ESSAYS ON LEARNING, INFORMATION, AND EXPECTATIONS IN MACROECONOMICS AND FINANCE

GENE PAUL GERARD CALDERON AMBROCIO

TESI DOCTORAL UPF / ANY 2015

DIRECTOR DE LA TESI
ALBERTO MARTIN and FERNANDO BRONER
Departament d'Economia i Empresa

Universitat
Pompeu Fabra
Barcelona

This thesis is dedicated to Eugenio Jr, Carmen, Eugenio III, Gene Phillip, Giancarla, Ma. Carmencita, and the late Emilio Antonio Jr.

Acknowledgements

This thesis would not have been possible without the continued support, guidance, and patience of my thesis advisers, Alberto Martin and Fernando Broner. I would also like to especially thank Jose-Luis Peydro, Filippo Ippolito, Kristoffer Nimark, Jaume Ventura, and Christian Brownlees. I am also very grateful to Barbara Rossi, Xavier Freixas, Javier Gomez-Biscarri, Giacomo Ponzetto, Andrea Polo, Vladimir Asriyan, and the rest of the faculty at the Universitat Pompeu Fabra Graduate Program in Economics, Finance, and Management and the Centre de Recerca en Economia Internacional.

I am similarly grateful for the help and support from my friends and colleagues at the UPF Ph.D. program namely, Tom Schmitz, Oriol Anguera, Fernando Perez, Miguel Martinez, Miguel Karlo de Jesus, Kizkitza Pastor-Biguri, Joerg Stahl, Stefan Pitschner, Jagdish Tripathy, Maria Paula Gerardino, Marc Goñi, Francesco Amodio, Bruno Caprettini, Andrei Potlogea, Francesc Rodriguez-Tous, Alain Schlaepfer, Tomas Williams, and many others.

Finally, I would also like to thank all the other participants of the International Lunch seminar at CREi and Finance Lunch at UPF as well as acknowledge financial and administrative support from the Universitat Pompeu Fabra with special thanks to Marta Araque.

Abstract

This dissertation focuses on learning and expectations formation in Macroeconomics and Finance and the role of information production in shaping macroeconomic fluctuations. The first chapter provides a theory of information production to explain two features of modern business cycles. In my theory information is produced along two dimensions, a pro-cyclical quantitative margin and a counter-cyclical qualitative margin, that generates both slow recoveries and episodes of ”rational exuberance” where optimistic booms tend to end in crises. The second chapter provides supporting evidence for the proposed cyclical variation in private information production using term loan data in the United States. Finally, the third chapter documents biases in terms of over-optimism and overconfidence in forecasts of real GDP growth from the survey of professional forecasters in the United States.

Resumen

Esta tesis estudia la formación de expectativas macroeconómicas y financieras y en el rol que tiene la producción de información sobre las fluctuaciones económicas. El primer capítulo desarrolla una teoría de la producción de información que genera una economía con recuperaciones lentas y ”booms” económicos que terminan en crisis económicas. En mi teoría, la información es producida a través de dos mecanismos, la cantidad de información producida es pro-cíclica y la calidad de la información es anti-cíclica. El segundo capitulo presenta evidencia usando datos de préstamos a plazos en EE.UU. para medir la producción de información privada durante el ciclo económico. Finalmente, el tercer capítulo documenta sesgos en términos de exceso de optimismo y exceso de confianza en las predicciones de crecimiento del PBI real usando la ”Survey of Professional Forecasters” en el Estados Unidos.

Preface

Modern macroeconomics has moved on from the full information rational expectations assumption. Although it is still a useful benchmark, alternative specifications for the expectations formation process have become necessary in order to match modern macroeconomic stylized facts and help provide relevant guidance to macroeconomic policy. However, today's macroeconomists are presented with a plethora of options and a key challenge is to provide tractable and insightful models which are also consistent with the data. This thesis provides both theoretical and empirical contributions to this endeavor.

The first chapter focuses on the role of time-varying information production in characterizing the dynamics of business cycles. In this chapter, I provide a theory of information production in which the private incentives to engage in activities that reveal information about the underlying state of the economy varies over phases of the business cycle. I develop a model of information production and learning in credit markets which generates booms fueled by optimism and slow recoveries hampered by pessimistic beliefs. In my theory, information about aggregate fundamentals is produced along two dimensions. First, optimistic beliefs lead to a fall in private investment in information reducing the *quality* of information available about aggregate fundamentals, an intensive margin. This gives rise to episodes of rational exuberance where optimism sustains booms even as fundamentals decline in the buildup to crises episodes. Second, the *quantity* of information is increasing in the level of economic activity, an extensive margin. Thus, recoveries are slow since pessimistic beliefs induce low levels of real investment and output which in turn provides little information about improvements in the state of the economy. Consistent with model predictions, I find supporting evidence in terms of a U-shaped pattern in macro-uncertainty measures over the course of the last two expansions for the United States.

In the second chapter, I provide a simple framework to assess the extent to which private information is used and revealed in credit markets using over two decades of term loan data from the United States. I decompose the variation in term loan spreads into that which can be explained by publicly observable borrower, lender, and macroeconomic conditions and a residual variation which I attribute to private information. I show that, and consistent with the theory proposed in the first chapter, my measure of private information production in the term loan market in the U.S. exhibits a hump-shaped pattern over the business cycle. In particular, I document evidence to suggest that private information production in these markets decline well before the start of the last two recessions and bottoms out about two years after the onset of the recession. I also find some evidence that the my measure of private information is associated with variation in the stock market returns of borrowers' equity around the days of the loan agreement date. These findings provide complementary evidence to the largely theoretical literature on credit cycles driven by variations in screening intensity and private information production.

The third chapter looks at the survey of professional forecasters in the United States to evaluate the extent of both optimism and confidence biases in their forecasts of real GDP growth. Timely and reliable forecasts of macroeconomic conditions are crucial inputs to economic agents and policy-makers alike. Surveys of professional forecasters provide such information and the rising popularity and use of density forecasts provide users of such surveys not only with mean predictions but also a measure of uncertainty, in terms of the predicted variance, of macroeconomic variables. In this chapter, I implement a simple framework which decomposes the predictive accuracy of density forecasts into biases according to their mean prediction (over or under-optimism) and variance (over- or under-confidence). First, at the individual-forecast horizon level, I find some evidence for optimism biases. On the other hand, the majority of professional forecasters appear to be overconfident. Second, the degree of overconfidence appears to be increasing in the forecast horizon. Sample estimates of the biases over a 5-year rolling window of survey responses also indicates that overconfidence seems to be more pronounced in the last business cycle.

Contents

List of Figures

List of Tables

Chapter 1

RATIONAL EXUBERANCE AND INFORMATION PRODUCTION

1.1 Introduction

Modern business cycles are asymmetric. Although recessions are typically sharp and short, the recovery process tends to be prolonged and gradual. One theory to account for this is that the speed of learning co-varies with economic activity.[1] Expansions are slow to start because they do so from periods of low economic activity which reveal little about the state of the economy. The opposite is true at the height of expansions, when production and investment are high and learning is quick. Although very insightful, this theory has an important limitation. In many instances, expansions that end in financial crises seem *ex post* to have been fueled by excessive optimism, which sustained economic activity and led agents to disregard negative signals about the state of the economy. The recent Financial Crisis is a case in point. According to the report of the Financial Crisis Inquiry Commission: *"this financial crisis was avoidable.... there were warning signs. The tragedy was that they were ignored or discounted"*. This is of course problematic for the theory of learning mentioned above. After all, should not the high levels of investment and production at the height of an expansion produce a very precise signal of any changes in the underlying economic fundamentals? In this paper,

[1]See for example Veldkamp (2005); Van Nieuwerburgh and Veldkamp (2006); Ordonez (2013), and Fajgelbaum et al. (2013).

I develop a model of information production to show why this may not be the case.

The main feature of my model is that the precision of information about the state of the economy depends on two channels. On the one hand, the state of the economy is partially revealed by observing the outcome from running investment opportunities. This is the channel emphasized in the literature, in which the production of information is increasing in the level of economic activity - *an extensive margin*. On the other hand, agents in my model can also invest in producing information about each project, and ultimately about the state of the economy, before undertaking them - *an intensive margin*. The key result is that this type of information production may decline once beliefs become highly optimistic. The reason is simple: economic agents have little reason to invest in more information about projects if they expect that most of them will do well. Therefore, a central prediction of my model is that the amount of information produced on the underlying state of the economy may decline as the economy approaches the peak of an expansion.

To formalize this theory, I embed these mechanisms in a simple model of credit markets with Bayesian learning about an unobserved aggregate fundamental and where there are no other distortions to financial intermediation. Resources and investment opportunities (projects) are separately endowed to two types of agents, financiers and entrepreneurs. A cycle emerges in that the fundamental, the aggregate quality of the pool of entrepreneurs, may be in one of two states, a high state with a large proportion of good projects and a low state with less. The fundamental is persistent and the transition between the two states is symmetric. Agents learn about the aggregate state by observing past credit market and production outcomes. In this environment I include a costly screening technology by which an individual lender may produce information about a particular borrower's type increasing the quality of information revealed in markets.

