bail-out expectations). Then, in addition to having a tax on short-term funding that graduates the contribution to systemic externalities of the banks in the non-gambling group (and also, though imperfectly, the banks in the gambling group), it might be socially valuable to introduce a complementary quantitative limit on short-term funding (say, via a stable funding requirement) so as to further control the externalities caused by the group of gambling banks.

As a simple illustration of this possibility, we consider the following variation on the parameterization described in prior examples: a measure 0.5 of non-gambling banks have θ_1 uniformly distributed in the interval $[0, 1]$ and $\theta_2 = 0$, while another measure 0.5 of gambling banks have $\theta_1 = 1$ and θ_2 uniformly distributed in the interval $[0, 1]$. Figure 4 depicts social welfare as a function of the quantitative limit \bar{x} on short-term funding for various values of the flat rate $\bar{\tau}$ of a Pigovian tax on short-term funding. The solid curve is the one whose maximum identifies the welfare maximizing combination of instruments, which in the example has $\bar{\tau} = 0.22$ and $\bar{x} = 1.59$. The other two curves are generated by varying \bar{x} , while keeping $\bar{\tau}$ fixed at values of 0 and 0.44, respectively. The flat sections of the curves identify situations in which the quantitative limit \bar{x} is large enough not to be binding for any bank. The optimal \bar{x} in Figure 4, identified by a big dot, is not far from the non-binding level, while when the flat Pigovian tax is zero (the dashed curve), the optimal quantitative limit (though numerically very similar) is binding for a larger set of banks. All in all, the example confirms that price and quantity instruments can only imperfectly substitute each other and, in general, using both will be optimal.

9 Dealing with heterogeneity in systemic importance

Suppose that factors such as interconnectedness, lack of substitutability, centrality or size makes some banks more “systemically important” than other in the very sense that the per-unit contribution of their short-term funding to the systemic risk measure X is larger than for other banks. Suppose in particular that systemic importance is captured by

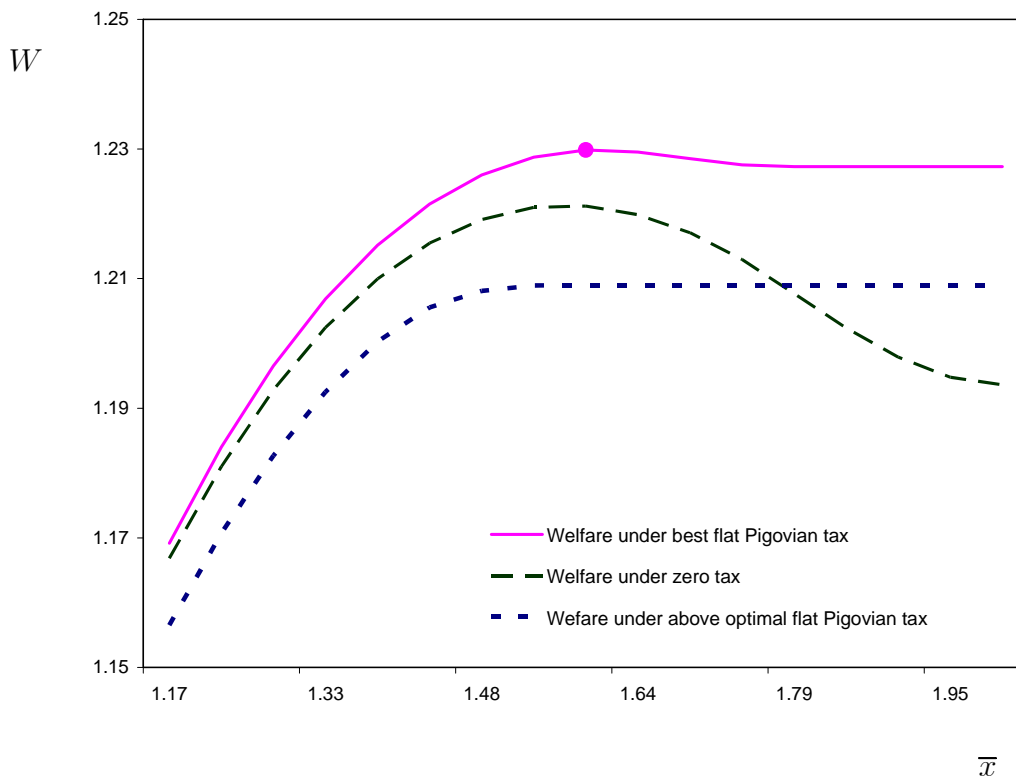


Figure 4. Social welfare under various combinations of instruments

a new dimension of heterogeneity, θ_3 , which only enters significantly into the equations of the economy through the following extended measure of systemic risk:

$$X = \int_0^1 \int_0^1 w(\theta_3)x(\theta_1, \theta_3)f(\theta_1, \theta_3)d\theta_1d\theta_3, \quad (29)$$

where $w(\theta_3)$ is the systemic risk weight of the banks of class θ_3 and $x(\theta_1, \theta_3)$ denotes the short-term funding used by banks characterized by the pair (θ_1, θ_3) . This quantity is written as a function θ_3 to account for the possibility that regulation is made contingent on θ_3 and, through it, systemic importance has an impact on the short-term funding used by each bank.

Extending our characterization of competitive equilibria (unregulated or with taxes) and the socially optimal allocation to deal with this case is immediate. Moreover, it can be shown that decentralizing the socially optimal allocation as a competitive equilibrium with taxes will only require setting $\tau(\theta_3) = w(\theta_3)\tau^*$, where $\tau^* = E_z(\varepsilon(x^*(z), z))c'(X^*)$ is a reference rate

set exactly like in (12), except because z should now be interpreted as the vector (z_1, z_3) of individual bank characteristics. So the presence of heterogenous systemic importance simply leads to the need to consider each bank’s systemic importance measure $w(\theta_3)$ in scaling up the reference tax rate τ^* . But $\tau(\theta_3)$ preserves the key property of being not directly dependent on the individual value of each bank’s lending opportunities as measured by θ_1 .

10 Conclusions

We have developed a formal analysis of the relative performance of realistic price-based and quantity-based approaches to the regulation of systemic externalities associated with banks’ short-term funding. The analysis suggests that, if the return to the lending (or investment) activities undertaken by the banks using this funding is heterogeneously distributed across banks (or, similarly, over time), a Pigovian tax on short-term funding will dominate a net stable funding ratio or a liquidity coverage ratio. If some (poorly capitalized or low charter value) banks have strong gambling incentives and expand their activity as a way to shift risk to outside stakeholders (e.g. the deposit insurer), quantity requirements may have better properties. In general terms, an optimal regulatory design may combine price and quantity-based instruments, and the emphasis on each of them will depend on what is the dominant dimension of heterogeneity across banks (or variation over time).

Going beyond the pure regulation of short-term funding, capital requirements—the most important regulatory instrument in banking—can be seen as a way to directly influence gambling incentives (and the extent to which banks differ in this dimension). Strengthening capital requirements, by ensuring shareholders internalize a larger part of the lower tail of the returns generated by the banks, will tend to produce a smaller measure of banks with strong inclinations for risk-shifting. In this sense, our results suggests that in a scenario with stronger capital regulation, such as that envisaged after the full implementation of Basel III, there will be even greater room for having a Pigovian-style tax, levy or charge on short-term funding as part of the second best regulatory mix.

References

- Acharya, Viral, and Ouarda Merrouche (2009), “Precautionary Hoarding of Liquidity and Inter-Bank Markets: Evidence from the Sub-prime Crisis,” mimeo, NYU Stern (revised August 2010).
- Acharya, Viral, and Sabri Öncü (2010), “The Repurchase Agreement (Repo) Market,” in Viral Acharya, Thomas Cooley, Matthew Richardson, and Ingo Walters, eds., *Regulating Wall Street: The Dodd-Frank Act and the New Architecture of Global Finance*, Chapter 11, Wiley.
- Acharya, Viral, Lasse Pedersen, Thomas Philippon, and Matthew Richardson (2010), “Measuring Systemic Risk,” FRB of Cleveland WP No. 10-02.
- Acharya, Viral, and S. Viswanathan (2011), “Leverage, Moral Hazard, and Liquidity” *Journal of Finance*, 66, 99-138.
- Adrian, Tobias, and Markus Brunnermeier (2009), “CoVaR,” Federal Reserve Bank of New York Staff Reports No. 348.
- Allen, Franklin, Anna Babus and Elena Carletti (2010), “Financial Connections and Systemic Risk,” mimeo, Wharton.
- Allen, Franklin, and Douglas Gale (2000), “Financial Contagion,” *Journal of Political Economy* 108(1), 1-33.
- Basel Committee on Banking Supervision (2010), “Basel III: International Framework for Liquidity Risk Measurement, Standards and Monitoring” Bank for International Settlements, December.
- Brunnermeier, Markus (2009), “Deciphering the Liquidity and Credit Crunch 2007-08,” *Journal of Economic Perspectives* 23(1), 77-100.
- Brunnermeier, Markus, and Martin Oehmke (2010), “The Maturity Rat Race,” mimeo, Princeton University.
- Calomiris, Charles, and Charles Kahn (1991), “The Role of Demandable Debt in Structuring Optimal Banking Arrangements,” *American Economic Review* 81(3), 497-513.
- Diamond, Douglas, and Philip Dybvig (1983), “Bank Runs, Deposit Insurance, and Liquidity,” *Journal of Political Economy* 91(3), 401-419.

