WP/18/123

IMF Working Paper

Does Financial Tranquility Call for Stringent Regulation?

by Deepal Basak and Yunhui Zhao

INTERNATIONAL MONETARY FUND

 WP/18/123

IMF Working Paper

Monetary and Capital Markets Department

Does Financial Tranquility Call for Stringent Regulation?

Prepared by Deepal Basak and Yunhui Zhao[1]

Authorized for distribution by Martin Cihak

May 2018

Abstract

Consistent with the Minsky hypothesis and the "volatility paradox" (Brunnermeier and Sannikov, 2014), recent empirical evidence suggests that financial crises tend to follow prolonged periods of financial stability and investor optimism. But does financial tranquility always call for more stringent regulation over time? We examine this question using a simple portfolio choice model that features the interaction between learning and externality. We evaluate the potential of a macroprudential policy to restore efficiency, and characterize the necessary and sufficient condition for the countercyclicality of the optimal regulation/macroprudential policy. Our paper implies that policymakers should not only consider the cyclical indicators "on the surface" (for example, credit growth), but also closely examine the deep structural change of the resilience of the system. The paper also highlights the importance of assigning the macroprudential policy function to independent agencies with technical expertise.

JEL Classification Numbers: G01, G18, G28

Keywords: Financial stability, Financial regulation, Learning, Externality, Macroprudential

Authors' E-Mail Addresses: deepal_basak@isb.edu; yzhao@imf.org

---

[1] Basak and Zhao are affiliated with the Indian School of Business and the IMF, respectively. We are extremely grateful for the helpful discussions with Viral Acharya, Tobias Adrian, Jorge Alvarez, Anil Ari, Woon Gyu Choi, Martin Cihak, Luis Cortavarria-Checkley, Udaibir Das, Alan Xiaochen Feng, Daniel Garcia-Macia, Gaston Gelos, Dan Greenwald, Javier Hamann, Xing Hong, Aaditya Iyer, James Morsink, Saptarshi Mukherjee, Maurice Obstfeld, Miguel Segoviano, Bowen Shi, Ennio Stacchetti, Di Wu, participants at NYU Stern Finance Student Seminar in Spring 2016, participants at the IMF MCM Policy Forum in July 2017, numerous internal reviewers at the IMF, and especially Alexander Murray, Tao Zha, and Divya Kirti. Mayur Choudhary has provided excellent research assistance. All remaining errors are our own.

Contents

List of Figures

I Introduction

A view associated with Minsky (1992) holds that prolonged periods of economic growth and financial stability lead investors to become overconfident about the prospects for future growth and stability. Overconfidence leads to the buildup of risks in the financial system as investors increase their leverage and invest in riskier portfolios. This buildup of risk sets the stage for a financial crisis. This view is confirmed by the "volatility paradox" proposed by Brunnermeier and Sannikov (2014), a phenomenon where low-risk environments are conducive to higher equilibrium leverage and greater buildup of systemic risk.[1] Consistent with this view, recent empirical evidence suggests that financial crises tend to follow prolonged periods of financial stability during which credit supply expands and asset prices rise (Borio, 2012; Dell'Ariccia and others, 2012; Drehmann and others, 2012; Danielsson and others, 2016; Espinoza and others, 2017).[2] [3]

How should this phenomenon be accounted for in the design of macroprudential policies, if at all? And does financial tranquility always call for more stringent regulation/macroprudential policies over time?[4] Our paper answers these questions using a simple learning model, and offers two contributions to the literature. Our *first* contribution is to illustrate the excessive risk-taking in financial markets through the interaction between learning and externality. A longer history of financial tranquility (that is, the absence of the crisis) builds the confidence of investors and leads each investor to take a larger position in the risky asset. This in turn raises the *aggregate* risky asset position and the probability of a systemic crisis, imposing a higher negative externality that the individual investor does not internalize (under the condition laid out in our paper). Hence, the learning process strengthens the negative externality, and aggravates the excessive risk-taking problem and the constrained inefficiency.

Our *second* contribution is to provide a simple framework to assess the efficiency of macroprudential regulation. We characterize the necessary and sufficient condition for the countercyclicality of optimal regulation. Our paper implies that optimal macroprudential policies are *not* always countercyclical; instead, they depend on the trade-off between the resilience effect and risk-taking effect modeled in our paper. Therefore, macroprudential policymakers should not only consider the cyclical indicators "on the surface" (for example, credit growth), but also closely examine the deep structural change of the resilience of the system, such as the dynamics of the industrial structure (for example, the economy's dependence on one par-

[1] One difference between our model and Brunnermeier and Sannikov (2014)'s is that the former focuses on the role of learning in a simple portfolio model, whereas the latter highlights the role of liquidity in a full macroeconomic model.

[2] Recessions associated with systemic financial crises tend to be particularly deep and long-lasting. Laeven and Valencia (2010) document that the median output loss of the recent financial crisis is 25 percent. Reinhart and Rogoff (2009) observe that in financial crises the unemployment rate increases by 7 percentage points and remains high for over four years on average.

[3] IMF (2017a) also warns that the longer booms last and the larger credit grows, the more dangerous they become.

[4] "More stringent" is meant as a statement about change over time, not about comparisons with other models or with real-world policy; that is, "more stringent" than when the market had not been tranquil for that long (the market confident was not so high).

ticular industry). The paper also highlights the importance of assigning the macroprudential policy function to independent agencies with technical expertise.

Summary of Model and Policy Implications

At any point in time, macroprudential policies to constrain investor risk-taking are justified if investors fail to internalize the effect of their portfolio choices on the probability of a systemic financial crisis (defined as an event where the returns of all financial assets are substantially reduced and such effects may spill over to the real sector). However, both investors and policymakers face uncertainty about the size of that externality, that is, about the extent to which investors' risky portfolio choices may raise the likelihood of a crisis. A long period of financial tranquility may induce investors to take more risk (risk-taking effect, or input effect, as in Figure 1), but it also provides evidence that the transmission mechanism between investor risk-taking and systemic crisis may not be as strong as previously believed (resilience effect, or transmission effect). From the perspective of a macroprudential regulator, it is therefore unclear *a priori* whether macroprudential policies should be made more or less stringent over a period of prolonged financial tranquility. The answer depends on the relative magnitudes of the two effects just described: the increase in investor risk-taking and the improvement in the regulator's perception of the resilience of the financial system.

In this paper, we formalize this intuition using a simple model of portfolio choice and Bayesian learning. In the model, investors have access to a safe asset and to a risky asset that, in addition to its conventional risk properties, may expose the financial sector to the risk of a systemic crisis. The investors do not know the extent to which the risky asset exposes the financial sector to this systemic crisis risk, but they can learn about it from the financial system's history of performance. A longer history of financial tranquility (that is, the absence of a crisis) builds up investors' confidence and leads them to take larger positions in the risky asset. However, a larger aggregate risky asset position raises the probability of a systemic crisis in the event that the investor confidence is mistaken.

Note that the learning process in our model is subject to a "complacency trap" in the following senses. First, the investors will become more confident and invest more in the risky asset as long as the crisis did not occur in the previous period. That is, the investors will revise down their confidence only after the crisis has actually hit ("*Cry only when death is staring one in the face*"). Second, the investors have a short memory about the crisis and will revise down their confidence in only one period, which is the first period right after the crisis; from the second post-crisis period on, as long as there is no crisis in the previous period, the investors will again become more confident and invest more in the same risky asset that has only recently led to the crisis ("*Once on shore, pray no more*"). Despite the simplicity of our model, these features are consistent with recent evidence that financial markets are rapidly increasing the investments in some credit products widely blamed for exacerbating the recent global financial crisis.[5]

[5] For example, *Financial Times* reported on August 23, 2017, that: "Hedge funds are embracing an esoteric credit product widely blamed for exacerbating the financial crisis a decade ago, as low volatility and near record prices for corporate debt tempt them into riskier areas to seek higher returns. The market for 'bespoke tranches' — bundles of credit default swaps that are tied to the risk of corporate defaults — has

Within this simple framework, we first study the relationship between investors' learning process and the degree of financial market inefficiency. The laissez-faire equilibrium is inefficient because each trader imposes a negative externality through the effect of his or her or her risky asset position on the systemic crisis probability. The equilibrium is characterized by excessive risk-taking and excessive financial instability; at a point in time, equilibrium risky asset positions exceed the positions chosen by a constrained benevolent planner who must also learn about the true systemic crisis probability. In our first main result, we derive the (necessary and sufficient) condition on the crisis probabilities under which rising investor confidence strengthens the negative externality, thereby aggravating the excessive risk-taking problem and increasing the inefficiency of the laissez-faire equilibrium.

We then turn to the question of prudential policy. We show that, at a point in time, a macroprudential regulator can reestablish efficiency by means of a capital income tax set at an appropriate level.[6] Our second main result is that, under the same condition derived in the first result, the optimal tax rate rises as the degree of investor confidence (and, hence, the degree of market inefficiency) rises.

The necessary and sufficient condition we derive captures precisely the intuition outlined earlier. As a period of tranquility persists, the constrained planner becomes more confident in the fundamental stability of the financial system and comes to believe that the transmission channel from investor risk-taking to systemic crisis is weaker (that is, the crisis probability per "unit" of risk-taking is smaller). This reduces the severity of the externality in the eyes of the planner, and leads the planner to increase its risky asset position relative to the investors' laissez-faire risky asset position (resilience effect, or transmission effect). At the same time, however, a higher risky asset position increases the severity of the externality directly (by "inputting" more risks into the system), leading the planner to rein in risky asset investment relative to that of the private investors under laissez-faire (risk-taking effect, or input effect). The overall effect of rising confidence on the degree of market inefficiency (and hence on the optimal stringency of macroprudential policy) depends on the relative sizes of these two countervailing effects. The condition derived in our paper captures this comparison. We deliberately keep our model very simple in order to focus on this point, which is obscured in

more than doubled in the first seven months of 2017." IMF (2017b) also warned that "Risk appetite has grown markedly as near-term stability risks have declined." In a speech at Jackson Hole in late August 2017, the Federal Reserve chair Janet Yellen also warned that memories of the last crisis "may be fading." (*Financial Times*, August 25, 2017).

[6] This is a form of Pigouvian taxation, and is similar to the systemic risk taxation proposed by Acharya and others (2010). For convenience, the capital income tax in our paper is levied based on the gross investment income rather than net investment income (that is, investment return), as in the standard capital income tax. Our derivations show that our results apply if we switch to the standard definition instead. For discussions of the benefits and costs associated with the capital income tax (under the standard definition), see Gordon and others (2004) and the references therein. Our results also apply if we use financial transaction tax instead. For discussions of the financial transaction tax, see Adam and others (2015).

other papers that feature more complex financial market models.[7] [8]

These arguments highlight that optimal macroprudential policies are not always countercyclical. Instead, they depend on the trade-off between the above resilience effect and risk-taking effect. Therefore, policymakers should not only consider the cyclical indicators "on the surface" (for example, credit growth), but also closely examine the deep structural change of the resilience of the system (for example, the high credit growth may be driven by lendings to non-oil industries in an oil-exporting economy; by diversifying the economy's structure, the credit growth may have fundamentally improved the resilience of the system). This view is consistent with the evidence in Dell'Ariccia and others (2012), who document that more than a third of credit booms are not followed by economic underperformance.[9]

These arguments also highlight a key challenge faced by macroprudential regulators: the optimal policy depends on the resilience of the financial system, but it is difficult to know the true resilience. This opens the door to "inefficient deregulation" or "inefficient regulation." On the one hand, a regulator who overestimates the resilience of the system will tend to reduce the stringency of macroprudential regulation as financial tranquillity persists, and this will induce a buildup of financial risk that the regulator would not tolerate if s/he knew the system's true resilience. Some commentators and policymakers (such as Alan Greenspan) have argued that this occurred in the United States in the 1990s and 2000s.[10] On the other hand, a regulator who underestimates the resilience of the financial system may repress financial activities needlessly. This could prevent the uptake of financial innovations that might actually be effective at delivering value to investors (for example, by diversifying risk more effectively). Thus, the regulatory challenge is especially stark for innovative financial industries, such as the rapidly growing online finance industries in many countries.

To further connect our paper with the practice of macroprudential policymaking, we also

---

[7] In particular, we do not model the mechanism by which investment in the risky asset may lead to a systemic crisis. Our focus is on the interaction between investors' learning process and the degree of excessive risk buildup, not on the underlying reasons for the crisis. Thus, rather than focus on a particular model of financial crises, we exogenously specify the crisis probability as a function of the aggregate risky asset position and ask how the dynamics of financial market efficiency (or inefficiency) depends upon the properties of this function (together with investors' learning process). Our framework may be regarded as a reduced-form version of some more fully developed models of financial crisis (for example, Biais and others, 2015; Gertler and Kiyotaki, 2015).

[8] All our results also hold if the underlying state is not fixed and follows a Markov process instead, that is, a good state may become bad or a bad state may become good. The main difference from the fixed-state case is that now agents may never learn the true state. See Section V for more details.

[9] Using cross-country data from primary debt capital markets, Kirti (2018) shows that lending standards help separate good credit booms from bad credit booms contemporaneously.