In this setting, the acquisition of private information today (a particular borrower's type) affects the quality of public information available tomorrow (the average quality of all borrowers). In turn prior beliefs determine how much information will be generated the next period. To simplify the narrative, I abstract from conventional frictions to financial intermediation such as asymmetric information and imperfect banking competition. In the model, agents make choices to maximize the expected returns from running these projects where the only constraint to achieving the first-best outcome is imperfect information about the aggregate state and the quality of individual projects.

As in the current literature, my model generates asymmetric business cycles with slow recoveries. Nevertheless, once the economy enters a period of high optimism, it tends to stay there. When fundamentals deteriorate in these periods, the low levels of investment in private information generate weak signals. Agents are less likely to perceive the fall in fundamentals and remain optimistic - what I refer to as rationally exuberant. These episodes are likely to occur in economies and periods of time where the relative cost to producing information is within intermediate levels. This implies that financial development, by gradually easing the cost to private information production, may entail a transition period where business cycles are more volatile and crises episodes occur more frequently. Further, information is under-produced in the competitive equilibrium. This occurs both at the onset of recoveries and at the height of optimistic booms. Thus, policies which promote credit and investment during recessions, by generating more information, can help speed up the recovery process. In addition, macro-prudential policies which effectively limits the expected volatility of financial intermediary profits particularly during credit booms may also induce more private information production and help mitigate the likelihood of rational exuberance and crises episodes.

I also document evidence on macro-uncertainty data from the U.S. that is consistent with the model mechanisms. I show that measures of macro-uncertainty exhibits a U-shaped pattern over expansions. Although initially declining, macro-uncertainty appear to increase several years prior to NBER-dated peaks in economic activity. This pattern is present in two widely used measures, forecaster disagreement and the *VIX*[2], which I complement with a forecast uncertainty index constructed from the average diffusion of individual forecasts. Through the lens of the model, such patterns in macro-uncertainty arise due to the low levels of information production both at the beginning of expansions and after they peak.

Related Literature. My model complements the literature on imperfect information and uncertainty as important drivers of short-run fluctuations.[3] My paper provides an information production mechanism as way to endogenize uncertainty *shocks*. Through the extensive margin to information production, I incorporate the positive feedback between

[2]The *VIX* corresponds to the option-implied expected 30-day volatility of the *S&P 500 Index*.

[3]Examples are Beaudry and Portier (2006); Collard et al. (2009); Blanchard et al. (2013); Bloom (2009); Barsky and Sims (2011); Eusepi and Preston (2011); Bloom et al. (2012); Christiano et al. (2014).

economic activity and the precision of information in current models of social learning over the business cycle such as Veldkamp (2005); Van Nieuwerburgh and Veldkamp (2006); Ordonez (2013), and Fajgelbaum et al. (2013). Unlike these models which imply a monotonic relationship between beliefs and information production, the inclusion of the intensive margin in my model allows for the occasional highly optimistic boom where information production falls and generates rising uncertainty.[4]

The intensive margin in my model draws from the literature on boom-bust episodes as arising from informational cascades and herding.[5] In my model, a similar feature appears during periods of high optimism at the heights of expansions. I refer to these as episodes of rational exuberance where agents rationally discount warning signs because they are thought to be less precise.This is in contrast to models with adaptive or rules-based learning such as Bullard et al. (2010) or behavioral biases as in Gennaioli et al. (2013). Further, agents in my model do not internalize the information benefits that accrue to agents in future periods and thus information is under-produced as in Burguet and Vives (2000) who abstract from business cycles. In my model, this externality is strongest both at the beginning and at the heights of expansions.

I use screening in credit markets as the means by which private information is produced. A reduction in the incentives to use this technology at the height of booms appears prominently in the credit screening literature (Ruckes, 2004; Berger and Udell, 2004; Dell'Ariccia and Marquez, 2006; Gorton and Ordonez, 2014a). The particular screening mechanism I employ in the model is adapted from learning about collateral in Gorton and Ordonez (2014a,b). My mechanism differs in two key points. First, I introduce heterogeneity across entrepreneurs such that for any given point in time a proportion of borrowers gets screened. This smooths the transition between periods of high and low screening. Second, I focus on producing information about the productivity of projects themselves and effectively endogenize the formation of beliefs about the aggregate fundamental. As such, this paper complements crises models arising from a shock to beliefs or investor sentiment.[6]

---

[4]An alternative view is that uncertainty rises prior to or during crashes as these are realizations of unusual or rare events as in Nimark (2014) and Orlik and Veldkamp (2014).

[5]Related examples from the literature are Bikhchandani et al. (1992),Chamley and Gale (1994), Caplin and Leahy (1994), Chari and Kehoe (2003), Chamley (2004), and Broner (2008).

[6]For example see Boz (2009) for crises in developing economies, and Boz and Mendoza (2014), Martin and Ventura (2011), Gennaioli et al. (2013), and Gorton and Ordonez (2014a) for the recent U.S. Financial crisis.

The rest of the chapter is organized as follows. The next section develops a model of credit markets to illustrate information production. Section 1.3 shows the business cycle implications and Section 1.4 provides some evidence. Finally, Section 1.5 concludes.

1.2 Credit markets and information production

1.2.1 Setup

Let us focus on a simple economy where resources and investment opportunities are separately endowed to two types of risk-neutral agents who maximize end-of-period consumption. We have a set $M \in \mathbb{N}$ financiers with an initial endowment W of a consumption good at the beginning of the period. Then we have $N \in \mathbb{N}$ entrepreneurs each of which have an investment opportunity - a project - in which some investment K of the consumption good made at the beginning of the period can potentially yield Y consumption goods at the end of the period. A credit market exists to facilitate the transfer of resources from financiers to entrepreneurs.

However, the return to investing in each project depends on two factors which differ across entrepreneurs. In particular, if an entrepreneur i borrows K_i then he may obtain Y_i according to the following production function:

$$Y_i = \begin{cases} A_i K_i^\alpha & \text{with probability } \theta_i \quad 0 < \alpha < 1 \\ 0 & \text{with probability } 1 - \theta_i \end{cases}$$

The productivity parameter A_i differs across entrepreneurs and comes from some distribution with an upper and lower bound $F(A_i; \underline{A}, \bar{A})$. The success probability θ_i also (independently) differs across entrepreneurs. Some projects are "good" and have a high probability of success (θ_G) and others are "bad" with a low probability of success (θ_B). In the aggregate a proportion μ of projects are of the good type and this is the aggregate fundamental in the economy. The productivity parameter A_i is publicly observed by all agents but the success probability is not. In particular both the financiers and the entrepreneurs do not know the type of the projects they have. In this section only the aggregate proportion of good projects μ is known at the beginning of the period.

Both entrepreneurs and financiers have alternative uses for their time and resources - outside opportunities. Financiers have a savings technology which allows them to convert some or all of their endowments at the beginning of the period to more resources at the end of the period at the rate $1 + r_f$. This is the opportunity cost of financing in the model. On the other hand, if entrepreneurs do not run their project then they may use their time on an outside activity that yields a fixed output of w consumption goods at the end of the period. To abstract from potential distortions to the supply of credit, I assume that there are substantially more financiers than entrepreneurs and that there are sufficient endowments to fully finance all of the projects at their desired levels of investment.[7]

Finally, financiers have a screening technology which allows them to perfectly reveal a project's type at a cost of γ goods. In the credit market, financiers offer two types of loan contracts. They may provide credit without screening the borrower and thus the loan rate is conditioned on the average success probability of all entrepreneurs. Alternatively, they offer a screening contract where they first screen the borrower and then the entrepreneur may choose how much to borrow at the loan rate which is contingent on the screening outcome.[8]

In the market for loans each entrepreneur who wants to invest and borrow faces a menu of contracts from all financiers who will have to compete for the loan. I denote the screening contract and a no-screening contract with subscripts S and I respectively. Each contract details the (gross) rate of interest on the loan $R(\theta_i, K_i, S)$ and $R(\mu, K_i, I)$. Output from projects that are run are costlessly verifiable and ensures no strategic default in equilibrium. Given this menu, the entrepreneur chooses the loan type and size of the loan K_i and invests in the project. If financiers make identical offers, the entrepreneur is randomly matched to one of them. At the end of the period, projects that are run either succeed or fail, repayments are made and consumption takes place.

[7]These assumptions ensure that entrepreneurs have all the bargaining power in setting the terms of credit such that the participation constraint for financiers will bind and pinned down by the cost of financing and the probability of repayment.

[8]To rule out potential cross-sectional information spillovers as in Petriconi (2012), I assume that screening and lending are a packaged deal and a commitment to borrow from the same financier comes with screening.