- Diamond, Douglas, and Raguram Rajan (2001), "Liquidity Risk, Liquidity Creation and Financial Fragility: A Theory of Banking," *Journal of Political Economy* 109(2), 287-327.
- Farhi, Emmanuel, and Jean Tirole (2010), "Collective Moral Hazard, Maturity Mismatch and Systemic Bailouts," forthcoming, *American Economic Review*.
- Financial Services Authority (2009), *Strengthening Liquidity Standards*, Policy Statement 09/16, October.
- Glaeser, Edward, and Andrei Shleifer (2001), "A Reason for Quantity Regulation," *American Economic Review: Papers and Proceedings* 91(2), 431-435.
- Gorton, Gary (2009), "Slapped in the Face by the Invisible Hand: Banking and the Panic of 2007", paper prepared for the Federal Reserve Bank of Atlanta's 2009 Financial Markets Conference, May.
- Gorton, Gary, and Andrew Metrick (2010), "Regulating the Shadow Banking System," mimeo, Yale University.
- Hellmann, Thomas, Kevin Murdock, Joseph Stiglitz (2000), "Liberalization, Moral Hazard in Banking, and Prudential Regulation: Are Capital Requirements Enough?," *American Economic Review* 90(1), 147-165.
- Holmstrom, Bengt, and Jean Tirole (1997), "Financial Intermediation, Loanable Funds, and the Real Sector," *Quarterly Journal of Economics* 112(3), 663-691.
- Huang, Rocco, and Lev Ratnovski (2011), "The Dark Side of Bank Wholesale Funding," *Journal of Financial Intermediation*, 20(2), 248-263.
- Huberman, Gur, and Rafael Repullo (2010) "Moral Hazard and Debt Maturity," mimeo, CEMFI.
- Jeanne, Olivier, and Anton Korinek (2010), "Managing Credit Booms and Busts: A Pigouvian Taxation Approach," mimeo, John Hopkins University.
- Jensen, Michael, and William Meckling (1976), "Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure," *Journal of Financial Economics* 3(4), 305-360.

- Kaplow, Louis, and Steven Shavell (2002), "On the Superiority of Corrective Taxes to Quantity Regulation," *American Law and Economics Review* 4(1), 1-17.
- Keeley, Michael C. (1990), "Deposit Insurance, Risk, and Market Power in Banking," *American Economic Review* 80(5), 1183-1200.
- Laeven, Luc, and Fabian Valencia (2010), "Resolution of Banking Crises: The Good, The Bad, and The Ugly," IMF WP No. 10/146.
- Lorenzoni, Guido (2008), "Inefficient Credit Booms," *Review of Economics Studies* 75, 809-833.
- Martin, Antoine, David Skeie, and Ernst Ludwig von Thadden (2010), "Repo Runs," mimeo, University of Mannheim.
- Perotti, Enrico, and Javier Suarez (2009), "Liquidity Insurance for Systemic Crises," CEPR Policy Insight No. 31, February.
- Poole, William (1970), "Optimal Choice of Monetary Policy Instruments in a Simple Stochastic Macro Model," *Quarterly Journal of Economics* 84(2), 197-216.
- Repullo, Rafael (2004), "Capital Requirements, Market Power, and Risk-Taking in Banking," *Journal of Financial Intermediation* 13, 156-182.
- Rochet, Jean-Charles, and Xavier Vives (2004), "Coordination Failures and the Lender of Last Resort: Was Bagehot Right After All?," *Journal of the European Economic Association* 2(6), 1116-1147.
- Segura, Anatoli, and Javier Suarez (2011), "Liquidity Shocks, Roll-over Risk, and Debt Maturity," CEPR DP 8324, April.
- Wagner, Wolf (2010), "Diversification at Financial Institutions and Systemic Crises," *Journal of Financial Intermediation* 19(3), 373-386.
- Weitzman, Martin L. (1974), "Prices vs. Quantities," *Review of Economic Studies* 41(4), 477-491.