[10] For example, in 2013 Robert Samuelson argued that "many of the institutions that came to grief — banks, investment banks — were regulated. But regulators shared the optimistic consensus concerning the economy's transformation. Complacency made regulation permissive. It was the Great Moderation that gave us the financial crisis and Great Recession." The famous admission of Alan Greenspan in 2008 that s/he had "made a mistake in presuming that the self-interest of organizations, specifically banks, is such that they were best capable of protecting shareholders and equity in the firms" was an *ex post* acknowledgement that regulators had overestimated the resilence of the financial system in the runup to the crisis.

consider the problem of an unconstrained planner who has full information about the true state of the financial sector. The motivation for this assumption is that in practice, macroprudential policymakers are likely to know more about the true state (strong or fragile) of the system than individual investors or the uninformed/constrained "planner." This is because macroprudential policymakers spend an enormous amount of time and effort trying to assess the resilience of the system, including through stress tests and indicators of balance sheet vulnerabilities. As a simple approximation, policymakers are modeled by assuming that they have full information about the true state. We show that when the true state is fragile, the choice of risky investment made by a constrained planner is inefficiently higher than that by an unconstrained planner. Therefore, it is desirable to assign the macroprudential policy function to independent agencies with technical expertise that allows them to gauge the underlying true resilience of the financial system. This is consistent with the viewpoints of IMF (2011, 2014).

In the remainder of this introduction, we discuss the related literature and situate our work within it. After that, the rest of the paper is structured as follows. Section II outlines the model, characterizes the competitive equilibrium and the social planner's solution, and presents our first main result. Section III characterizes the optimal capital income tax, presents our second main result, and discusses the possibility of inefficient regulation. Section IV presents some parametric examples. Section V further examines the dynamics of learning, and extends our model to where the true state follows a Markov process. Section VI considers the case of an unconstrained planner and its policy implications. Section VII concludes. The Appendices collect the figures and some technical proofs.

Related Literature

Our paper is related to four strands of literature. The first is the literature on countercyclical macroprudential regulation. It has been widely believed that desirable macroprudential policies should be countercyclical: during economic booms, the financial sector is lenient in monitoring borrowers (Jimenez and Saurina, 2006) or is plagued by overborrowing (Bianchi, 2011), so more stringent macroprudential policies should be employed to build up buffers and curb excessive borrowing; during economic busts, more loose macroprudential policies should be employed to help financial institutions absorb losses. Such more stringent macroprudential policies can be higher capital requirements (Gordy and Howells, 2006; Gordy, 2009; Drehmann and others, 2010; Bank for International Settlements, 2010; Basel Committee on Banking Supervision, 2010a, 2010b), higher provisioning (Packer and Zhu, 2012; Fernandez de Lis and Garcia-Herrero, 2012; Jiménez and others, 2017), higher borrowing cost (Bianchi, 2011), capital inflow taxation (Jeanne and Korinek, 2010), or lower caps on loan-to-value and debt-to-income ratios (Wong and others, 2011; Krznar and Morsink, 2014).[11] Our pa-

[11] A related literature discusses some practical issues while implementing countercyclical regulations. For example, in the context of countercyclical capital requirements, Kowalik (2011) points out that the effectiveness of countercyclical regulations depends on how they are implemented, and finds that the rule-based approach has more advantages than the discretion-based approach. In addition, McCoy (2016) points out five challenges to the successful execution of countercyclical regulation: the federal government's data collection initiatives; how to track threats from new financial products and respond to them; how to justify intervention through rules when risks are small; regulatory capture and inertia; and the likelihood of regulatory arbitrage.

per emphasizes that the optimal cyclical adjustment of macroprudential policy — that is, whether it should become more or less stringent as a period of financial tranquility persists — is *a priori* unclear under uncertainty about the resilience of the financial system.

Second, it is related to the literature on learning. One part of this literature focuses on learning about an individual manager or firm and examines the issue of manager compensation in the finance industry.[12] Another, more relevant part of the literature focuses on learning about aggregate parameters of the financial industry. Biais and others (2015) consider the dynamics of an innovative sector in which agents learn the sector's exposure to negative shocks and managers' risk management efforts are subject to a moral hazard problem. They show that rising confidence leads to lower risk-abatement effort by managers. Boz and Mendoza (2014) construct a model with a collateral constraint, in which financial innovation acts as a structural change that introduces a regime with a higher leverage limit. In the model, learning about the risk of a new financial environment predicts large increases in household debt. Bianchi and others (2012) construct a macroeconomic model in which agents learn about the transition probabilities between states with tight and loose borrowing constraints. In this framework, they examine how the effectiveness of macroprudential policy depends on the information available to the planner.

Perhaps the paper most similar to ours is Bhattacharya and others (2015). In their framework, the economy switches between two aggregate states — "good" and "bad" — and asset returns in a period depend on the realization of the state. Investors do not know the true probability that the good state will be realized in the next period, but they know that it takes one of two possible values and they use the history of asset returns to update their beliefs about it. Investors have access to two assets, one of which is riskier than the other (that is, has a larger return variance) under any probability distribution. They finance their investment using their net worth and by issuing debt in the credit market. The authors use a three-period version of the model to study the implications of learning for investors' leverage and risk-taking behavior.

While the modelling approach of Bhattacharya and others (2015) bears many similarities to ours, there are important differences that reflect different research priorities. Their interest in the interaction between portfolio risk and the leverage cycle leads them to include debt finance and micro-founded credit market transactions in their model. Moreover, the complexity of their framework makes their welfare analysis rather opaque; they propose macroprudential policies that they show to be welfare-enhancing using simulations, but they cannot provide much intuition for those results. In contrast, we focus specifically on the implications of learning, over many periods, for the degree of excessive risk-taking (and hence on the optimal stringency of macroprudential regulation). Our simpler framework allows us to obtain analytical characterizations of our welfare results and to highlight the countervailing effects of rising investor risk-taking and increasing regulator confidence in the stability of the financial system over periods of financial tranquility. As a side note, our model does not rely on leverage to generate the crisis; instead, the crisis is generated through the misallocation of resources (for example, human capitals) among the real (conventional) and financial

[12] Examples include Diamond (1991), Berk and Green (2004), Noe and Rebello (2012), Axelson and Bond (2015), Celerier and Vallee (2015), and Bolton and others (2016).

(innovative) sectors.

Third, our paper is related to the literature on inefficient risk-taking arising from externalities in financial markets. Using a model of collateralized international borrowing, Jeanne and Korinek (2010) considers a negative externality that arises as declining collateral values, tightening financial constraints and falling consumption mutually reinforce each other. Bianchi (2011) identifies a systemic credit externality, which arises because private agents fail to internalize the financial amplification effects of carrying a large amount of debt when credit constraints bind. Focusing on economies with nominal rigidities in goods and labor markets (and subject to constraints on monetary policy), Farhi and Werning (2016) identify an aggregate demand externality: ex post, the distribution of wealth across agents affects aggregate demand and output; but ex ante, these effects are not internalized in private financial decisions by atomistic agents.[13] Our paper also has externality since each atomistic investor does not internalize the negative impact of his or her or her risky investment on others (through increasing the probability of crisis). One difference between our paper and these papers is that we have learning in our model, and focus on the interaction between learning and externality.[14]

Finally, the paper is related to the broader literature on macroprudential policies.[15] Besides the aforementioned macroprudential literature that addresses the time dimension of the systemic risk (that is, the dynamics of the macroprudential policies over the boom-bust cycle), our paper is also related to macroprudential literature that addresses the cross-section dimension of the systemic risk at any given point in time. Financial institutions manage their own risks, but do not consider their impact on the system as a whole, imposing a negative externality on other institutions (Acharya, 2009; Acharya and others, 2009; Acharya and others, 2010). Systemic risk can also be understood as a network externality resulting from contagion effects (Acemoglu and others, 2013). To cope with this externality, policymakers can use either taxation (reducing the gap between public and private costs of systemic risk) or regulation (imposing direct restrictions and requirements on financial institutions) (Masciandaro and Passarelli, 2013). Several studies advocate for a taxation of systemic risk (Acharya and others, 2009; Acharya and others, 2010; Zlatic and others, 2015). To implement the systemic risk tax, we need to properly quantify systemic risk. To this end, Segoviano (2006) proposes the CIMDO approach; several studies propose measures that focus on statistics of losses, accompanied by a potential shortfall during periods of synchronized behavior when many institutions are simultaneously distressed (Acharya and others, 2009;

[13] Farhi and Werning (2016) also propose an extended framework to incorporate both pecuniary externality and aggregate demand externality, and characterize the optimal macroprudential policy that can correct for these externalities.

[14] We follow this literature in viewing excessive risk-taking as the result of a financial market externality. Other mechanisms that can generate excessive risk-taking include neglected disaster risk (Gennaioli and others, 2012, 2013; Baron and Xiong, 2016), extrapolative expectations (Greenwood and Shleifer, 2014; Barberis and others, 2015; Barberis and others, Forthcoming), diagnostic expectations (Bordalo and others, 2016), and "this-time-is-different" thinking (Reinhart and Rogoff, 2009). The key difference is that our mechanism does not rely on the behavioral assumptions.

[15] Our paper falls into the macroprudential literature also in the sense that the macroprudential orientation treats the aggregate risk as an endogenous variable that depends on the collective behavior of all financial institutions rather than being exogenously given by the market (Kahou and Lehar, 2017).

Adrian and Brunnermeier, 2016; Brownlees and Engle, 2012; Huang and others 2012).[16]

II Benchmark Model

A Decentralized Problem

A.1 Set up and Learning

There are a continuum of investors, each indexed by $i \in [0, I]$ and endowed with 1 unit of investment good. Each investor can invest in two types of assets: a safe asset in the real sector, and a risky asset in the financial sector. We can also view both of these assets as financial assets: one is a "safe," conventional financial asset; the other is a risky, innovative financial asset. The second view allows us to apply our model to study financial innovation. Time t is discrete and infinite, that is, $t = 1, 2, 3, ...$ At the beginning of period t, agent i invests $\alpha_{it} \in [0, 1]$ in the risky asset and $1 - \alpha_{it}$ in the safe asset. At the end of period t, agent i consumes and dies/exits.

A systemic financial crisis can occur at any point of time. If the financial crisis does not occur, the safe (real) asset pays μ_S with probability 1, and the risky (financial) asset pays μ_R in expectation, with $\mu_R > \mu_S > 0$. Figure 2 illustrates the payoff structure of the model. If the financial crisis does occur, the safe asset pays $\tau_S \mu_S$ with probability 1, and the risky asset pays $\tau_R \mu_R$ with probability 1 where $0 \leq \tau_R \leq \tau_S \leq 1$. Note that τ_S and τ_R capture the post-crisis recovery value of the safe asset and the risky asset, respectively. Moreover, τ_S captures the *spillover* effect from the financial sector to the real sector: if $\tau_S = 1$, then even if the financial crisis hits, the payoff of the real asset is still unaffected, and thus there is no spillover from the financial sector to the real sector; if $\tau_S = \tau_R$, then the spillover is at the highest possible level.

The financial sector can be either strong (in state "G") or fragile (in state "B"). Importantly, the probability of the financial crisis depends on which state the financial sector is in, as well as the aggregage position in the risky asset $\alpha_t \equiv \int \alpha_{it} di$. Denote the crisis probability at time t in state $j \in \{G, B\}$ by $p_j(\alpha_t)$. We make the following assumptions about the $p(.)$ functions:

$$p_j(\alpha_t) \in [0, 1], \forall \alpha_t, \forall j$$
$$p_G(\alpha_t) < p_B(\alpha_t)$$
$$p_j^{'}(\alpha_t) > 0, \forall \alpha, \forall j$$

Let $\pi_t \in [0, 1]$ be the probability at which the period t investors believe the financial sector is strong. At the beginning of period t, all investors "inherit" the belief π_{t-1} in the previous period (that is, π_{t-1} is the prior belief), and then update the belief based on the outcomes they observe from period $t - 1$. That is,

$$\pi_t \equiv P(j = G | C_{t-1})$$

[16] For a comprehensive literature review on systemic risk tax, see Poledna and Thurner (2016).

where

$$C_{t-1} = \begin{cases} 1, & if\, crisis\, occurred\, at\, t-1 \\ 0, & otherwise \end{cases}$$

For any risky investment α_{t-1}, the probability that the crisis occurs (that is, $C_{t-1} = 1$) is higher if the true state is B. That is, $p_B(\alpha_{t-1}) > p_G(\alpha_{t-1})$. Thus, if a crisis occurs in $t-1$, the agents will update their beliefs downwards; if not, they will update their beliefs upwards. Bayesian rule implies that the probability of being a strong sector conditional on no crisis at $t-1$ equals:

$$\begin{aligned} \pi_t(C_{t-1} = 0) &= \frac{P(j = G)P(C_{t-1} = 0|j = G)}{P(C_{t-1} = 0)} \\ &= \frac{\pi_{t-1}[1 - p_G(\alpha_{t-1})]}{\pi_{t-1}[1 - p_G(\alpha_{t-1})] + (1 - \pi_{t-1})[1 - p_B(\alpha_{t-1})]} \\ &> \frac{\pi_{t-1}[1 - p_G(\alpha_{t-1})]}{\pi_{t-1}[1 - p_G(\alpha_{t-1})] + (1 - \pi_{t-1})[1 - p_G(\alpha_{t-1})]} \\ &= \pi_{t-1}, \end{aligned}$$

where the inequality uses $p_B(\alpha_{t-1}) > p_G(\alpha_{t-1})$.