Equilibrium definition and derivation

Equilibrium is defined as Sub-game Perfect Nash. Given the state of the economy μ_t, equilibrium is given by the set of choices:

1. A menu of contract offers consisting of screening loan rates $R_j(\theta_G, K_i, S)$, $R_j(\theta_B, K_i, S)$ for revealed good and bad types, and the no-screening interest rate $R_j(\theta_t, K_i, I)$ set by each each financier $j \in [1 : M]$ and for every potential match with the set of entrepreneurs $i \in [1 : N]$.

2. A set of participation, contract type, and loan size K_i choices for all entrepreneurs $i \in [1 : N]$

Here, θ_t reflects the average success probability given by the proportion μ_t of good projects. The time subscript t is used to denote an aggregate variable which agents will be learning about in the dynamic version of the model. These choices are made such that each financier and entrepreneur maximize expected end-of-period consumption.

1.2.2 Credit market equilibrium

Optimal contracts

The menu of screening (S) and no-screening (I) loan contracts set by each financier yield interest rates such that she is at least as well of as investing the loaned funds into her savings technology:

$$R(\theta_t, I) = \frac{1 + r_f}{\theta_t} \tag{1.1}$$

$$R(\theta_i, S) = R(\theta_i, I) + \frac{(1 + r_f)\gamma}{\theta_i K_i} \quad i \in \{G, B\} \tag{1.2}$$

Competition guarantees that this participation constraint holds with equality. Moreover, in the absence of information asymmetry and strategic default, the expected probability of repayment is equivalent to the expected success probability of the average entrepreneur. For the screening contract we have that the participation constraint for the financier, which now includes the cost of screening, must hold for both realizations.

Screening and participation choices

With this menu of interest rates, the entrepreneur may then decide on the optimal size of borrowing:

$$K_i^* = \left[\frac{\alpha\theta_i A_i}{1+r_f}\right]^{\frac{1}{1-\alpha}} \tag{1.3}$$

where $\theta_i \in \{\theta_t, \theta_G, \theta_B\}$ depending on whether the entrepreneur chooses the screening or no-screening contract and the screening outcome.[9] In turn, expected profits from either the screening or the no-screening contracts are:

$$\mathbb{E}[\pi_i|\mu_t, I] = \theta_t^{\frac{1}{1-\alpha}} A_i^{\frac{1}{1-\alpha}} \Lambda \tag{1.4}$$

$$\mathbb{E}[\pi_i|\mu_t, S] = (\mu_t\theta_G^{\frac{1}{1-\alpha}} + (1-\mu_t)\theta_B^{\frac{1}{1-\alpha}}) A_i^{\frac{1}{1-\alpha}} \Lambda - \gamma(1+r_f) \tag{1.5}$$

where $\Lambda \equiv \left[\frac{\alpha}{(1+r_f)^\alpha}\right]^{\frac{1}{1-\alpha}} \left[\frac{1-\alpha}{\alpha}\right]$. Figure 1.1 plots expected profits under both contracts across different levels of productivities on the horizontal axis. The solid line (**I**) is for the no-screening contract whereas the dashed line (**S**) refers to the screening contract profits. Screening profits are initially lower due to the additional cost of screening but have a steeper incline as higher productivity implies a larger loan size to leverage the information gains.

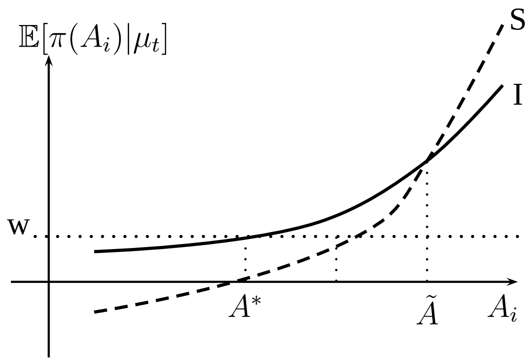

Figure 1.1: Profits over productivity levels

The point of intersection identifies a threshold productivity $\tilde{A}$ above which screening is preferred. Similarly, there is another threshold A^* such that investing in the project

[9] In the case of the screening contract, expected profits are not always positive given the screening cost. In these cases, optimal loan sizes are zero.

yields more profits than the outside opportunity w. Combining equations 1.4 and 1.5 yield the following:

$$\tilde{A}(\mu_t) = \left[\frac{\gamma(1+r_f)}{Z(\mu_t)\Lambda}\right]^{1-\alpha} \tag{1.6}$$

$$A^*(\mu_t) = \frac{1}{\theta_t\Lambda} min\{w, \frac{w+\gamma(1+r_f)}{\zeta(\mu_t)}\} \tag{1.7}$$

where $Z(\mu_t)$ and $\zeta(\mu_t)$ reflects the private value of information.[10] In this economy, the entrepreneurs themselves do not know their own type. We can think of such information as market- or sector-relative characteristics which financiers, by interacting with multiple entrepreneurs, are better able to evaluate. Screening gains in this model is driven by the convexity of expected profits. Rather than investing the 'average' loan size, one can choose to be screened and thus condition the size of investment to their revealed type.[11] The model generates a hierarchy in the chosen activity given productivities. Productive entrepreneurs will engage in borrowing and the most productive can afford to acquire information while doing so. The proportion of agents who engage in these types of action determines the information revealed about the aggregate fundamental. Engaging in the outside activity and not running the project produces no information.

How does the aggregate state affect the decision to be screened? Figure 1.2 plots expected profits for both the screening contract (equation 1.5 as thick dashed line) and the no-screening contract (equation 1.4 as solid line) over the range of μ_t on the horizontal axis for some productivity level A_i. When there is a large proportion of good projects, the value of screening is relatively low and does not merit the cost. The no-screening interest rate would be very similar to the interest rate that would be charged if one turns out to be of the good type. On the other hand, if the proportion is low then the likelihood of being the bad type is high and entrepreneurs are likely to get worse loan terms when screened.

[10] $Z(\mu_t)$ is given by $\mathbb{E}[\theta_i^{\frac{1}{1-\alpha}}|\mu_t] - \mathbb{E}[\theta_i|\mu_t]^{\frac{1}{1-\alpha}}$ and $\zeta(\mu_t) = \frac{\mathbb{E}[\theta_i^{\frac{1}{1-\alpha}}|\mu_t]}{\mathbb{E}[\theta_i|\mu_t]^{\frac{1}{1-\alpha}}}$. Note that $A^*(\mu_t) = A^*(\mu_t, I)$ whenever the screening threshold is above the participation threshold. This is the case whenever $\gamma \geq \zeta(\mu_t)\left(\frac{\Lambda}{1+r_f}\right)$

[11] This is an extended version of the same mechanism found in (Gorton and Ordonez, 2014a). A key difference is that in our model screening changes the loan rate and allows for the investment size to optimally adjust to both better information about productivity and the new interest rate.

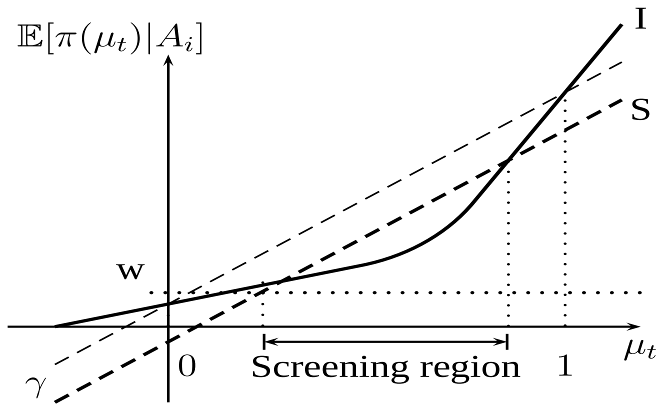

Figure 1.2: Profits over beliefs

To see how the proportion of entrepreneurs engaging in each of these activities vary over values of the fundamental, in Figure 1.3 I plot the screening and participation cutoffs over a range of states from μ_L to μ_H on the horizontal axis. On the vertical axis we have the range of observable productivities. The area above the dashed line (representing the cutoff productivity level to participate A^*) correspond to the measure of firms engaging in production and the area above the solid line (representing the cutoff productivity such that screening is optimal $\tilde{A}$) correspond to the measure of firms who would choose to be screened.

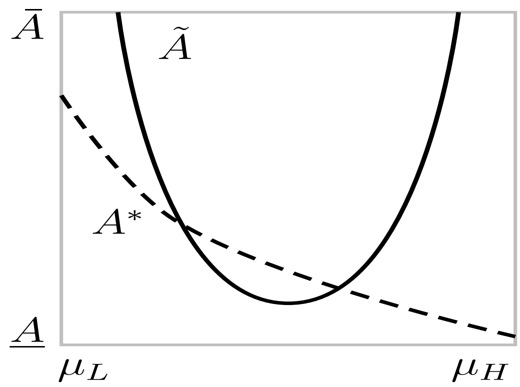

Figure 1.3: Relative values of cutoffs

In my economy, participation in credit markets is increasing in the state of the economy (the proportion of good projects). On the other hand, screening intensity, the proportion of screened projects, has a U-shaped pattern. Very few projects are screened when the aggregate fundamental is 'too' low or 'too' high. In the next section, the relative proportion of projects which are (i) not run, (ii) run and are not screened, and (iii) projects which are both run and screened, determine the precision of information about the aggregate fundamental.