And the probability of being a strong sector conditional on crisis at $t-1$ equals:

$$\begin{aligned} \pi_t(C_{t-1} = 1) &= \frac{\pi_{t-1}p_G(\alpha_{t-1})}{\pi_{t-1}p_G(\alpha_{t-1}) + (1 - \pi_{t-1})p_B(\alpha_{t-1})} \\ &< \frac{\pi_{t-1}p_G(\alpha_{t-1})}{\pi_{t-1}p_G(\alpha_{t-1}) + (1 - \pi_{t-1})p_G(\alpha_{t-1})} \\ &= \pi_{t-1} \end{aligned}$$

Thus, agents' confidence grows as they see a longer history of financial tranquility (that is, no crisis). These arguments establish the following proposition, which is the key driving force to the main results explained later.

PROPOSITION 1: A longer history of tranquility (that is, no crisis) in the financial sector tends to build investor confidence that the sector is strong. That is, π_t strictly increases in the number of no-crisis periods.

A.2 Characterization of the Competitive Equilibrium

The problem of investor i at period t is:

$$\begin{aligned} \max_{\alpha_{it}\epsilon[0,1]} \quad & \pi_t\{(1 - p_G)[(1 - \alpha_{it})\mu_S + \alpha_{it}\mu_R] & + \\ & p_G[(1 - \alpha_{it})\tau_s\mu_S + \alpha_{it}\tau_R\mu_R]\} & + \\ & (1 - \pi_t)\{(1 - p_B)[(1 - \alpha_{it})\mu_S + \alpha_{it}\mu_R] & + \\ & p_B[(1 - \alpha_{it})\tau_s\mu_S + \alpha_{it}\tau_R\mu_R]\} \end{aligned}$$

For convenience, let $\Delta_0 \equiv \mu_R - \mu_S$ be the risky asset's (expected) excess payoff conditional on no-crisis at $t-1$, and $\Delta_1 \equiv \tau_R\mu_R - \tau_S\mu_S$ be its excess payoff conditional on crisis at $t-1$. In order for the risky asset to be meaningful, it is natural to assume $\Delta_0 > 0$. In addition, we make the following assumption:

A1: $\Delta_0 > 0$, and $\Delta_1 < 0$, that is, $\tau_R\mu_R < \tau_S\mu_S$.

Moreover, let

$$\theta_t \equiv \theta(\pi_t, \alpha_t) \equiv \pi_t p_G(\alpha_t) + (1-\pi_t)p_B(\alpha_t)$$

be the unconditional probability of a crisis at period t, which also captures the vulnerability of the system *perceived* by investors (and the constrained social planner, as discussed later). Using these notations, the problem of investor i at period t can be rewritten as:

$$\max_{\alpha_{it}\epsilon[0,1]} \quad [1-\theta(\pi_t,\alpha_t)](\mu_S + \alpha_{it}\Delta_0) + \theta(\pi_t,\alpha_t)(\tau_S\mu_S + \alpha_{it}\Delta_1) \tag{1}$$

The first-order condition (FOC) with respect to α_{it} is given by:

$$[1-\theta(\pi_t,\alpha_t)]\Delta_0 + \theta(\pi_t,\alpha_t)\Delta_1 = 0 \tag{2}$$

Note that although the individual α_{it} does not show up explicitly in the FOC, it shows up implicitly through the aggregate α_t in $\theta(\alpha_t,\pi_t) \equiv \pi_t p_G(\alpha_t) + (1-\pi_t)p_B(\alpha_t)$. The intuition of the above FOC (equation (2)) is as follows: in any interior equilibrium, the aggregate risky asset position α_t must be such that the expected excess payoff of the risky asset (relative to the safe asset) is 0, where the expectation is taken over the occurrence of the crisis. This is actually the no-arbitrage condition between the risky and safe assets. Although each individual investor is indifferent between any $\alpha_{it} \in [0,1]$, the aggregate α_t must satisfy the no-arbitrage condition (2). Rewrite equation (2) as:

$$\frac{\Delta_0}{\Delta_0 - \Delta_1} - \theta(\pi_t,\alpha_t) = 0 \tag{3}$$

For a given value of π_t, it is easy to see that an interior equilibrium exists as long as the function $\theta(\pi_t,\alpha_t)$ crosses the value $\frac{\Delta_0}{\Delta_0-\Delta_1}$ as α_t varies. Since $\theta_t \equiv \pi_t p_G(\alpha_t) + (1-\pi_t)p_B(\alpha_t)$, this amounts to a condition on the functions p_G and p_B. Since we have already assumed that $p_B(\alpha) > p_G(\alpha)\,\forall\alpha$ and that both p_B and p_G are strictly increasing and differentiable (and hence continuous), a sufficient condition for the existence of a unique interior equilibrium conditional on *any* π_t is that $p_B(0) < \frac{\Delta_0}{\Delta_0-\Delta_1}$ and $p_G(1) > \frac{\Delta_0}{\Delta_0-\Delta_1}$. If these conditions are not satisfied, agents could learn their way to a corner solution. We do not examine this possibility in this paper.

The FOC implies that in any (interior) competitive equilibrium, investors will adjust the risky asset position α_t such that θ_t is always a constant regardless of the belief π_t. Moreover, since $\theta_t \equiv \pi_t p_G(\alpha_t) + (1-\pi_t)p_B(\alpha_t)$, $p_G(\alpha_t) < p_B(\alpha_t)$, $p_G^{'}(\alpha_t) > 0$, and $p_B^{'}(\alpha_t) > 0$, this implies that as π_t increases (which puts a higher weight to $p_G(\alpha_t)$), investors will need to increase α_t and thus $p_B(\alpha_t)$ in order to keep θ_t constant. Hence, α_t strictly increases in π_t under a competitive equilibrium. More formally, we have the following proposition:

PROPOSITION 2: The aggregate risky investment α_t in the competitive equilibrium strictly increases in investor confidence π_t, provided $p'_G(\alpha) > 0$ and $p'_B(\alpha) > 0$.

PROOF: See Appendix 1. *Q.E.D.*

Note that only the first-order derivatives of $p_G(.)$ and $p_B(.)$ are involved for this proposition to hold. Based on Lemma 1 and Proposition 1, we can make the following analogue to the set-up in our model. The decision on α_t is like deciding the speed of a car while driving in the fog. As each driver becomes more optimistic that the car ahead is far away from his or her current location (higher π_t), s/he will drive faster. The higher speed has raised the *actual* systemic risk that his or her car will hit the car ahead (higher $p_G(\alpha_t)$, $p_B(\alpha_t)$), but it is still the optimal behavior of each driver because the probability of hitting the car ahead (θ_t) *perceived* by him or her remains constant.

B Constrained Planner's Problem

In this subsection, we consider the problem of a social planner who is subject to the same constraint as the decentralized investors. That is, the planner cannot observe the true state of the financial sector either, and also has to update its belief through the realized occurrence of financial crises.

Since each investor only lives for one period, it is sufficient to consider the planner's problem in one single period. Without loss of generality, the problem of the constrained planner at period t is:

$$\begin{aligned}\max_{\alpha_{it}\epsilon[0,1],\,\forall i} \quad & \pi_t\{(1-p_G)\textstyle\int_i[(1-\alpha_{it})\mu_S+\alpha_{it}\mu_R]di \quad + \\ & p_G\textstyle\int_i[(1-\alpha_{it})\tau_s\mu_S+\alpha_{it}\tau_R\mu_R]di\} \quad + \\ & (1-\pi_t)\{(1-p_B)\textstyle\int_i[(1-\alpha_{it})\mu_S+\alpha_{it}\mu_R]di \quad + \\ & p_B\textstyle\int_i[(1-\alpha_{it})\tau_s\mu_S+\alpha_{it}\tau_R\mu_R]di\}\end{aligned}$$

That is,

$$\begin{aligned}\max_{\alpha_t\epsilon[0,1]} \quad & \pi_t\{(1-p_G)[(1-\alpha_t)\mu_S+\alpha_t\mu_R] \quad + \\ & p_G[(1-\alpha_t)\tau_s\mu_S+\alpha_t\tau_R\mu_R]\} \quad + \\ & (1-\pi_t)\{(1-p_B)[(1-\alpha_t)\mu_S+\alpha_t\mu_R] \quad + \\ & p_B[(1-\alpha_t)\tau_s\mu_S+\alpha_t\tau_R\mu_R]\}\end{aligned}$$

Using the same definitions for θ, Δ_0, and Δ_1, we can rewrite the problem as

$$\max_{\alpha_t\epsilon[0,1]} \quad [1-\theta(\pi_t,\alpha_t)](\mu_S+\alpha_t\Delta_0)+\theta(\pi_t,\alpha_t)(\tau_S\mu_S+\alpha_t\Delta_1) \tag{4}$$

The FOC of the planner is as follows:

$$[1-\theta(\pi_t,\alpha_t)]\Delta_0+\theta(\pi_t,\alpha_t)\Delta_1-\theta_\alpha(\pi_t,\alpha_t)[(\mu_S+\alpha_t\Delta_0)-(\tau_S\mu_S+\alpha_t\Delta_1)]=0 \tag{5}$$

Comparision between equation (2) and equation (5) makes it clear that the difference between the decentralized problem (1) and the constrained planner's problem (4) is that the planner takes into account the impact of the risky asset position on the conditional crisis probabilities $p_G(\alpha_t)$ and $p_G(\alpha_t)$ (equivalently, on the unconditional crisis probability or the *perceived* vulnerability θ_t). Specifically, a higher risky position α_t raises the crisis probability θ_t and lowers the excess payoff of the risky asset as well as the welfare of the investors. This is the negative externality that individual investors fail to consider in the decentralized equilibrium. Specifically, the externality term consists of two elements: the first is the probability term $\theta_\alpha(\pi_t, \alpha_t)$, which is the increase in crisis probability due to the increase in α; the second is the payoff term $(\mu_S + \alpha_t\Delta_0) - (\tau_S\mu_S + \alpha_t\Delta_1)$, which is the reduction in total payoff in case the crisis does occur.

Normalize (5) by dividing by $\Delta_0 - \Delta_1$:

$$\frac{\Delta_0}{\Delta_0 - \Delta_1} - \theta(\pi_t, \alpha_t) - \xi(\alpha_t, \pi_t) = 0 \tag{6}$$

where $\xi(\pi_t, \alpha_t)$ captures the normalized externality of the system, defined as

$$\xi(\pi_t, \alpha_t) \equiv \theta_\alpha(\pi_t, \alpha_t)[\alpha_t + \frac{(1-\tau_S)\mu_S}{\Delta_0 - \Delta_1}] \equiv \theta_\alpha(\pi_t, \alpha_t)\frac{1}{X}$$

Comparision between the normalized FOC of the competitive equilibrium with that of the planner (equation (3) versus equation (6)) makes it clear that while the market adjusts α_t to keep $\theta(\pi_t, \alpha_t)$ constant in equilibrium, the constrained planner adjusts α_t to keep $\theta(\pi_t, \alpha_t) + \xi(\pi_t, \alpha_t)$ constant.

The following proposition characterizes the constrained planner's equilibrium:

PROPOSITION 3: The aggregate risky investment α_t in the constrained planner's equilibrium strictly increases in the planner's confidence π_t, provided $p_B^{'}(\alpha) > p_G^{'}(\alpha) > 0$, $p_B^{''}(\alpha) > 0$, and $p_G^{''}(\alpha) > 0$.

PROOF: See Appendix 2. *Q.E.D.*

C Comparison between the Decentralized and Constrained Planner's Equilibria

C.1 Constrained Inefficiency of the Competitive Equilibrium

To distinguish between the two equilibria, hereafter we denote the competitive equilibrium and the constrained planner's equilibrium by α_t^C and α_t^P, respectively. The following proposition compares the levels of these two equilibria:

PROPOSITION 4: The competitive equilibrium is constrained inefficient, and is characterized by excessive risk-taking. That is, $\alpha_t^C(\pi_t) > \alpha_t^P(\pi_t)$ for any level of belief π_t.

PROOF: Recall that the competitive equilibrium and the constrained planner's equilibrium are given respectively by

$$\begin{aligned} F(\pi_t, \alpha_t^C) &= 0 \\ \Phi(\pi_t, \alpha_t^P) = F(\pi_t, \alpha_t^P) - \xi(\pi_t, \alpha_t^P) &= 0 \end{aligned}$$

where $\xi(\pi_t, \alpha_t^P) > 0$ by the proof of Proposition 3.

It follows that

$$F(\pi_t, \alpha_t^P) = \xi(\pi_t, \alpha_t^P) > 0 = F(\pi_t, \alpha_t^C)$$

Since $F(\pi_t, \alpha_t)$ is strictly decreasing in α_t for any π_t, we have $\alpha_t^P < \alpha_t^C$ for any π_t. *Q.E.D.*

C.2 Interaction between Learning and Inefficiency

Although the negative externality discussed above exists even in a model without learning, this subsection will examine the interaction between the negative externality and learning, and will establish that learning can aggravate the negative externality problem as long as the perceived vulnerability function $\theta(\pi_t, \alpha_t)$ satisfies some conditions.

To this end, we will first define the "sufficiently convex condition:"

Sufficiently convex condition (SCC): $\theta(\pi_t, \alpha_t)$ is sufficiently convex in α_t for any π_t, if and only if $\theta_{\alpha\alpha}(\pi_t, \alpha_t) > \frac{\theta_\alpha(\pi_t,\alpha_t)}{|\theta_\pi(\pi_t,\alpha_t)|}|\theta_{\alpha\pi}(\pi_t, \alpha_t)| - \theta_\alpha X$.

For the convenience of illustration, we make two more definitions:

Iso-vulnerability curve (IV): a locus of (π, α) along which the perceived vulnerability $\theta(\pi, \alpha)$ is constant, that is, $\pi p_G(\alpha) + (1-\pi)p_B(\alpha) = \bar{\theta}$ for some constant $\bar{\theta} \in [0, 1]$.

Iso-externality curve (IE): a locus of (π, α) along which the normalized externality $\xi(\pi, \alpha)$ is constant, that is, $\theta_\alpha(\pi_t, \alpha_t)\frac{1}{X}(\alpha_t) = \bar{\xi}$ for some constant $\bar{\xi} > 0$.

Based on these definitions, we have the following lemma:

LEMMA 1: The SCC is satisfied if and only if the slope of the iso-vulnerability curve is larger than that of the iso-externality curve.

PROOF: See Appendix 3. *Q.E.D.*

Figure 3 illustrates an example of IV and IE curves that satisfy the SCC. Based on this lemma, we can establish the following proposition, which compares the responses of α (with respect to π_t) by the competitive equilibrium and the constrained planner's equilibrium:

PROPOSITION 5: As investors become more optimistic after observing a longer history of financial tranquility, the investment in the competitive equilibrium α_t^C increases faster than that in the constrained efficient equilibrium α_t^P, if and only if the SCC is satisfied. In other words, when the SCC is satisfied, the learning process strengthens the negative externality, and aggravates the excessive risk-taking problem and the inefficiency.

PROOF: Recall that

$$\frac{d\alpha_t^C}{d\pi_t} = \frac{-\theta_\pi(\pi_t, \alpha_t)}{\theta_\alpha(\pi_t, \alpha_t)}$$

$$\frac{d\alpha_t^P}{d\pi_t} = \frac{-\theta_\pi(\pi_t, \alpha_t) - \xi_\pi(\pi_t, \alpha_t)}{\theta_\alpha(\pi_t, \alpha_t) + \xi_\alpha(\pi_t, \alpha_t)}$$

Therefore, $\frac{d\alpha_t^C}{d\pi_t} > \frac{d\alpha_t^P}{d\pi_t}$ if and only if

$$\frac{-\theta_\pi(\pi_t, \alpha_t)}{\theta_\alpha(\pi_t, \alpha_t)} > \frac{-\theta_\pi(\pi_t, \alpha_t) - \xi_\pi(\pi_t, \alpha_t)}{\theta_\alpha(\pi_t, \alpha_t) + \xi_\alpha(\pi_t, \alpha_t)}$$

Equivalently,

$$\frac{-\theta_\pi(\pi_t, \alpha_t)}{\theta_\alpha(\pi_t, \alpha_t)} > \frac{-\xi_\pi(\pi_t, \alpha_t)}{\xi_\alpha(\pi_t, \alpha_t)} \tag{7}$$

By Lemma 1, condition (7) is exactly the same as the necessary and sufficient condition for the function $\theta(\pi_t, \alpha_t)$ to be "sufficiently convex" in α_t. *Q.E.D.*

The intuition of the proposition can be understood as a "two-round" problem (note that this phrase is used only for the convenience of illustrations, and does not mean "Round One" occurs before "Round Two"). In "Round One," as the market and the constrained planner become more confident (π increases by $\Delta\pi$), two effects will occur: first, both the market and the constrained planner will assign a higher weight to p_G while computing the unconditional crisis probability (that is, perceived vulnerability) θ, and thus both will perceive a lower θ, which induces both the market and the planner to increase their aggregate risky investment α^C and α^P; second, the planner will assign a higher weight to $p_G^{'}(\alpha_t)$ while computing θ_α (that is, the planner believes the system becomes more resilient), and thus the externality $\xi \equiv \theta_\alpha \frac{1}{X}$ will decrease by $|\xi_\pi|\Delta\pi$, which induces the planner to further *increase* α^P relative to α^C by $|\alpha_\xi^P||\xi_\pi|\Delta\pi$ ("resilience effect, RE).[17] This effect appears counter-intuitive, but is – in our view – reasonable, and it is a new channel identified by our paper. Also note that the market will ignore the second effect (the externality) associated with the increase in π.

Then in "Round Two," as the aggregate risky investment α increases, the planner will want to decrease α to account for the externality. Specifically, the higher α will raise the unconditional crisis probability (that is, perceived vulnerability) θ, which in turn raises the externality ξ by $\xi_\alpha\Delta\alpha$ (recall that $\xi_\alpha = \theta_{\alpha\alpha}\frac{1}{X} + \theta_\alpha$). The higher externality will induce the planner to *decrease* α^P relative to α^C by $|\alpha_\xi^P|\xi_\alpha\Delta\alpha$ (risk-taking effect, RTE). This second effect is the traditional externality argument. Again, the market will ignore this second effect associated with the increase in α.

Therefore, the ultimate comparision between $\Delta\alpha^P$ and $\Delta\alpha^C$ depends on the comparison between the relative increase in α^P in "Round One" due to the RE and the relative decrease in α^P in "Round Two" due to the RTE. The SCC exactly captures this comparision: if θ is sufficiently convex in α, then the subsequent decrease in α^P due to the RTE will dominate

[17]Note that the parameter $|\xi_\pi|$ (that is,$|\frac{\partial\theta_\alpha}{\partial\pi}|$) is crucial for determining the magnitude of the RE.

the initial increase in α^P due to the RE, and thus α^P will increase less than α^C. Figure 4 illustrates this trade-off. We can also see this mathematically: On the one hand, the RTE dominates the RE if and only if $|\alpha^P_\xi|\xi_\alpha\Delta\alpha > |\alpha^P_\xi||\xi_\pi|\Delta\pi$, that is, $\frac{\Delta\alpha}{\Delta\pi} > \frac{|\xi_\pi|}{\xi_\alpha}$ (slope of the IE curve). On the other hand, all competitive equilibria correspond to the same θ (and thus any adjustment to α and π must be on the same IV curve), so $\frac{\Delta\alpha}{\Delta\pi}$ is actually the slope of the IV curve. Therefore, RTE dominates RE if and only if the SCC holds.

C.3 An Economic Example for SCC

We can use the following fire-sale example to illustrate the circumstances under which the SCC holds. A larger α strengthens fire-sale effects following a bad shock. Therefore, when α is small, a fire sale does not affect other asset values in the economy much. But when α is large, a fire sale has bigger effects on balance sheets in the economy, and may lead to a systemic crisis. Other kinds of contagion effects might work in a similar way.

One recent paper by Gertler and Kiyotaki (2015) provides support for this example. In that paper, financial crises become possible when, in the event of a run, the recovery value of bank assets falls below the bank's debt obligations to depositors. The recovery value can be low because bank runs necessitate fire sales in which banks sell their capital to households, who are less efficient managers of capital; that is, they have to pay a management fee that is convex in the amount of capital they manage (banks do not have to pay this fee). The authors assume that in states where a run is possible, the probability of a run is an exogenously specified function of the expected recovery rate. The expected recovery rate depends on exogenous states (for example, productivity) and on the parameters of the household's asset management cost function.

Now, suppose that households can invest either in banks (modeled as described above) or in some safe asset with a lower expected return. Let α be the share of their funds they choose to invest in the banks. In any aggregate state where a run is possible, the probability of a run is higher when α is larger. This is because a higher α means that banks are managing more capital, which in turn means that a fire sale would reduce capital prices by a larger amount given the convexity of the household's asset management cost function. This would be broadly consistent with our model. This example also implies that one way to empirically test our SCC condition would be to test whether the nonbank households' management fee is indeed convex (and how convex it is) in the amount of capital they manage.

III Competitive Equilibrium with Tax

In this section, we will show the existence of the optimal capital income tax that can restore the constrained efficiency of the competitive equilibrium. Moreover, we will derive an analytical solution for the optimal tax as a function of the belief, and discuss its properties from a macroprudential policy perspective.

A Decentralized Equilibrium with Tax

Specifically, the capital income tax scheme works as follows: for every dollar of the payoff from the risky asset at period t, the investor will pay D_t dollar.[18] Under such a tax scheme, the decentralized problem of investor i at period t is:

$$\begin{aligned}\max_{\alpha_{it}\epsilon[0,1]} \quad & \pi_t\{(1-p_G)[(1-\alpha_{it})\mu_S+\alpha_{it}\mu_R(1-D_t)] & + \\ & p_G[(1-\alpha_{it})\tau_S\mu_S+\alpha_{it}\tau_R\mu_R(1-D_t)]\} & + \\ & (1-\pi_t)\{(1-p_B)[(1-\alpha_{it})\mu_S+\alpha_{it}\mu_R(1-D_t)] & + \\ & p_B[(1-\alpha_{it})\tau_S\mu_S+\alpha_{it}\tau_R\mu_R(1-D_t)]\}\end{aligned}$$

Using the previous notations for the two constants $\Delta_0 \equiv \mu_R - \mu_S$ and $\Delta_1 \equiv \tau_R\mu_R - \tau_S\mu_S$, we can rewrite the above problem of investor i at period t as:

$$\max_{\alpha_{it}\epsilon[0,1]} \quad (1-\theta_t)[\mu_S+\alpha_{it}(\Delta_0-\mu_R D_t)]+\theta_t[\tau_S\mu_S+\alpha_{it}(\Delta_1-\tau_R\mu_R D_t)] \tag{8}$$

The FOC of the above problem is (taking p_G and p_B as given):

$$\Delta_0-\mu_R D_t+[(\Delta_1-\tau_R\mu_R D_t)-(\Delta_0-\mu_R D_t)]\theta_t=0 \tag{9}$$

B Optimal Tax and Macroprudential Policy Implications

Now we derive the analytical solution of the optimal tax rate. For this purpose, denote the social planner's risky position as α^P. Corresponding to this α^P, there is an associated $\theta^P \equiv \pi_t p_G(\alpha^P)+(1-\pi_t)p_B(\alpha^P) \in (0,1)$. This θ^P is the level of θ prevailing in the constrained efficient equilibrium. It is also the target level of θ that the competitive equilibrium needs to achieve (by adjusting the tax D_t) in order to be constrained efficient.

Plug the target level of θ (that is, θ^P) into the competitive equilibrium's FOC, and we get the unique optimal tax rate:

$$D^* = \frac{\Delta_0+(\Delta_1-\Delta_0)\theta^P}{\mu_R+(\tau_R\mu_R-\mu_R)\theta^P} \tag{10}$$

Using $\Delta_0 = \mu_R - \mu_S$ and $\Delta_1 = \tau_R\mu_R - \tau_S\mu_S$ and after some rearrangements, we have

$$D^* = 1-\frac{\mu_S[1-(1-\tau_S)\theta^P]}{\mu_R[1-(1-\tau_R)\theta^P]} \tag{11}$$

The proposition below formally presents the existence of a well-defined optimal tax rate:

[18] To ensure there is no trivial welfare implication, it is assumed that tax revenues are rebated lump-sum to investors. Since this does not affect the optimization problem, we abstract it from the objective function.

PROPOSITION 6: For any level of belief π_t, there exists a unique optimal tax rate $D^* \in (0, 1)$ given by equation (11), which can restore the constrained efficiency of the competitive equilibrium.

PROOF: See Appendix 4. *Q.E.D.*

Moreover, the optimal tax D^* has an important relationship with the belief π_t, which is summarized in the following proposition:

PROPOSITION 7: As investors and the constrained planner become more confident (that is, as π_t increases), the optimal tax rate D_t^* needed to restore the constrained efficiency will be higher, if and only if the SCC is satisfied.

PROOF: See Appendix 5. *Q.E.D.*

Combining Proposition 1 and Proposition 7, we have the following important macroprudential policy implication: if and only if the SCC holds, the optimal macroprudential policy should be countercyclical; that is, as the market tranquility lasts for one more period, the policymaker should raise the capital income tax in order to curb the excessive risk-taking and lower the systemic risk.

The qualification of this proposition (the perceived vulnerability function $\theta(\alpha_t, \pi_t)$ needs to be sufficiently convex in α_t) has an important policy implication. For this purpose, let us interpret the constrained planner as the financial regulator. When the regulator becomes more confident that the financial sector is strong, s/he will assign a higher weight to $p_G(\alpha_t)$ (while computing $\theta(\alpha_t, \pi_t)$) as well as to $p_G^{'}(\alpha_t)$ (while computing $\theta_\alpha(\alpha_t, \pi_t)$). That is, as the regulator becomes more optimistic (π_t becoming higher), s/he will believe that not only the crisis would occur with a lower probability (corresponding to a lower *level* of the perceived vulnerability $\theta(\alpha_t, \pi_t)$), but also the financial system will be more resilient to any additional buildup in the aggregate position of the risky investment (corresponding to a lower marginal *derivative* of the perceived vulnerability $\theta(\alpha_t, \pi_t)$ with respect to α_t). As a result, the higher confidence tends to induce the regulator to initially *increase* the aggregate risky position α^P more than the increase of α^C by the market (resilience effect).