1.3 Business cycle dynamics

We now move on to the dynamic version of the model and characterize the evolution of business cycles arising from endogenous persistence in beliefs. Consider now that the aggregate state of the economy follows a Hidden Markov Chain and agents may learn about it by observing aggregate credit market outcomes. In particular the proportion of good projects may either be high (μ_H) or low (μ_L). Further, the aggregate fundamental is persistent. With probability $\lambda > 0.5$, the proportion of good projects today will be the same as in the previous period. Equilibrium in credit markets for each period is as defined in the previous section and I now impose Bayesian learning on the process by which agents learn about the state of the economy. Equilibrium is defined as Subgame Perfect Nash given a sequence of realizations of shocks for the state of the world, entrepreneur types, and project outcomes. For any period t, prior beliefs about the state of the economy $\mathbb{E}[\mu_t|\mathcal{I}_{t-1}]$ are formed given all past information.

At the end of the period, aggregate statistics on credit market contracts and project outcomes become available. First, they observe the number of entrepreneurs who chose to be screened and the aggregate screening outcome. Second, they observe the number of projects which are run for each type of contract chosen and the number of projects which succeed within each category. For entrepreneurs that are screened, the relevant statistic is the proportion of those who were found to have θ_G. For the rest of the entrepreneurs that participate in credit markets, project successes are used to infer the state of the world. Denote the aggregate number of projects run and their composition into screened (S) and un-screened projects (I) as $n_t = n_t^S + n_t^I$. Define the number of successful projects in n_t^I as s_t^I and the number of screened projects (in n_t^S) with success probability θ_G as s_t^S. We then have two sets of signals which may be used to infer the state of the world at time t which we denote with $\Sigma_t \equiv \{\frac{s_t^S}{n_t^S}; \frac{s_t^I}{n_t^I}\}$.[12]

Beliefs about the current state are formed given the full history of past information. At the end of the period, these beliefs are updated with information generated at time t to form a posterior belief given the new information. That is, belief-formation has a dynastic flavor wherein agents born at time t have the information set $\mathcal{I}_{t-1}$ which is recursively updated with outcomes in each period $\mathcal{I}_t = \mathcal{I}_{t-1} \cup \Sigma_t$.

Given some prior probability that $\mu_t = \mu_H$ denoted by $p_{t|t-1}$, beliefs are updated using the set of signals from the unscreened and screened outcomes (Σ_t) to form posterior beliefs. The information content of current signals is quantified by the likelihood ratio between the high and low states. Let $\mathcal{L}(\mu_H|\Sigma_t)$ be the likelihood of a high state given period t signals and $\mathcal{L}(\mu_L|\Sigma_t)$ the corresponding likelihood of a low state. Then, posterior beliefs combine priors and period t signals using Bayes' theorem[13]:

$$p_{t|t} = \frac{\mathcal{L}(\mu_H|\Sigma_t)p_{t|t-1}}{\mathcal{L}(\mu_H|\Sigma_t)p_{t|t-1} + \mathcal{L}(\mu_L|\Sigma_t)(1-p_{t|t-1})} \tag{1.8}$$

The numerator is the joint likelihood of the high state and the denominator normalizes the likelihood with the sum of both joint likelihoods. Finally the optimal forecast for the aggregate state of the world is given by the posterior belief and the persistence parameter:

$$p_{t+1|t} = \lambda\, p_{t|t} + (1-\lambda)(1-p_{t|t}) \tag{1.9}$$

Equations 1.8 and 1.9 characterize the evolution of the state variable (prior beliefs) given a sequence of realizations $\{\Sigma_i\}_{i=1}^t$ and an initial prior $p_{1|0}$. My learning mechanism generates unbiased forecasts whose precision depend on endogenous information production.[14] In turn, the precision of information determine how much beliefs adjust with respect to period t shocks. In the next section I describe how the two margins to information production determine the precision of period t information.

[12]Both sets of signals are informative about the aggregate state since $s_t^S \sim Bin(\mu_t, n_t^s)$ and $s_t^I \sim Bin(\theta_B + \mu_t(\theta_G - \theta_B), n_t^I)$.

[13]For two competing hypotheses (H and L), evidence Σ, and some prior probability of one hypothesis $Pr(H) = p$, Bayes' theorem is $Pr(H|\Sigma, p) = \frac{Pr(\Sigma|H)p}{Pr(\Sigma|H)p + Pr(\Sigma|L)(1-p)}$

[14]Note that although the ratios $\frac{s_t^S}{n_t^S}$ and $\frac{s_t^I}{n_t^I}$ are exogenously driven by the true state of the world μ_t, the precision of period t information given by n_t and the ratio $\frac{n_t^S}{n_t}$ are endogenous and depend on beliefs $\mu_{t|t-1}$.

1.3.1 Information Production

At any given period, the cutoffs $A^*(\mu_{t|t-1})$ and $\tilde{A}(\mu_{t|t-1})$ are sufficient to characterize equilibrium in credit markets. In turn, the cutoffs are given by the state variable $\mu_{t|t-1}$. Finally, the next generation's beliefs are going to be updated with new information Σ_t. We can then examine how beliefs about the state of the world affect the number and type of signals the economy generates. First, the quantity of information is monotonically increasing in beliefs.

Proposition 1 (The extensive margin to information production)**.** *Given a constant proportion of screening to undertaken projects, information production is non-decreasing in beliefs about the state of the world* $\mu_{t|t-1}$

Proofs for all the propositions are in the Appendix. The extensive margin to information production generates a negative relationship between persistence and optimism. This makes business cycles asymmetric with slower recoveries than expansions as the quantity of information produced is pro-cyclical. These economies are characterized by asymmetry in persistence and would exhibit "*slow booms*" and "*sudden crashes*" (e.g. Veldkamp 2005). The key innovation of the paper is that the quality of information is also changing over time. In this economy, two types of signals may be produced and one is more informative about the state than the other. This is what I refer to as the intensive margin. Along this dimension, informatio production is hump-shaped over prior beliefs.

Proposition 2 (The intensive margin to information production)**.** *Given a fixed number of projects undertaken and whenever* $\underline{\gamma} < \gamma < \bar{\gamma}$*, information production is hump-shaped in beliefs about the state of the world.*

Corollary 1. *When* $\gamma \geq \bar{\gamma}$*, screening never takes place and* $n_t = n_t^I \quad \forall \mu_{t|t-1}$*. When* $\gamma \leq \underline{\gamma}$ *then all credit market participants choose to be screened and* $n_t = n_t^S \quad \forall \mu_{t|t-1}$*.*

The intensive margin follows the literature on credit screening with hump-shaped incentives to screen. The value of screening and acquiring private information is inversely proportional to the 'strength' of prior beliefs about being in a particular state which in this case is when $\mu_{t|t-1}$ tends to either μ_H or μ_L. This allows for the possibility of rational exuberance episodes where information is seemingly abundant and yet appears to be discounted. As screening intensity is decreasing in γ, there exists an upper and lower bound to the cost of screening such that screening is either always preferred or never engaged in for any set of beliefs. That is, a corollary of proposition 2 is that there exists a $\bar{\gamma}$ such that no screening ever takes place. Conversely, there exists a $\underline{\gamma}$ such that all agents who participate in credit markets are always screened. For intermediate ranges of the screening cost, the intensive margin kicks in and the precision of date t information about the state is jointly determined by both the quantity and quality of signals. The next proposition outlines the overall evolution of information production in my economy depending on the relative cost to producing information.

Proposition 3 (Dynamics of belief persistence)**.** *The speed of learning, and hence the persistence of beliefs, is characterized by the following:*

1. *When $\gamma \leq \underline{\gamma}$, screening always occurs and the persistence of beliefs is monotonically decreasing in $\mu_{t|t-1}$. Similarly, when $\gamma \geq \bar{\gamma}$, no screening takes place and the persistence of beliefs is also monotonically decreasing in $\mu_{t|t-1}$. However, information production (persistence) is everywhere lower (higher) than in the previous case.*

2. *When $\underline{\gamma} < \gamma < \bar{\gamma}$, and under some parameter conditions on the relative informativeness of screening to no-screening signals, $\exists \mu^*$ for which persistence is decreasing in beliefs within the range $\mu_{t|t-1} \in [1-\lambda \quad \mu^*]$ and increasing in beliefs for the range $\mu_{t|t-1} \in [\mu^* \quad \lambda]$.*[15]

Persistence is inversely proportional to the precision of period t information which is jointly determined by the extensive and intensive margins. Depending on parameter values, my model allows for persistence to increase when beliefs become too optimistic giving rise to protracted booms where beliefs remain relatively optimistic even when negative shocks appear.

Discussion

Before characterizing business cycles with endogenous learning, I first discuss some of the key assumptions behind my model of credit markets. In this economy, variations

[15]The parameter conditions require that the screening signals are sufficiently more informative than the no-screening signals such that a fall in the relative proportion of screening to total signals may generate a reduction in the overall precision of information despite the increase in the total number of signals. See the proof in the appendix for details.

in the credit market equilibrium over time depend solely on differences in the expected proportion of good to bad projects. I have assumed constant other factors that may change the proportion of screening and participating agents. For instance, if the outside opportunity cost w were pro-cyclical then the extensive margin would be attenuated. The same would be true if the cost of funding $1 + r_f$ were to be counter-cyclical. One may also think that the screening cost γ varies over time and is likely to decline as the financial sector becomes more developed. Nevertheless, as the recent Financial Crisis has shown, financial innovation - through the introduction of new and more complex instruments and assets - may keep pace with financial development and keep the cost of acquiring further information relevant even in today's financial markets.

Unlike the bulk of the literature on credit market screening, I have not included a role for externalities arising from competition in the banking sector or information asymmetry between borrowers and lenders. This is done largely to simplify the model and isolate the effects of my mechanism. In my setup, the choices in the competitive equilibrium maximize the value of investing in projects given the information that these agents have and allocating bargaining power to financiers will lead to the same results. Further, introducing asymmetric information is likely to produce the same predictions. If borrowers knew their type, then Bad types would always prefer not to be screened and only the low productivity good types would join the pool of unscreened borrowers. When aggregate fundamentals improve then a larger proportion of the unscreened borrowers will be of the good type which lowers the no-screening interest rate and may further increase the fraction of Good types in the unscreened pool.

1.3.2 The evolution of business cycles with endogenous learning

In this chapter I focus on an economy satisfying case 2 of proposition 3. Under this setting, the model generates the following predictions. First, my economy tends to stay longer in periods of highly optimistic beliefs relative to moderately optimistic ones. This results in a relatively higher frequency of periods where beliefs are and remain at a higher threshold level of optimism. Second, when fundamentals deteriorate during these highly optimistic periods it takes a while before a recession appears. Further, uncertainty in the model would already have been rising prior to the start of the recession. Finally, these occur on top of the pro-cyclical asymmetry in business cycles that the extensive margin to information production generates.

I illustrate these main predictions by simulating my economy and examining the evolution of expansions and recessions. To document how the two margins to information production affect the dynamics of business cycles, I compare the model simulation against two benchmarks, an economy $Extensive$ which holds the proportion of screened to unscreened projects constant (the average in $Model$) but vary the total number of projects run as per the extensive margin, and $Constant$ where both the proportion and the total number of projects run are constant.

I chose the parameters such that in the high state 60 percent to be of the good type and 98.6 percent of all projects succeed where the same figures are 50 percent and 98.5 percent for the low state respectively. The expected duration of a state is 20 periods. The values for outside opportunity w and screening cost γ where chosen such that all projects which are invested in are also screened for all but highly optimistic beliefs.[16] Finally, I use a uniform distribution for A_i and limit the maximum number of projects to 20.[17]

I simulate for 50,000 periods with an initial prior of one half and estimate business cycle turning points by filtering out high frequency fluctuations[18]. Given the filtered output I define a peak (or the start of a recession) as a point in time with the local maximum of filtered output. A trough (end of the recession) is identified similarly. This gives us between 992 to 1,104 full cycles across simulations. In Table 1.1 I first report the skewness of the growth rate of investment and changes in the no-screening interest rate.

[16] In the simulation, screening intensity begins to fall when the average success probability is about 98.57 percent.

[17] I allocated one agent each to 20 equally-spaced points within $\underline{A}$ and $\bar{A}$ including the boundaries. The relatively small difference between the high and low state proportions were chosen so as to allow for a sufficiently large number of agents N in the simulation while still generating variations in the precision of signals across ranges of prior beliefs. See for example the calibration of the endogenous learning model in Veldkamp (2005) and Ordonez (2013).

[18] This was done using a Hodrick-Prescott filter set to 14400.

Table 1.1: Skewness estimates

	Investment	Loan rates
Model	-0.48	0.25
Extensive	-0.49	0.48
Constant	0.04	0.06

First, constant information production generates symmetric changes in these variables whereas in simulations $Extensive$ and $Model$, interest rate changes are right-skewed and investment growth is negatively skewed. These are features of asymmetric business cycles consistent with slow recoveries generated by the extensive margin to information production. This asymmetry also appears in the frequency of pessimistic periods relative to optimistic ones. In Figure 1.4, I report the frequency (in percent) of periods with beliefs given on the horizontal axis.

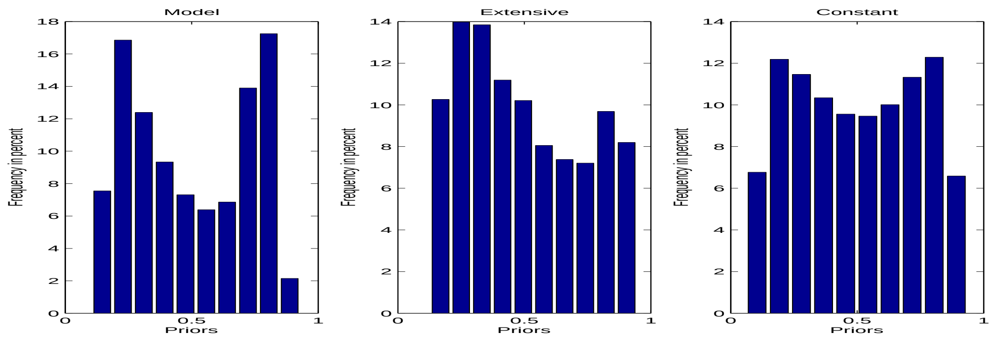

Figure 1.4: Frequency of pessimistic and highly optimistic beliefs

The distribution of beliefs for $Constant$ is symmetric whereas pessimistic beliefs ($p_{t|t-1} < 0.5$) occur two to five percent more often in $Model$ and $Extensive$. Through the intensive margin, highly optimistic beliefs also occur more frequently in $Model$ relative to $Extensive$. Beliefs higher than 75 percent probability of the High state occur four percent more often in $Model$ relative to $Constant$ or $Extensive$.

We now turn to how expansions end in my economy. I collect all peaks in filtered output and compute the average evolution of participation, screening, and uncertainty in the simulations. As the model predicts, in $Model$ participation increases as we near the peak in economic activity while screening falls. This generates a U-shaped pattern in average uncertainty. In Figure 1.5 I plot the average evolution of uncertainty for

$Model$, and $Extensive$ relative to that in $Constant$. On the horizontal axis we have periods before the peak in economic activity and a value of zero on the vertical axis means that uncertainty would be the same as in $Constant$.[19]

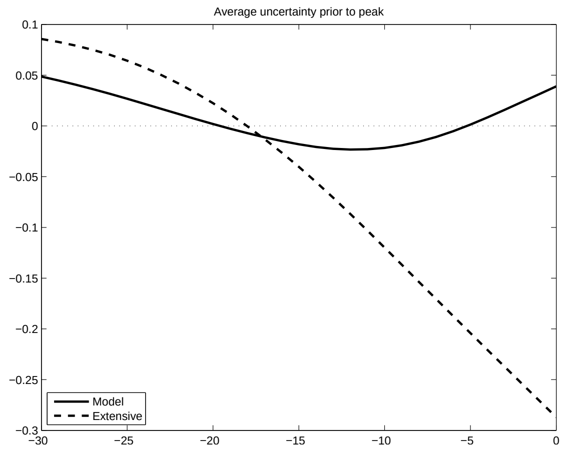

Figure 1.5: Average path of uncertainty

Here we see how the two margins contribute to macro-uncertainty. In $Extensive$, average uncertainty is monotonically decreasing as we approach the peak in economic activity. In $Model$, where both margins to information production are active, we get a U-shaped pattern. This, along with the predicted decline in private information production, is one of the predictions of the model for which we find supporting evidence from U.S. data.

The consequences of optimism

In this section, I simulate an episode of rational exuberance. I start the simulation from the most optimistic belief and then simulate a fall in fundamentals by hitting the economy with shocks from the Low state. I do this 1,000 times and report the average evolution of beliefs and entrepreneurial profits in the top left and right panel of Figure 1.6. I also report the actual and expected default rates in the bottom right panel and average uncertainty in the bottom left panel.

[19]Here it is important to evaluate uncertainty relative to $Constant$ as the two-state assumption in the model necessarily implies that uncertainty has a hump-shaped pattern over beliefs. If agents believe that it is the high state with probability one or zero then expected uncertainty is also necessarily zero.