However, there is a countervailing effect: as α^P increases, the regulator understands that the externality will increase, which would induce the regulator to subsequently decrease α^P more than the decrease of α^C by the market (risk-taking effect). If the risk-taking effect dominates the resilience effect, the ultimate increase in α^P will be lower than that in α^C, and the regulator should *increase* the optimal tax to achieve this (countercyclical policy); otherwise, the regulator should decrease the optimal tax (procyclical policy).

The risk-taking effect is usually associated with cyclical indicators "on the surface." For example, a sustained credit boom usually suggests more risk-taking by investors. However, the resilience effect can only be assessed by closely examining the deep structural change of the financial system. Proposition 7 implies that macroprudential policymakers should go beyond the cyclical indicators to look into the underlying structural change of the system. For example, suppose the high credit growth is mainly driven by lendings to non-oil industries in an economy that used to heavily depend on oil exports. Then by diversifying the economy's structure, the credit growth may have fundamentally improved the resilience of the system to shocks. In this case, even though cyclical indicators suggest a sustained credit boom, tightened regulation may not be desirable. Such an approach is important also because financial cycles are much longer than business cycles (currently about 15 years in the U.S. and the U.K., according to Strohsal and others, 2015; and about 15 years for an average G-7 country, according to Schüler and others, 2017). Hence, during the long period of credit *decline*, it may *not* be desirable to keep *loosening* macroprudential policies; similarly, during the long period of credit boom, it may not be desirable to keep tighening macroprudential policies without gauging the change of the resilience of the system.

C Inefficient Deregulation/Regulation and Further Comparative Statics

The qualification of Proposition 7 raises an interesting question: what if the financial regulator mistakenly believes that the SCC fails? This subsection examines this question using a simple extension of the previous model.

Specifically, introduce to the previous model a "regulator" who has the following features: first, like the constrained planner, the regulator has to learn about the true state of the financial sector and cares about the welfare of all investors; second, unlike the constrained planner, the regulator mechanically[19] overestimates $|\xi_\pi|$ such that $\frac{|\theta_\pi(\pi_t,\alpha_t)|}{\theta_\alpha(\pi_t,\alpha_t)} \leq \frac{|\xi_\pi(\pi_t,\alpha_t)|}{\xi_\alpha(\pi_t,\alpha_t)}$ for some π_t and some α_t (that is, s/he believes that the SCC fails), even though the SCC actually holds. Moreover, the regulator makes the decision based on the overestimated $|\xi_\pi|$. This second assumption captures the possibility that some regulators in reality tend to overestimate the resilience of the financial system.

Consider an initial situation where the belief is π and the optimal capital income tax needed to restore constrained efficiency is D^*. As π increases to $\pi + \Delta\pi$, Proposition 7 implies that the capital income tax chosen by this regulator will decrease to $D^{Reg}(\pi + \Delta\pi) = D^* - \Delta D$. However, because the SCC actually holds, Proposition 7 also implies that the actual optimal capital income tax that would be chosen by the constrained planner, denoted by $D^P(\pi+\Delta\pi)$, is larger than D^*. Hence,

$$D^{Reg}(\pi + \Delta\pi) < D^* < D^P(\pi + \Delta\pi)$$

And thus

$$\alpha^{Reg}(\pi + \Delta\pi) > \alpha^P(\pi + \Delta\pi)$$

More formally, we have the following corollary to Proposition 7 and Proposition 1:

COROLLARY 1: In case the regulator overestimates the resilience of the financial system and mistakenly believes the SCC fails ($\frac{|\theta_\pi|}{\theta_\alpha} \leq \frac{|\xi_\pi|}{\xi_\alpha}$): as the regulator observes one more no-crisis period and becomes more confident, s/he will lower the capital income tax and induce an inefficiently high aggregate risky investment position.

[19]"Mechanically" means that the regulator does not learn about $|\xi_\pi|$, but instead makes a mechanical and persistent judgment about it.

The flip side is inefficient regulation. In case the regulator underestimates the resilience of an innovative financial sector and mistakenly believes the SCC holds ($\frac{|\theta_\pi|}{\theta_\alpha} > \frac{|\xi_\pi|}{\xi_\alpha}$): as the regulator observes one more no-crisis period and understands that the market becomes more confident, s/he will tighten the regulation and induce an inefficiently low aggregate investment in this industry.

COROLLARY 2: If the crisis damages the safe real sector more (lower τ_S) and/or the risky financial sector less (higher τ_R), then:

(i) the market will invest more in the risky sector.

(ii) the planner will be more likely to raise the tax when the market has been tranquil for a longer time.

The intuition of the corollary is as follows: a lower τ_S and/or higher τ_R imply that the safe sector is less attractive than the risky one. This induces investors to invest more in the risky sector, hence Part (i) of the corollary. Moreover, given the stronger incentive of the market to invest in the risky sector, α^C is more responsive to the confidence π under a lower τ_S and/or higher τ_R; to contain the higher externality, the planner would be more likely to raise the tax as the financial tranquility persists for one more period (and the confidence builds up further).

IV Parametric Examples

Consider the following parametric specification:

$$\begin{aligned} p_G(\alpha) =& a_G + b_G \cdot \alpha^k \text{ and} \\ p_B(\alpha) =& a_B + b_B \cdot \alpha^k. \end{aligned}$$

We will assume the followings: (1) $a_B \geq a_G \geq 0$, $b_B \geq b_G > 0$; (2) $a_i + b_i \leq 1$ for any $i = B, G$ and (3) $k \geq 1$. The first two assumptions guarantee that the probability of crisis in either state $p_G, p_B \in [0, 1]$ and are increasing in the aggregate risky investment α. The third assumption implies that the probability of crisis in either state is convex in α.

Thus,

$$\theta = (\pi a_G + (1-\pi)a_B) + (\pi b_G + (1-\pi)b_B)\,\alpha^k.$$

This implies

$$\begin{aligned} -\theta_\pi = (a_B - a_G) + \alpha^k(b_B - b_G),\, \theta_\alpha = k\alpha^{k-1}(\pi b_G + (1-\pi)b_B), \\ -\theta_{\pi\alpha} = k\alpha^{k-1}(b_B - b_G),\, \theta_{\alpha\alpha} = k(k-1)\alpha^{k-2}(\pi b_G + (1-\pi)b_B). \end{aligned}$$

Recall that the slope of IV is

$$\frac{-\theta_\pi}{\theta_\alpha} = \frac{\alpha(b_B - b_G) + (a_B - a_G)\alpha^{1-k}}{k(\pi b_G + (1-\pi)b_B)}.$$

On the other hand,

$$\xi(\pi_t, \alpha_t) = \theta_\alpha(\pi_t, \alpha_t)\frac{1}{X} = \theta_\alpha(\pi_t, \alpha_t)[\alpha_t + \frac{(1-\tau_S)\mu_S}{\Delta_0 - \Delta_1}]$$

and slope of IE is

$$\frac{-\xi_\pi}{\xi_\alpha} = \frac{-\theta_{\pi\alpha}}{\theta_{\alpha\alpha} + \theta_\alpha X} = \frac{\alpha(b_B - b_G)}{((k-1) + \alpha X)(\pi b_G + (1-\pi)b_B)}.$$

Therefore, condition SCC holds true, that is, the slope of IV is greater than the slope of IE iff

$$\begin{aligned}
\frac{\alpha(b_B - b_G) + (a_B - a_G)\alpha^{1-k}}{k(\pi b_G + (1-\pi)b_B)} &> \frac{\alpha(b_B - b_G)}{((k-1) + \alpha X)(\pi b_G + (1-\pi)b_B)}, \\
\alpha(b_B - b_G)(k - (1 - \alpha X)) + (a_B - a_G)\alpha^{1-k}(k - (1 - \alpha X)) &> \alpha(b_B - b_G)k \\
(a_B - a_G)\alpha^{-k}(k - (1 - \alpha X)) &> (b_B - b_G)(1 - \alpha X) \\
k &> \left(1 + \frac{b_B - b_G}{a_B - a_G}\alpha^k\right)(1 - \alpha X).
\end{aligned}$$

Assuming strict inequalities $a_B > a_G$ and $b_B > b_G$, we can say that under this parametric specification, if $k > 1 + \frac{b_B - b_G}{a_B - a_G}$, then condition SCC holds, but otherwise condition SCC may not hold.

Suppose that a marginal increase in aggregate risky investment α increases the probability of crisis exactly the same regardless of the states being G or B, that is, $p'_G = p'_B$. In this parametric example, this is equivalent to $b_B = b_G$. Note that the RHS of the above inequality becomes $1 - \alpha X < 1$. Thus, condition SCC holds for any $k > 1$ (however small). Therefore, the planner should optimally increase the tax when the market has been tranquil for a longer time.

Let us consider another extreme case. Suppose that at zero aggregate risky investment, the probability of crisis is the same regardless of the state - that is, $a_B = a_G$. Then the RHS of the above inequality becomes ∞. Thus, condition SCC does not hold regardless of k (however large). Therefore, the planner should optimally decrease the tax when the market has been tranquil for a longer time.

V Dynamics of Learning and Markov Switching

This section further examines the dynamics of learning under the assumption of a fixed true state, as well as extends our model to the case where the true state follows a Markov process.

A Dynamics of Learning

Proposition 1 has the following corollary:

COROLLARY 3: The dynamics of investors' learning process have two properties:

First, investors will not revise down their confidence until the crisis has actually hit.

Second, the occurrence of crisis only affects investor confidence in the immediately next period alone, and the confidence will again improve from the second post-crisis period on as long as there is no crisis in the previous period.

PROOF: The first property directly follows the proof of Proposition 1. Below is the proof of the second property.

Suppose the first crisis occurs in Period t_1^C. Proposition 1 implies that at the beginning of the immediately next period $t = t_1^C + 1$, investors' posterior belief $\pi_{t_1^C+1} < \pi_{t_1^C}$.

At the beginning of any $t_1^C + 2 \leq t \leq t_2^C$, where t_2^C is the period during which the second crisis occurs, investors' posterior belief is:

$$\begin{aligned}
\pi_t &= P(j = G | C_{t-1} = 0) \\
&= \frac{\pi_{t-1}[1 - p_G(\alpha_{t-1})]}{\pi_{t-1}[1 - p_G(\alpha_{t-1})] + (1 - \pi_{t-1})[1 - p_B(\alpha_{t-1})]} \\
&> \frac{\pi_{t-1}[1 - p_G(\alpha_{t-1})]}{\pi_{t-1}[1 - p_G(\alpha_{t-1})] + (1 - \pi_{t-1})[1 - p_G(\alpha_{t-1})]} \\
&= \pi_{t-1}.
\end{aligned}$$

That is, the first crisis only lowers the confidence at the immediately next period; after then, the confidence will still increase in the number of tranquil periods. The above proof applies to all crises. *Q.E.D.*

In some sense, the two properties in this corollary suggest that investors are subject to a complacency trap: they will lower the confidence only after they actually go through a crisis, and they will (rationally) become "complacent" again two periods after the crisis. The top-left panel of Figure 5 (Figure 6) plots the dynamics of the investors' posterior belief when the financial sector's true state is fragile (strong), under some parametrizations of the $p_G(.)$ and $p_B(.)$ functions. In these figures, one sharp decline corresponds to the occurrence of one crisis. Note that even though the posterior belief converges to 0 when the true state is fragile, the peak of the learning cycle does not decrease over time as the system is hit by more and more crises: investors can still be more confident at the peak of a subsequent cycle than they were at the peak of a previous cycle. This further confirms the complacency trap.

B Markov Switching

If the underlying state is fixed, then the agents will eventually learn the underlying state, as the simulation shows (Figures 5 and 6). Note that under the social planner solution, the crisis occurs less frequently regardless of the state, due to the fact that the planner will contain the excessive risk-taking *ex ante*.

If the underlying state is not fixed, that is, a good state may become bad or a bad state may become good, then agents may never learn the true state. Our model can be extended to such a setting. For example, consider a Markov transition, where the underlying state remains the same with probability $q \in (\frac{1}{2}, 1)$ and switches with probability $1 - q$.[20] While updating their beliefs that the state is G, agents will consider both possibilities that (1) the state j was G and it remains G, and (2) the state j was B but it has become G. Let π_t be the interim updated belief that $j_{t-1} = G$ and $\pi_t^{'}$ be the final updated belief that $j_t = G$.

$$
\begin{aligned}
P(j_t = G | C_{t-1} = 0) = & \; P(j_t = G,\, j_{t-1} = G | C_{t-1} = 0) + P(j_t = G,\, j_{t-1} = B | C_{t-1} = 0) \\
= & \; P(j_t = G | j_{t-1} = G) P(j_{t-1} = G | C_{t-1} = 0) \\
& + P(j_t = G | j_{t-1} = B) P(j_{t-1} = B | C_{t-1} = 0).
\end{aligned}
$$

$$
\begin{aligned}
\Rightarrow \pi_t'(C_{t-1} = 0) &= q\pi_t(C_{t-1} = 0) + (1 - q)[1 - \pi_t(C_{t-1} = 0)] \\
&= (2q - 1)\pi_t(C_{t-1} = 0) + (1 - q),
\end{aligned}
$$

where $\pi_t(C_{t-1} = 0)$ is exactly the same as defined before (that is, in the case where the true state is fixed). Similarly, we can also define and obtain

$$
\begin{aligned}
P(j_t = G | C_{t-1} = 1) &= \pi_t'(C_{t-1} = 1) \\
&= (2q - 1)\pi_t(C_{t-1} = 1) + (1 - q).
\end{aligned}
$$

The simulation in Figure 7 shows the dynamics of belief and risky investment of the decentralized market and those of the social planner.

Importantly, given that $q > \frac{1}{2}$ and thus $2q - 1 > 0$, it follows that $\pi_t'(C_{t-1} = 0)$ (under Markov switching) is monotonically increasing in $\pi_t(C_{t-1} = 0)$ (under fixed true states). Therefore, Proposition 1 still holds under Markov switching, that is, a longer history of financial tranquility builds investor confidence that the sector is strong. Consequently, all other propositions hold under Markov switching. The same arguments apply to $\pi_t'(C_{t-1} = 1)$.

VI Unconstrained Planner's Problem and Informed Policymaker

In this section, we consider the problem of an unconstrained social planner who has full information about the true state of the financial sector. The motivation for this assumption is that in practice, macroprudential policymakers are likely to know more about the true state (G or B) of the system than investors or the uninformed/constrained "planner." This is because they spend an enormous amount of time and effort trying to assess the resilience of the system, including through stress tests and indicators of balance sheet vulnerabilities. Examples of such effort include the Bank of England's annual cyclical stress tests and biennial exploratory stress tests, the Federal Reserve's Dodd-Frank Act Stress Tests and Comprehensive Capital Analysis and Review, etc. Such informed policymakers could more

[20] The assumption that $q > \frac{1}{2}$ makes more sense than $q < \frac{1}{2}$ because the underlying state is likely to display some persistency.

realistically be modeled by assuming that the policymaker receives an additional independent signal about the state, but as a simple approximation, they are modeled by assuming that they have full information about the true state.

Consider the case where the true state is B (fragile). Then $\theta(\pi_t, \alpha_t) = p_B(\alpha_t)$, $\xi(\alpha_t, \pi_t) = p_B^{'}(\alpha)[\alpha_t + \frac{(1-\tau_S)\mu_S}{\Delta_0 - \Delta_1}]$, and the problem of the unconstrained planner at period t is:

$$\max_{\alpha_t \epsilon [0,1]} \quad [1 - p_B(\alpha_t)](\mu_S + \alpha_t \Delta_0) + p_B(\alpha_t)(\tau_S \mu_S + \alpha_t \Delta_1) \tag{12}$$

The FOC of the unconstrained planner is

$$\frac{\Delta_0}{\Delta_0 - \Delta_1} - p_B(\alpha_t) - p_B^{'}(\alpha)[\alpha_t + \frac{(1-\tau_S)\mu_S}{\Delta_0 - \Delta_1}] = 0 \tag{13}$$

Comparision between the normalized FOC of the constrained planner's equilibrium with that of the unconstrained planner (equation (6) versus equation (13)) makes it clear that both planners adjust α_t to keep $\theta(\pi_t, \alpha_t) + \xi(\pi_t, \alpha_t)$ constant at $\frac{\Delta_0}{\Delta_0 - \Delta_1}$, with the only difference being about the definition of $\theta(\pi_t, \alpha_t)$.

The following proposition characterizes the unconstrained planner's equilibrium:

PROPOSITION 8: In case the true state is B (fragile), the aggregate risky investment in the constrained planner's equilibrium (α_t^P) is strictly higher than that in the unconstrained/informed planner's equilibrium (α_t^{IP}).

PROOF: The normalized FOCs of the two planners imply that

$$p_B(\alpha_t^{IP}) + p_B^{'}(\alpha_t^{IP})[\alpha_t^{IP} + \frac{(1-\tau_S)\mu_S}{\Delta_0 - \Delta_1}] = \frac{\Delta_0}{\Delta_0 - \Delta_1} = \theta(\pi_t, \alpha_t^P) + \theta_\alpha(\pi_t, \alpha_t^P)[\alpha_t^P + \frac{(1-\tau_S)\mu_S}{\Delta_0 - \Delta_1}]$$

Since $p_B(\alpha_t) > \theta(\pi_t, \alpha_t)$ and $p_B^{'}(\alpha_t) > \theta_\alpha(\pi_t, \alpha_t^P)$ for any given α_t, the LHS would be larger than the RHS if $\alpha_t^{IP} = \alpha_t^P$.

Since both sides strictly increase in α_t, it follows that $\alpha_t^{IP} < \alpha_t^P$. *Q.E.D.*

This proposition has an important policy implication. It implies that due to information asymmetry, the choice of risky investment made by the uninformed/constrained "planner" is inefficiently high, even though this choice is more efficient than decentralized investors (since the constrained planner endogenizes the externality; see Proposition 4). Therefore, it is desirable to assign the macroprudential policy function to independent agencies with technical expertise that allows them to gauge the underlying true resilience of the financial system (central banks or dedicated macroprudential agencies, such as the Financial Stability Oversight Council and Office of Financial research in the US). This is consistent with the viewpoints of IMF (2011, 2014).[21]

[21] Since the aggregate risky investment in the unconstrained planner's equilibrium α_t^{IP} does not depend on the belief π_t, both α_t^{IP} and the optimal tax chosen by the unconstrained planner remain constant as the market and the constrained planner become more confident.

VII Conclusion

Recent empirical evidence suggests that financial crises tend to follow prolonged periods of financial stability and investor optimism. In this paper, we examine how to account for the cycles of optimism and pessimism in the design of macroprudential policies. Our first contribution is to illustrate the excessive risk-taking in financial markets through the *interaction* between the negative externality and learning. A longer history of financial tranquility (that is, the absence of crisis), as observed in the lead-up to the global crisis of 2007-2008, builds the confidence of investors and leads them to take larger positions in the risky asset. However, a larger aggregate risky asset position raises the probability of the crisis. Moreover, each trader imposes a negative externality through the effect of his or her risky asset position on the systemic crisis probability. As a result, the decentralized equilibrium is characterized by constrained inefficiency and excessive risk-taking.

Our second contribution is to provide a simple framework to assess the efficiency of macroprudential regulation. We characterize conditions for the degree of constrained inefficiency to be increasing in investor confidence, which also turn out to be the conditions for the countercyclicality of optimal macroprudential policies. We evaluate the potential of a macroprudential policy in the form of a capital income tax (similar to the tax proposed by Acharya and others, 2009 and Acharya and others, 2010) to restore constrained efficiency and reduce systemic risk, and find that under the same conditions the optimal tax rate should be higher as the market tranquility persists. This result contributes to the literature by highlighting that optimal macroprudential policies are *not* always countercyclical; instead, they depend on the trade-off between the resilience effect and risk-taking effect. If the policymaker incorrectly assesses this trade-off, then s/he may engage in an "inefficient regulation." The inefficiency could also go in the other direction: a regulator who underestimates the resilience of the financial system may repress financial activities needlessly. This could prevent the uptake of financial innovations that might really be effective at delivering value to investors. Our framework could be used to shed light on these discussions.

There are two avenues for future research. First, it can be directed at providing a micro foundation for the crisis probability functions ($p_G(\alpha)$ and $p_B(\alpha)$). Second, it can be directed at testing the implications of our model using rigorous empirical methods, possibly along the lines of the approach indicated in Subsection II.C.3.

Appendices

Figure 1: Input Effect and Transmission Effect

Source: Authors.

Figure 2: Payoff Structure

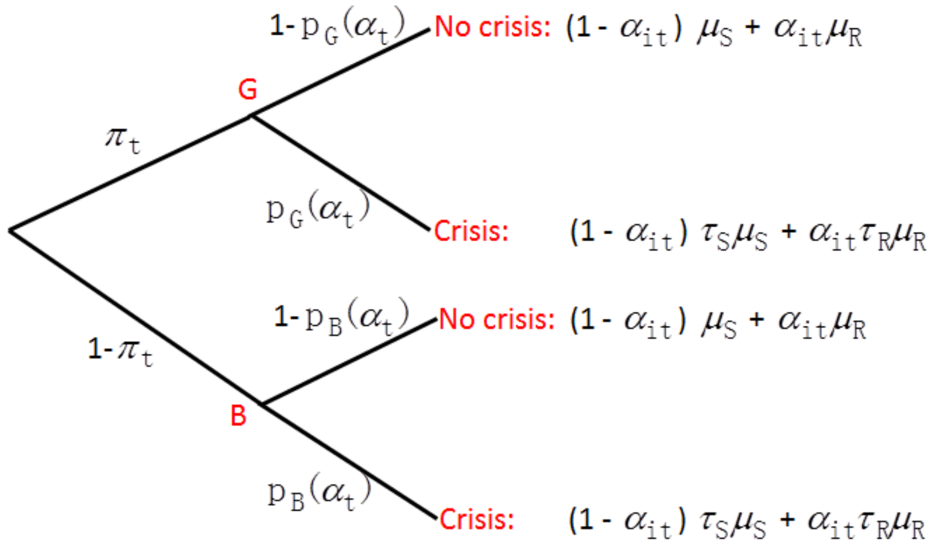

Source: Authors.

Figure 3: Iso-Vulnerability and Iso-Externality Curves

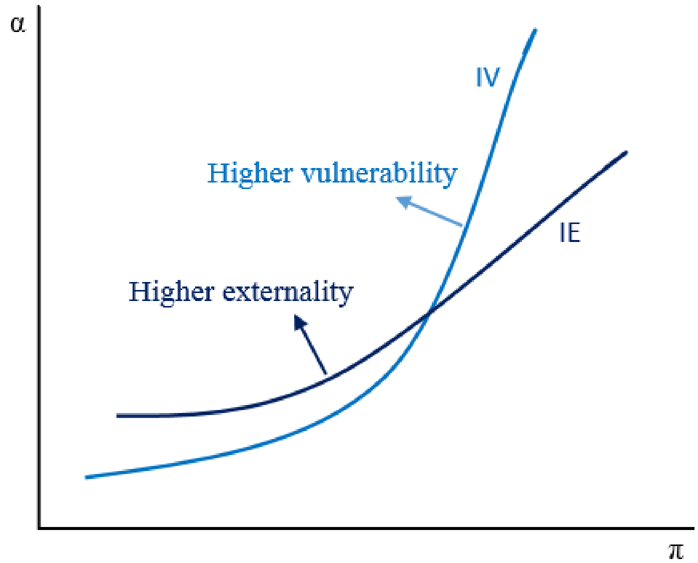

Source: Authors.

Figure 4: Trade-off between Risk-Taking Effect and Resilience Effect

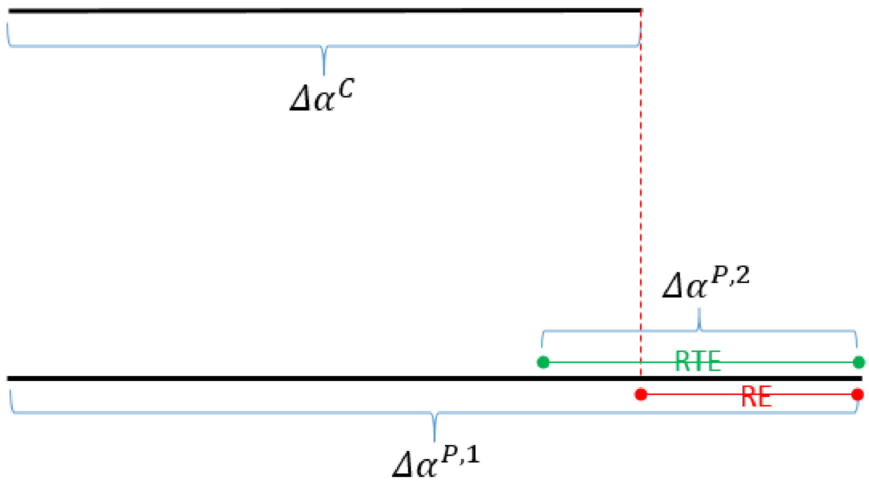

Source: Authors.

Note: $\Delta\alpha^C$ is the ultimate change of the market's risky position; $\Delta\alpha^{P,1}$ and $\Delta\alpha^{P,2}$ are the change of the planner's risky position in the first and the second "round," respectively; RE represents resilience effect, and RTE represents risk-taking effect.