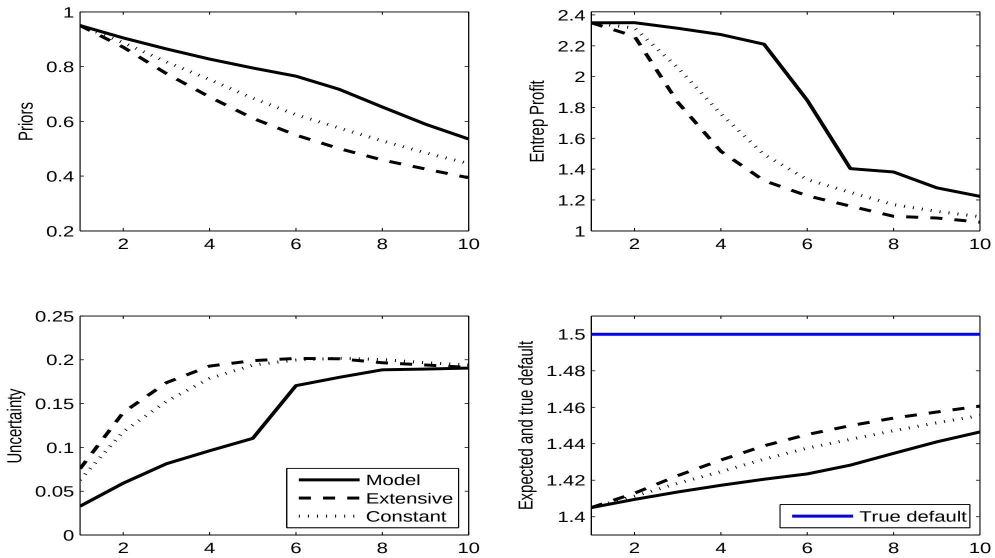

Figure 1.6: Simulated rational exuberance

For $Extensive$, high optimism leads to a lot of information production and learning is quick. On the other hand, optimistic beliefs slow down learning and induce persistence in our model and this implies two things. First, as the bottom right panel shows, on average financiers experience persistent and sizable surprises in the expected and actual default rates. Nevertheless, optimism also implies that a wider range of projects get financed and output does not fall significantly until about five periods into the simulation for $Model$. Relative to $Constant$ or $Extensive$, the recession would have been dated much later in our model and in this intermediate period, uncertainty would be rising, default rates are larger than expected, and yet the fall in output is not too large.

To an econometrician who observes and treats data for each period as being equally informative, our model generates an economy where agents appear to disregard warning signs and are *irrationally exuberant*. Who pays for this over-optimism? First, entrepreneurs who are screened get to know their types which makes prior beliefs irrelevant for their investment decisions. Thus optimism, by reducing the proportion of agents who are screened, also reduce the proportion of investments insulated from informational inefficiency. What about those who are not screened? The first implication of over-optimism is that the participation threshold is low. That is, a certain mass of entrepreneurs who would have taken the outside opportunity are instead investing in their projects. Nevertheless, the no-screening interest rate is also low. This leads to larger investment sizes for all unscreened borrowing. Further, as repayment is contingent on production being successful, it turns out that entrepreneurs do not suffer from overly optimistic beliefs. Instead, those who succeed in production would be making more profits than they would otherwise have. It is the financiers who shoulder the costs of being overly optimistic in our economy. The low interest rate regime brought about by optimistic beliefs effectively transfers wealth from financiers to successful entrepreneurs.[20] The opposite would be true in the recovery, when beliefs are pessimistic relative to the true state of the fundamental.

1.3.3 Socially optimal information production

One question that emerges is whether these types of fluctuations are inefficient. In the appendix, I consider the case where a constrained social planner maximizes the expected stream of current and (discounted) future profits by choosing the participation and screening thresholds. The social planner cannot influence belief-formation in any other way nor is she allowed to change the optimal loan contracts I have earlier derived for the economy. In my economy, information about the aggregate fundamental is increasing in participation and screening activity. However, the incentives to do either only depend on the private benefits from doing so. The information generated about aggregate fundamentals is a positive externality in my economy.

I show that the constrained social planner would choose lower thresholds than in the competitive equilibrium. By doing so, she makes beliefs less persistent. In the competitive equilibrium, agents take beliefs about the aggregate state as given and do not internalize the fact that more participation and screening activity generates more precise prior beliefs for the next period. Instead, the cutoffs that arise in the competitive equilibrium are the ones that maximize only the current period's expected profits. This creates space for a policy to improve on the competitive allocation by providing subsidies on information production and may be financed by a lump-sum tax on all agents. In particular, a policy which subsidizes both investment during pessimistic periods and screening during optimistic ones can implement the constrained social planner's solution.

[20] In the context of an open economy where financing comes from external sources, this would imply that the rest of the world is subsidizing the excesses of optimism in my economy. This temporarily keeps the economy in a boom until such a time that a sufficient set of negative surprises triggers information production and financing dries up leading to a large crash and a sudden stop or reversal in capital flows.

This result suggests that current macro-prudential policies which stimulate credit and economic risk-taking during recoveries and limit financial risk-taking by financial intermediaries, especially at the heights of optimistic credit booms, may generate informational gains that are currently not considered in the policy debate. Stimulating credit during recoveries increases the precision of information through the extensive margin at a time when financial intermediaries are already likely to be screening borrowers but where credit provision in the aggregate is limited. On the other hand, through the lens of the model, macro-prudential policies which limit the expected volatility of the performance of financial intermediaries' loan portfolios (or their profits) especially during periods of high credit growth would induce private information production. This would reduce the likelihood of crises following from rational exuberance episodes I describe in this paper.

1.4 Supporting Evidence

In this section I document several features of uncertainty over the business cycle using quarterly U.S. data. I find that both the beginning and the end of expansions are preceded by periods of rising macro-uncertainty. These facts are corroborated by evidence on business cycles and crises episodes documented elsewhere. With regard to asymmetric business cycles for example, Van Nieuwerburgh and Veldkamp (2006) and Ordonez (2013), and Jurado et al. (2013) document asymmetries in investment, lending rates, and real activity. For crises episodes, Gorton and Ordonez (2014b) show that credit booms that end in crises feature a larger initial productivity shock and a faster fall in total factor productivity over the boom. Further, excessive optimism seems to appear in many facets of the recent Financial Crisis. In their analysis of the 2006 housing bubble in the U.S., Cheng et al. (2014) find that misguided optimism was an important factor in the run-up to the crises over alternative explanations such as *bad incentives* or *bad luck*. Piazzesi and Schneider (2009) document an increase in a *momentum cluster* of households believing that house prices will continue to rise towards the end of the housing boom. Dell'Ariccia et al. (2012) link declining lending standards in the sub-prime market to the housing boom which led to an increase in delinquency rates. Similarly, Loutskina and Strahan (2011) show that the share of informed lending[21] in mortgage markets fell during the housing boom and such a fall in information production may have been an important factor in the crash that followed. Griffin and Tang (2012) look at credit ratings of collateralized debt obligations where they find subjective positive adjustments to model-implied ratings which predict future downgrades.

Macro-uncertainty is quite difficult to measure and various proxies are present in the literature.[22] Two of the most common are survey forecast disagreement and the option-implied expected volatility of the *S&P 500* index (VIX). In Figure 1.7 I plot the evolution of these two over the last four decades.

Though largely counter-cyclical, these two measures were rising well before the end of the last two expansions. The data appears to be noisier in the previous cycles and perhaps other factors played a larger role in generating the observed fluctuations. Nevertheless, one concern is that these are noisy measures of uncertainty and may not reflect the type of uncertainty consistent with the concept driving information production in the model. I introduce another measure of uncertainty which identifies forecast uncertainty at the individual level to augment these existing measures. In particular, I use the average dispersions of individual density forecasts (as against the dispersion of mean forecasts *across* forecasts) of annual average real GDP growth from the Survey of Professional Forecasters provided by the Federal Reserve Bank of Philadelphia from the first quarter of 1992 to the second quarter of 2013.[23] In a panel regression framework, I compute averages over time taking into account individual and forecast horizon differences. Further, I propose a simple way to correct for a potential bias in the measurement

[21]These are mortgage lenders that concentrate in a few markets and thus invest in more information about their borrowers.

[22]See Baker et al. (2013) for a measure of economic policy uncertainty based on news reports; Bloom (2009) and Bloom (2014) for uncertainty measures based on stock market volatility and micro-level dispersion and a comparison of various measures respectively; Jurado et al. (2013) for a factor-based approach using a host of macro time series; Orlik and Veldkamp (2014) focusing on uncertainty measures allowing for parameter uncertainty in forecast models; and Scotti (2013) for a higher frequency measure based on the size of surprises in real-time macro variables. More recently Rossi and Sekhoposyan (2015) constructs an uncertainty index based on the probability of observing a realized forecast error given the historical distribution of forecast errors.