Figure 5: Dynamics of Learning and Investments: Bad State

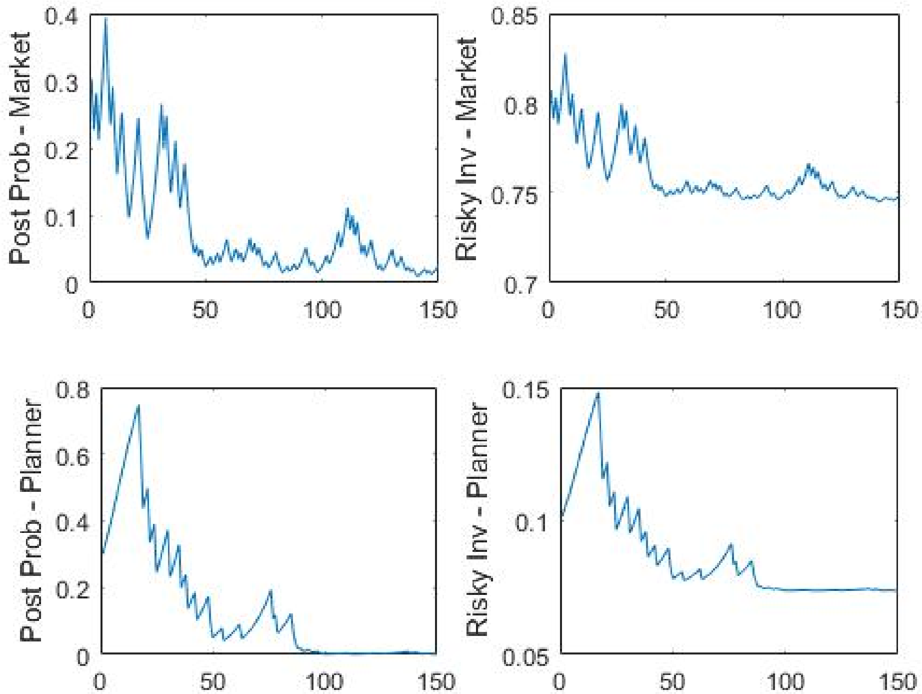

Source: Authors.

Figure 6: Dynamics of Learning and Investments: Good State

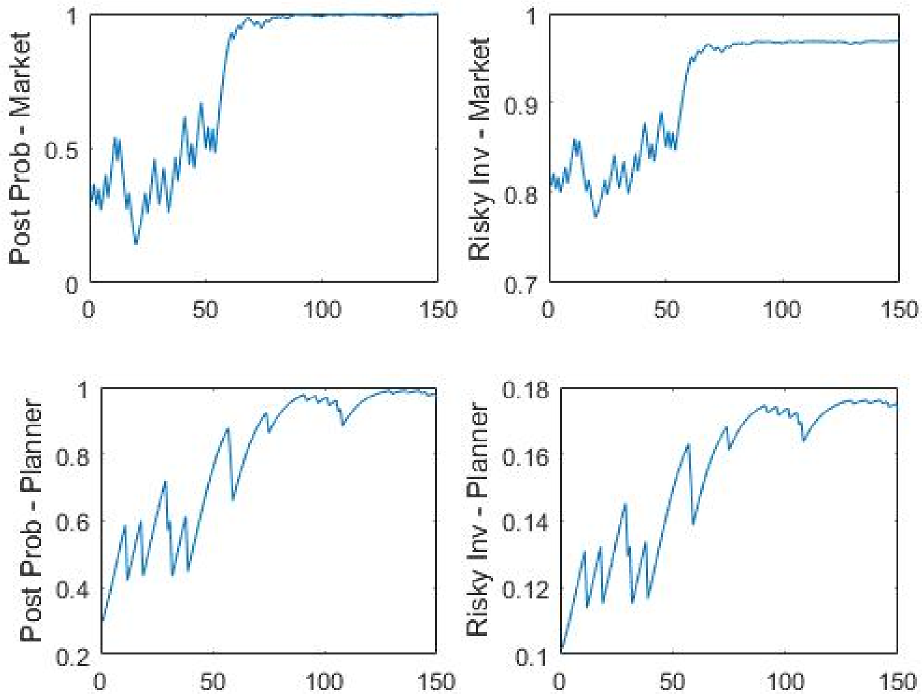

Source: Authors.

Figure 7: Dynamics of Learning and Investments under Market Switching

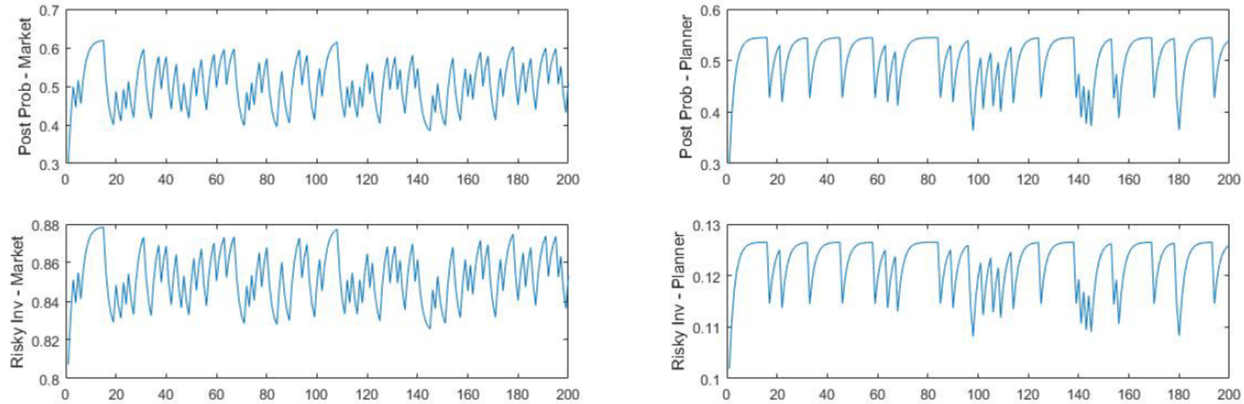

Source: Authors.

Appendix 1: Proof of Proposition 2

PROOF: Rewrite the FOC (3) as

$$F(\alpha_t, \pi_t) = 0$$

It follows that[22]

$$F_\alpha(\pi_t, \alpha_t) = -\theta_\alpha(\pi_t, \alpha_t) = -[\pi_t p_G^{'}(\alpha_t) + (1-\pi_t)p_B^{'}(\alpha_t)]$$

$$F_\pi(\pi_t, \alpha_t) = -\theta_\pi(\pi_t, \alpha_t) = p_B(\alpha_t) - p_G(\alpha_t) > 0$$

Then by implicit function theorem, we have:

$$\begin{aligned} \frac{d\alpha_t}{d\pi_t} &= -\frac{F_\pi(\pi_t, \alpha_t)}{F_\alpha(\pi_t, \alpha_t)} \\ &= \frac{-\theta_\pi(\pi_t, \alpha_t)}{\theta_\alpha(\pi_t, \alpha_t)} \end{aligned} \tag{14}$$

Since $p_B > p_G$, we have $\frac{d\alpha_t}{d\pi_t} > 0$ provided $p_G^{'}(\alpha) > 0$ and $p_B^{'}(\alpha) > 0$. *Q.E.D.*

Appendix 2: Proof of Proposition 3

PROOF: Rewrite the planner's FOC (6) as

$$\Phi(\pi_t, \alpha_t) = 0$$

It follows that

$$\Phi_\alpha(\pi_t, \alpha_t) = -\theta_\alpha(\pi_t, \alpha_t) - \xi_\alpha(\pi_t, \alpha_t)$$

$$\Phi_\pi(\pi_t, \alpha_t) = -\theta_\pi(\pi_t, \alpha_t) - \xi_\pi(\pi_t, \alpha_t)$$

where $\xi_\alpha(\pi_t, \alpha_t) = \theta_{\alpha\alpha}(\pi_t, \alpha_t)\frac{1}{X} + \theta_\alpha(\pi_t, \alpha_t) > 0$, and $\xi_\pi(\pi_t, \alpha_t) = \theta_{\alpha\pi}(\pi_t, \alpha_t)\frac{1}{X}$.

By implicit function theorem, we have

$$\frac{d\alpha_t}{d\pi_t} = -\frac{\Phi_\pi(\pi_t, \alpha_t)}{\Phi_\alpha(\pi_t, \alpha_t)} = \frac{-\theta_\pi(\pi_t, \alpha_t) - \xi_\pi(\pi_t, \alpha_t)}{\theta_\alpha(\pi_t, \alpha_t) + \xi_\alpha(\pi_t, \alpha_t)} \tag{15}$$

for the constrained planner's equilibrium.

The proof of Proposition 2 indicates that $\theta_\pi(\pi_t, \alpha_t) < 0$ and $\theta_\alpha(\pi_t, \alpha_t) > 0$, provided $p_G^{'}(\alpha) > 0$ and $p_B^{'}(\alpha) > 0$. Note that

$$\xi_\pi(\pi_t, \alpha_t) = \theta_{\alpha\pi}(\pi_t, \alpha_t)\frac{1}{X} = [p_G^{'}(\alpha_t) - p_B^{'}(\alpha_t)]\frac{1}{X} < 0$$

[22] Following the standard notations in calculas, $F_x(x, y)$ denotes the partial derivative of $F(x, y)$ with respect to x, treating y constant. Similar comments apply to $F_y(x, y)$, $F_{xy}(x, y)$, $F_{xx}(x, y)$, etc.

provided $p_B^{'}(\alpha_t) > p_G^{'}(\alpha_t)$; and

$$\theta_{\alpha\alpha}(\pi_t, \alpha_t) = \pi_t p_G^{''}(\alpha_t) + (1 - \pi_t) p_B^{''}(\alpha_t) > 0$$

provided $p_G^{''}(\alpha_t) > 0$ and $p_B^{''}(\alpha_t) > 0$. Therefore, $\frac{d\alpha_t}{d\pi_t} > 0$ provided $p_B^{'}(\alpha) > p_G^{'}(\alpha) > 0$, $p_B^{''}(\alpha) > 0$, and $p_G^{''}(\alpha) > 0$. *Q.E.D.*

Appendix 3: Proof of Lemma 1

PROOF: Take the total derivative of the IV equation with respect to π:

$$\theta_\pi + \theta_\alpha \frac{d\alpha}{d\pi} = 0$$

So the slope of IV equals

$$\frac{d\alpha}{d\pi}|_{IV} = \frac{-\theta_\pi}{\theta_\alpha}$$

Notice that $-\theta_\pi = p_B(\alpha) - p_G(\alpha) > 0$ and $\theta_\alpha = \pi p_G^{'}(\alpha) + (1 - \pi) p_B^{'}(\alpha) > 0$, so $\frac{d\alpha}{d\pi}|_{IV} > 0$, that is, the IV curve slopes upward.

Similarly, the slope of IE equals

$$\frac{d\alpha}{d\pi}|_{IE} = \frac{-\xi_\pi}{\xi_\alpha} = \frac{-\theta_{\alpha\pi}\frac{1}{X}}{\theta_{\alpha\alpha}\frac{1}{X} + \theta_\alpha} = \frac{-\theta_{\alpha\pi}}{\theta_{\alpha\alpha} + \theta_\alpha X}$$

Also, $-\theta_{\alpha\pi} = p_B^{'}(\alpha) - p_G^{'}(\alpha) > 0$, $\theta_{\alpha\alpha} = \pi p_G^{''}(\alpha) + (1 - \pi) p_B^{''}(\alpha) > 0$, $\theta_\alpha > 0$, and $X > 0$, so $\frac{d\alpha}{d\pi}|_{IE} > 0$, that is, the IE curve slopes upward as well.

Then we have

$$\frac{d\alpha}{d\pi}|_{IV} = \frac{-\theta_\pi}{\theta_\alpha} > \frac{-\theta_{\alpha\pi}}{\theta_{\alpha\alpha} + \theta_\alpha X} = \frac{d\alpha}{d\pi}|_{IE}$$

if and only if

$$\theta_{\alpha\alpha} > \frac{\theta_\alpha}{|\theta_\pi|}|\theta_{\alpha\pi}| - \theta_\alpha X,$$

which is the SCC. *Q.E.D.*

Appendix 4: Proof of Proposition 6

PROOF: Since $0 < (1 - \tau_S)\theta^P < 1$, $0 < (1 - \tau_R)\theta^P < 1$, $\mu_S > 0$, and $\mu_R > 0$, we have $D^* < 1$.

Now suppose $D^* \leq 0$, that is, $\mu_S[1 - (1 - \tau_S)\theta^P] \geq \mu_R[1 - (1 - \tau_R)\theta^P]$. Using $\mu_R - \mu_S = \Delta_0$ and $\tau_R \mu_R - \tau_S \mu_S = \Delta_1$, this implies:

$$-\mu_S(1 - \tau_S)\theta^P \geq -\mu_R(1 - \tau_R)\theta^{SP} + \mu_R - \mu_S$$

$$\Delta_0 + (\Delta_1 - \Delta_0)\theta^P \leq 0 \tag{16}$$

However, the planner's FOC (5) implies that

$$\Delta_0 + (\Delta_1 - \Delta_0)\theta^P = \theta_\alpha(\alpha_t, \pi_t)[\alpha_t(\Delta_0 - \Delta_1) + (1 - \tau_s)y] > 0$$

where the last inequality used $\theta_\alpha(\alpha_t, \pi_t) = \pi_t p_G^{'}(\alpha) + (1 - \pi_t)p_B^{'}(\alpha) > 0$. This contradicts inequality (16). Therefore, we have $D^* > 0$ for any π_t. *Q.E.D.*

Appendix 5: Proof of Proposition 7

PROOF: The analytical solution of D^*, equation (11), implies that

$$\begin{aligned}
\frac{dD^*}{d\pi} &= -\frac{-\mu_S(1-\tau_S)\frac{d\theta^P}{d\pi}\mu_R[1-(1-\tau_R)\theta^P] + \mu_S[1-(1-\tau_S)\theta^P]\mu_R(1-\tau_R)\frac{d\theta^P}{d\pi}}{\{\mu_R[1-(1-\tau_R)\theta^P]\}^2} \\
&= \mu_S\mu_R\frac{d\theta^P}{d\pi}\frac{[1-(1-\tau_R)\theta^P](1-\tau_S) - [1-(1-\tau_S)\theta^P](1-\tau_R)}{\{\mu_R[1-(1-\tau_R)\theta^P]\}^2} \\
&= \mu_S\mu_R\frac{d\theta^P}{d\pi}\frac{\tau_R - \tau_S}{\{\mu_R[1-(1-\tau_R)\theta^P]\}^2}
\end{aligned}$$

Since $\mu_S\mu_R > 0$, we have

$$Sign(\frac{dD^*}{d\pi}) = Sign[\frac{d\theta^P}{d\pi}(\tau_R - \tau_S)]$$

Assumption $\Delta_1 < 0$ implies that $\tau_R < \tau_S\frac{\mu_S}{\mu_R} < \tau_S$ (using $0 < \mu_S < \mu_R$), so $\tau_R - \tau_S < 0$. Hence, $\frac{dD^*}{d\pi} > 0$ if and only if $\frac{d\theta^P}{d\pi} < 0$.

Recall that $\theta_t \equiv \pi_t p_G(\alpha_t) + (1 - \pi_t)p_B(\alpha_t)$, we have

$$\frac{d\theta^P}{d\pi} = \theta_\pi(\alpha_t, \pi_t) + \theta_\alpha(\alpha_t, \pi_t)\frac{d\alpha^P}{d\pi} \tag{17}$$

The first term of equation (17), $\theta_\pi(\alpha_t, \pi_t) = p_G(\alpha^P) - p_B(\alpha^P)$, is negative and captures the *direct* effect of π on θ^P: as investors and the constrained planner become more optimistic that the financial industry is strong (in which case the crisis would be less likely to occur), the unconditional crisis probability θ^P *perceived* by them tends to be lower. The second term $\theta_\alpha(\alpha_t, \pi_t)\frac{d\alpha^P}{d\pi}$ is positive and captures the *indirect* effect of π on θ^P: a more optimistic belief also induces investors and the constrained planner to increase their positions in the risky asset (α becomes higher), which in turn raises the unconditional crisis probability θ.

Rearrange equation (17) and we get:

$$\begin{aligned}
\frac{d\theta^P}{d\pi} &= \theta_\alpha(\alpha_t, \pi_t)[\frac{\theta_\pi(\alpha_t, \pi_t)}{\theta_\alpha(\alpha_t, \pi_t)} + \frac{d\alpha^P}{d\pi}] \\
&= \theta_\alpha(\alpha_t, \pi_t)(\frac{d\alpha^P}{d\pi} - \frac{d\alpha^C}{d\pi})
\end{aligned}$$

Since $\theta_\alpha(\alpha_t, \pi_t) > 0$, we have $\frac{d\theta^P}{d\pi} < 0$ if and only if $\frac{d\alpha^C}{d\pi} > \frac{d\alpha^P}{d\pi}$, that is, the SCC is satisfied (by Proposition 5). *Q.E.D.*

References

Acemoglu, D., Ozdaglar, A., Tahbaz-Salehi, A., 2013. Systemic Risk and Stability in Financial Networks. Technical report, National Bureau of Economic Research.

Acharya, V.V., 2009. A Theory of Systemic Risk and Design of Prudential Bank Regulation. *Journal of Financial Stability*, 5, 224–255.

Acharya, V., Pedersen, L., Philippon, T., Richardson, M.P., 2009. Regulating Systemic Risk. In *Restoring Financial Stability: How to Repair a Failed System*, edited by Acharya, V.V., Richardson, M.P.. John Wiley & Sons: Hoboken, NJ.

Acharya, V., Pedersen, L., Philippon, T., Richardson, M.P., 2010. Taxing Systemic Risk. In *Regulating Wall Street: The Dodd-Frank Act and the New Architecture of Global Finance*, edited by Acharya, V.V., Cooley, T., Richardson, M.P., Walter, I.. John Wiley & Sons: Hoboken, NJ.

Adam, K., Beutel, J., Marcet, A., Merkel, S., 2015. Can A Financial Transaction Tax Prevent Stock Price Booms? *Journal of Monetary Economics*, 76, S90–S109.

Adrian, T., Brunnermeier, M., 2016. CoVaR. *American Economic Review*, 106 (7): 1705–1741.

Axelson, U., Bond, P., Wall Street Occupations. 2015. *Journal of Finance*, 70 (5), 1949–1996.

Bank for International Settlements, 2010. 80th Annual Report. BIS.

Barberis, N., Greenwood, R., Jin, L., Shleifer, A., 2015. X-CAPM: An Extrapolative Capital Asset Pricing Model. *Journal of Financial Economics*, 115 (1), 1–24.

Barberis, N., Greenwood, R., Jin, L., Shleifer, A. Extrapolation and Bubbles. *Journal of Financial Economics*, Forthcoming.

Baron, M., Xiong, W., 2016. Credit Expansion and Neglected Crash Risk. *NBER Working Paper No. 22695.*

Basel Committee on Banking Supervision, 2010a. Countercyclical Capital Buffer Proposal.

Basel Committee on Banking Supervision, 2010b. Guidance for National Authorities Operating the Countercyclical Capital Buffer.

Berk, J., Green, R., 2004. Mutual Fund Flows and Performance in Rational Markets. *Journal of Political Economy*, 112, 1269–95.

Bhattacharya, S., Goodhart, C.A.E., Tsomocos, D., Vardoulakis, A., 2015. A Reconsideration of Minsky's Financial Instability Hypothesis. *Journal of Money, Credit and Banking*, 47 (5), August, 931–973.

Biais, B., Rochet, J.C., Woolley, J., 2015. Dynamics of Innovation and Risk. *Review of Financial Studies*, 28 (5), 1353–1380.

Bianchi, J., 2011. Overborrowing and Systemic Externalities in the Business Cycle. *American Economic Review*, 101 (7), 3400–3426.

Bianchi, J., Boz, E., Mendoza E., 2012. Macro-prudential Policy in a Fisherian Model of Financial Innovation. *IMF Economic Review*, 60 (2), 223–269.

Bolton, P., Santos, T., Scheinkman, J.A., 2016. Cream-Skimming in Financial Markets. *Journal of Finance*, 71 (2), 709–736.

Borio, C., 2012. The Financial Cycle and Macroeconomics: What Have We Learnt? *BIS Working Paper No. 395*. http://www.bis.org/publ/work395.pdf.

Boz, E., Mendoza, E.G., 2014. Financial Innovation, the Discovery of Risk, and the U.S. Credit Crisis. *Journal of Monetary Economics*, 62, 1–22.

Brownlees, C.T., Engle, R.F., 2012. Volatility, Correlation and Tails for Systemic Risk Measurement. *SSRN Working Paper 1611229.*

Brunnermeier, M., Sannikov, Y., 2014. A Macroeconomic Model with a Financial Sector. *American Economic Review*, 104 (2), 379–421.
Celerier, C., Vallee, B., 2015. Returns to Talent and the Finance Wage Premium. Working Paper.
Danielsson, J., Valenzuela, M., Zer, I., 2016. Learning from History: Volatility and Financial Crises. *Federal Reserve Board Finance and Economics Discussion Series*.
Dell'Ariccia, G., Igan, D., Laeven, L., Tong, H., Bakker, B., Vandenbussche, J., 2012. Policies for Macrofinancial Stability: How to Deal with Credit Booms. *IMF Staff Discussion Note*.
Drehmann, M., Borio, C., Gambacorta, L., Jiménez, G., Trucharte, C., 2010. Countercyclical Capital Buffers: Exploring Options. *BIS Working Paper No. 137*.
Drehmann, M., Borio, C., Tsatsaronis, K., 2012. Characterising the Financial Cycle: Don't Lose Sight of the Medium Term! *BIS Working Paper No. 380*. http://www.bis.org/publ/work380.pdf.
Espinoza, R.A., Segoviano, M.A., Yan, J., 2017. Systemic Risk: Bringing Together Theory and Measurement. Working Paper.
Farhi, E., Werning, I., 2016. A Theory of Macroprudential Policies in the Presence of Nominal Rigidities. *Econometrica*, 84 (5), 1645–1704.
Fernandez de Lis, S., Garcia-Herrero, A., 2012. Dynamic Provisioning: A Buffer Rather Than A Countercyclical Tool? Working Paper. BBVA Bank, Economic Research Department.
Gertler, M., Kiyotaki, N., 2015. Banking, Liquidity, and Bank Runs in An Infinite Horizon Economy. *American Economic Review*, 105 (7), 2011–2043.
Gordon, R., Kalambokidis, L., Slemrod, J., 2004. Do We Now Collect Any Revenue from Taxing Capital Income? *Journal of Public Economics*, 88, 981–1009.
Gordy, M.B., 2009. First, Do No Harm: A Hippocratic Approach to Procyclicality in Basel II. In: *Paper presented at the Conference Procyclicality in the Financial System*, jointly organised by the Netherlands Bank and the Bretton Woods Committee, 9–10 February.
Gordy, M.B., Howells, B., 2006. Procyclicality in Basel II: Can We Treat the Disease without Killing the Patient? *Journal of Financial Intermediation*, 15, 395–417.
Greenwood, R., Shleifer, A., 2014. Expectations of Returns and Expected Returns. *Review of Financial Studies*, 27 (3), 714-746.
Huang, X., Zhou, H., Zhu, H., 2012. Systemic Risk Contributions. *Journal of Financial Services Research*, 42(1–2), 55–83.
International Monetary Fund Staff Discussion Note (prepared by E. W. Nier, J. Osinski, L. I. Jacome, and P. Madrid), 2011. *Institutional Models for Macroprudential Policy*.
International Monetary Fund, 2014. *Staff Guidance Note on Macroprudential Policy*.
International Monetary Fund, 2017a. Getting the Policy Mix Right. *Global Financial Stability Report*, Chapter 1.
International Monetary Fund, 2017b. Is Growth at Risk? *Global Financial Stability Report*, Chapter 1.
Jeanne, O., Korinek, A., 2010. Excessive Volatility in Capital Flows: A Pigouvian Taxation Approach. *American Economic Review: Papers & Proceedings*, 100 (2), 403–407.
Jiménez, G., Ongena, S., Peydró, J.-L., Saurina, J., 2017. Macroprudential Policy, Countercyclical Bank Capital Buffers and Credit Supply: Evidence from the Spanish Dynamic Provisioning Experiments. *Journal of Political Economy* (forthcoming).
Jimenez, G., Saurina, J., 2006. Credit Cycles, Credit Risk, and Prudential Regulation. *International Journal of Central Banking*, 2, 65–98.

Kahou, M.E., Lehar, A., 2017. Macroprudential Policy: A Review. *Journal of Financial Stability*, 29, 92–105.
Kirti, D., 2018. Lending Standards and Output Growth. *IMF Working Paper No. 18/23.*
Kowalik, M., 2011. Countercyclical Capital Regulation: Should Bank Regulators Use Rules or Discretion? *Federal Reserve Bank of Kansas City Economic Review.*
Krznar, I., Morsink, J., 2014. With Great Power Comes Great Responsibility: Macroprudential Tools at Work in Canada. *IMF Working Paper No. 14/83.*
Laeven, L., Valencia, F., 2010. Resolution of Banking Crises: The Good, the Bad and the Ugly. *IMF Working Paper No. 10/146.*
McCoy, P., 2016. Countercyclical Regulation and Its Challenges. *Arizona State Law Journal*, 47 (5), 1181–1237.
Masciandaro, D., Passarelli, F., 2013. Financial Systemic Risk: Taxation or Regulation? *Journal of Banking and Finance*, 37(2), 587–596.
Minsky, H.P., 1992. The Financial Instability Hypothesis. *Working Paper No. 74, Jerome Levy Economics Institute.*
Packer, F., Zhu, H., 2012. Loan Loss Provisioning Practices of Asian Banks. *BIS Working Paper No. 375.*
Poledna, S., Thurner, S., 2016. Elimination of Systemic Risk in Financial Networks by Means of A Systemic Risk Transaction Tax. *Quantitative Finance*, 16 (10), 1599–1613.
Reinhart, C.M., Rogoff, K., 2009. *This Time Is Different: Eight Centuries of Financial Folly.* Princeton University Press, Princeton, NJ.
Schüler, Y.S., Hiebert, P.P., Peltonen, T.A., 2017. Coherent Financial Cycles for G-7 Countries: Why Extending Credit Can Be An Asset. *European Systemic Risk Board Working Paper Series No. 43.*
Segoviano, M.A., 2006. Consistent Information Multivariate Density Optimizing Methodology. Financial Markets Group, London School of Economics and Political Science.
Strohsal, T., Proano, C.R., Wolters, J., 2015. Characterizing the Financial Cycle: Evidence from A Frequency Domain Analysis. *Discussion Papers 22/2015, Deutsche Bundesbank.*
Zlatic, V., Gabbi, G., Abraham, H., 2015. Reduction of Systemic Risk by Means of Pigouvian Taxation. PLoS ONE 10(7): e0114928. https://doi.org/10.1371/journal.pone.0114928