[23]In the Appendix I also include estimates using forecast of the GDP price index growth. Other works derive a similar measure are Zarnowitz and Lambros (1987),Rich and Tracy (2010), and Bloom (2014) for the U.S. survey and Boero et al. (2008), Abel et al. (2015), and Boero et al. (2014) using the Bank of England and ECB surveys of professional forecasters. Of these, only Bloom (2014) ask how this measure of macro-uncertainty evolves over the business cycle.

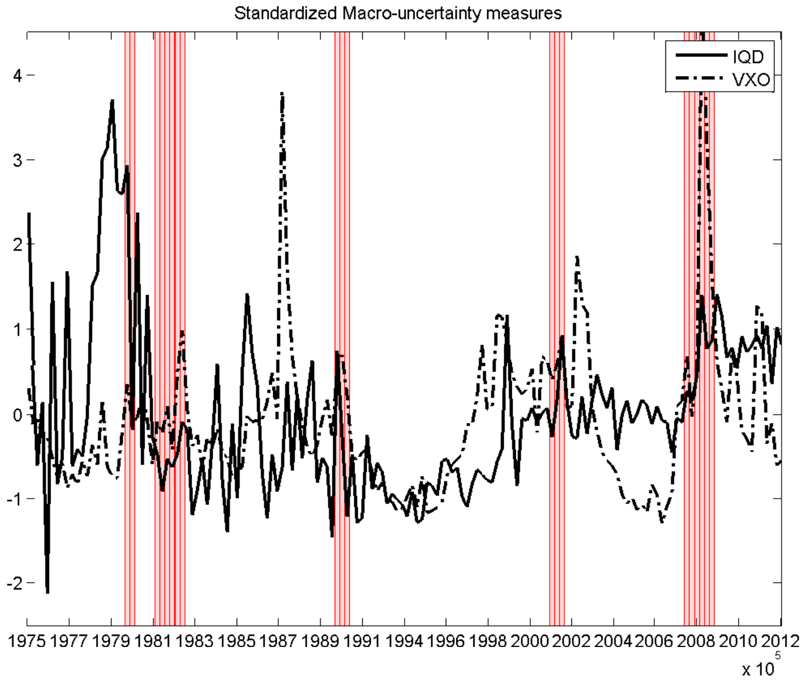

Figure 1.7: U.S. Macro-uncertainty proxies over time
IQD is the 75th less 25th percentile of four-quarter-ahead quarter-on-quarter real GDP growth forecasts from the Survey of Professional Forecasters. VXO is the quarter average of daily options-implied volatility of the S&P 500. Prior to 1986 I use realized volatility of the log difference of the S&P500 index. Between 1986 and up to 1989 I use the old definition of the VIX (VXO) and from 1990 onwards I use the current definition of the VIX. All measures have been standardized and a linear trend was taken out of IQD.

of uncertainty due to the truncation of survey responses.

These survey responses encode a coarse measure of the density of individual predictions about annual real GDP growth. Respondents are asked to provide probability values to specific ranges of outcomes for the target variable. I compute the entropy of a forecast $ENT_{i,t,h}$ as a measure of the diffusion of a forecast made by individual i at time t with a forecast horizon h.[24] A more dispersed distribution for the forecast yields a higher entropy which I interpret as forecast uncertainty. To obtain average forecast uncertainty over time, I run a regression with the following specification:

$$y_{i,t,h} = \alpha_i + \gamma_h + \delta_t + \epsilon_{i,t,h}$$

[24]See the Appendix for the data description. In contrast, Bloom (2014) compute the implied standard deviation of forecasts using the midpoints of bins and averages standard deviations per year.

The coefficients δ_t reflect the average individual forecast uncertainty for each quarter in the sample after controlling for individual and forecast-horizon average effects. I also estimate a specification controlling for a potential downward bias to the entropy measure when the survey responses place significant probabilities to the outer bins. In Figure 1.8 I plot the estimated time fixed effects for the entropy measure. My measure appears to co-move with the other two earlier shown and tend to be higher during recessions. Second, we also observe a U-shaped pattern over expansions. That is, macro-uncertainty begins to rise well before the end of an expansion from 1997 to 2001 and from 2006 to 2009.

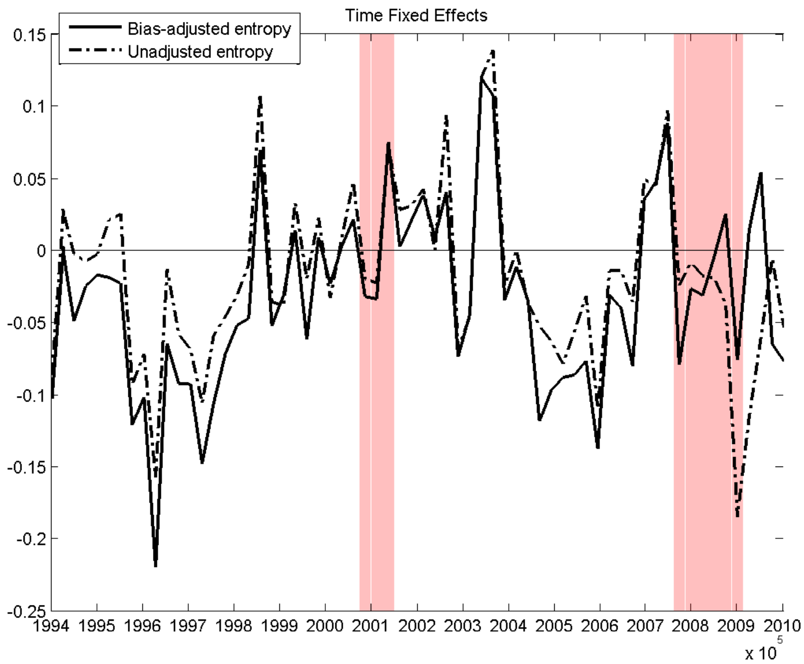

Figure 1.8: Average forecast uncertainty over time

The values reflect difference from the average level for the omitted time dummies for 1992. The adjusted series includes a bias-adjustment factor which depends on the median probability bin for each survey response. When the median probability bin is closer to the outer edges, the adjustment factor is larger.

My hypothesis for this pattern is that the flow of information is not simply pro-cyclical as the current literature would suggest. A fall in private information production at the heights of optimism in conjunction with pro-cyclical abundance of information may account for the U-shaped pattern of macro-uncertainty.

1.5 Conclusion

I have outlined a theory of learning and information production over the business cycle which regularly generates sharp recessions and gradual recoveries and produces the occasional boom that is sustained by optimistic beliefs even in the presence of warning signs that fundamentals have already began to deteriorate. In my theory, there is nothing irrational about these episodes. Instead, they arise because individuals have invested little in private information production. There is not much that can be learned from the actions of others in these booms and markets are content with being spared the details. These rational exuberance episodes endogenously arise in the model when beliefs become highly optimistic during peaks in economic activity. I have also documented suggestive evidence on measures of macro-uncertainty in the U.S. consistent with the proposed mechanisms.

My theory has several implications regarding the role of financial innovation and development on financial crises. Economies and periods in time with very high (or low) levels of financial development have a limited role for the intensive margin to information production and reduces the likelihood of rational exuberance episodes. On the other hand this also suggests that financial innovation, by generating complex and opaque assets, contributed to the fall in information production in the boom preceding the recent Financial Crisis. Furthermore, information is under-produced in the competitive equilibrium and one policy implication is that a subsidy on screening and lending may attain the constrained social planner's allocation. In the model, deviations of beliefs from the true state of fundamentals manifest in cyclical fluctuations of financier profits. This suggests that current proposals on macro-prudential policies that limit bank risk-taking, and hence induce private information production, and those that encourage credit and economic risk-taking during recoveries may generate additional information gains.

However, a more detailed and in-depth analysis of the policy implications of my learning mechanisms on macroeconomic fluctuations is limited by the simplistic nature of financial markets in the model. I have abstracted from several features of financial markets which we think are important drivers of crises and credit cycles. These include information asymmetries, the positive feedback between asset prices and credit, and leverage. This is an area for future research. More work also needs to be done in documenting and estimating private information production in financial markets. Finally, over-reliance on public information also generates correlated risk whose implications I do not explore in this paper. This is another area for future research.

Bibliography

Abel, J., Rich, R., Song, J., and Tracy, J. (2015). The measurement and behaviour of uncertainty: evidence from the ecb survey of professional forecasters. *Journal of Applied Econometrics*, forhcoming.

Baker, S., Bloom, N., and Davis, S. (2013). Measuring economic policy uncertainty. *Working Paper*.

Barsky, R. and Sims, E. (2011). News shocks and business cycles. *Journal of Monetary Economics*, 58 (3):273–289.

Beaudry, P. and Portier, F. (2006). Stock prices, news, and economic fluctuations. *American Economic Review*, 96 (4):1293–1307.

Berger, A. and Udell, G. (2004). The institutional memory hypothesis and the procyclicality of bank lending behavior. *Journal of Financial Intermediation*, 13:458–495.

Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992). A theory of fads, fashion, custom and cultural change as informational cascades. *Journal of Political Economy*, 100 (5):992–1026.

Blanchard, O., L'Huillier, J.-P., and Lorenzoni, G. (2013). News, noise, and fluctuations: An empirical exploration. *American Economic Review*, 103 (7):3045–3070.

Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica*, 77 (3):623–685.

Bloom, N. (2014). Fluctuations in uncertainty. *Journal of Economic Perspectives*, 28 (2):153–176.

Bloom, N., Floetotto, M., Jaimovich, N., Itay, S.-E., and Terry, S. (2012). Really uncertain business cycles. *NBER Working Paper*, 18245.

Boero, G., Smith, J., and Wallis, K. (2008). Uncertainty and disagreement in economic prediction: the bank of england survey of external forecasters. *The Economic Journal*, 118:1107–1127.

Boero, G., Smith, J., and Wallis, K. (2014). The measurement and characteristics of professional forecasters' uncertainty. *Journal of Applied Econometrics*, forthcoming.

Boz, E. (2009). Can miracles lead to crises? the role of optimism in emerging market crises. *Journal of Money, Credit and Banking*, 41 (6):1189–1215.

Boz, E. and Mendoza, E. (2014). Financial innovation, the discovery of risk, and the u.s. credit crisis. *Journal of Monetary Economics*, 62:1–22.

Broner, F. (2008). Discrete devaluations and multiple equilibria in a first generation model of currency crises. *Journal of Monetary Economics*, 55:592–605.

Bullard, J., Evans, G., and Honkapohja, S. (2010). A model of near-rational exuberance. *Macroeconomic Dynamics*, 14 (2):166–188.

Burguet, R. and Vives, X. (2000). Social learning and costly information acquisition. *Economic Theory*, 15 (1):185–205.

Caplin, A. and Leahy, J. (1994). Business as usual, market crashes, and wisdom after the fact. *American Economic Review*, 84 (3):548–565.

Chamley, C. (2004). Delays and equilibria with large and small information in social learning. *European Economic Review*, 48 (3):477–501.

Chamley, C. and Gale, D. (1994). Information revelation and strategic delay ina model of investment. *Econometrica*, 62 (5):1065–1085.

Chari, V. and Kehoe, P. (2003). Hot money. *Journal of Political Economy*, 111 (6):1262–1292.

Cheng, I.-H., Raina, S., and Xiong, W. (2014). Wall street and the housing bubble. *American Economic Review*, 104 (9):2797–2829. Preliminary. Presented in the Workshop in Behavioral Finance: Housing Bubbles, April 14, 2012.

Christiano, L., Motto, R., and Rostagno, M. (2014). Risk shocks. *American Economic Review*, 104 (1):27–65.

Collard, F., Dellas, H., and Smets, F. (2009). Imperfect information and the business cycle. *Journal of Monetary Economics*, 56:538–556.

Dell'Ariccia, G., Igan, D., and Laeven, L. (2012). Credit booms and lending standards: Evidence from the subprime mortgage market. *Journal of Money, Credit and Banking*, 44(2-3):367–384.

Dell'Ariccia, G. and Marquez, R. (2006). Lending booms and lending standards. *Journal of Finance*, 61 (5):2511–2546.

Eusepi, S. and Preston, B. (2011). Expectations, learning, and business cycle fluctuations. *American Economic Review*, 101 (6):2844–2872.

Fajgelbaum, P., Schaal, E., and Taschereau-Dumouchel, M. (2013). Uncertainty traps. *NBER Working Paper*, 19973.

Gennaioli, N., Shleifer, A., and Vishny, R. (2013). A model of shadow banking. *Journal of Finance*, 68 (4):1331–1363.

Gorton, G. and Ordonez, G. (2014a). Collateral crises. *American Economic Review*, 104 (2):343–378.

Gorton, G. and Ordonez, G. (2014b). Crises and productivity in good booms and bad booms. *Working Paper*.

Griffin, J. and Tang, D. Y. (2012). Did subjectivity play a role in cdo credit ratings? *Journal of Finance*, 67(4):1293–1328.

Jurado, K., Ludvigson, S., and Ng, S. (2013). Measuring uncertainty. *NBER Working Paper*, 19456.

Loutskina, E. and Strahan, P. (2011). Informed and uninformed investment in housing: the downside of diversification. *Review of Financial Studies*, 24 (5):1447–1480.

Martin, A. and Ventura, J. (2011). Theoretical notes on bubbles and the current crisis. *IMF Economic Review*, 59 (1).

Nimark, K. (2014). Man-bites-dog business cycles. *American Economic Review*, 104 (8):2320–2367.

Ordonez, G. (2013). The asymmetric effects of financial frictions. *Journal of Political Economy*, 121 (5):844–895.

Orlik, A. and Veldkamp, L. (2014). Understanding uncertainty shocks and the role of black swans. *NBER Working Paper*, 20445.

Patton, A. and Timmerman, A. (2010). Why do forecasters disagree? lessons from the term structure of cross-sectional dispersion. *Journal of Monetary Economics*, 57 (7):803–820.

Patton, A. and Timmerman, A. (2011). Predictability of output growth and inflation: A multi-horizon survey approach. *Journal of Business and Economic Statistics*, 29 (3):397–410.

Petriconi, S. (2012). Bank competition, information choice, and inefficient lending booms. *Job Market Paper*.

Piazzesi, M. and Schneider, M. (2009). Momentum traders in the housing market: Survey evidence and a search model. *American Economic Review: Papers & Proceedings*, 99 (2):406–411.

Rich, R. and Tracy, J. (2010). The relationships among expected inflation, disagreement, and uncertainty: Evidence from matched point and density forecasts. *The Review of Economics and Statistics*, 92 (1):200–207.

Rossi, B. and Sekhoposyan, T. (2015). Macroeconomic uncertainty indices based on nowcast and forecast error distributions. *American Economic Review Papers & Proceedings*, 105 (5):650–655.

Ruckes, M. (2004). Bank competition and credit standards. *The Review of Financial Studies*, 17 (4):1073–1102.

Scotti, C. (2013). Surprise and uncertainty indexes: Real-time aggregation of real-activity macro surprises. *FRB International Finance Discussion Paper*, 1093.

Van Nieuwerburgh, S. and Veldkamp, L. (2006). Learning asymmetries in real business cycles. *Journal of Monetary Economics*, 53 (4):753–772.

Veldkamp, L. (2005). Slow boom, sudden crash. *Journal of Economic Theory*, 124:230–257.

Zarnowitz, V. and Lambros, L. (1987). Consensus and uncertainty in economic prediction. *Journal of Political Economy*, 95 (3):591–621.

1.A Model Derivations

All aggregate variables which are subject to the learning mechanisms I propose have the time subscript $t|t-1$ denoting expectation of its value at time t given information up to time $t-1$.

1.A.1 Optimal contracts

Given a prior belief $p_{t|t-1}$ we then have the average expectation for borrower quality $\theta_{t|t-1}$ which we use to solve for credit market equilibrium where:

$$\theta_{t|t-1} = \theta_B + \mu_{t|t-1}(\theta_G - \theta_B)$$

First, we solve for loan rates. Given that there are no information asymmetries in the model, the probability of repayment is equal to the success probability of a project. If the borrower asks for the no-screening contract (I), then the participation constraint for each financier is given by:

$$\begin{aligned} \theta_{t|t-1} R_i(\theta_{t|t-1}, I) K_{i,\theta_{t|t-1}} &\geq (1+r_f) K_{i,\theta_{t|t-1}} \\ R_i(\theta_{t|t-1}, I) = R(\theta_{t|t-1}, I) &\geq \frac{1+r_f}{\theta_{t|t-1}} \end{aligned} \tag{1.10}$$

That is, without screening, the rate of interest must be such that the financier is at least as well of as investing the loaned funds into her savings technology. On the other hand, for the screening contract we have that the participation constraint for the financier must hold for both realizations of screening and thus the participation constraints of financiers are given by:

$$\begin{aligned} &\text{with Probability } \mu_{t|t-1}: \\ \theta_G R_i(\theta_G, S) K_{i,\theta_G} &\geq (1+r_f)(K_{i,\theta_G} + \gamma) \\ R_i(\theta_G, S) &\geq \frac{1}{\theta_G}[(1+r_f)(1+\frac{\gamma}{K_{i,\theta_G}})] \\ &\geq R(\theta_G, I) + \frac{(1+r_f)\gamma}{\theta_G K_{i,\theta_G}} \end{aligned} \tag{1.11}$$

with Probability $1 - \mu_{t|t-1}$:

$$R_i(\theta_B, S \quad \geq \quad R(\theta_B, I) + \frac{(1+r_f)\gamma}{\theta_B K_{i,\theta_B}} \tag{1.12}$$

Perfect competition among financiers ensure that these constraints bind and pins down interest rates.

1.A.2 Loan sizes and expected profits

With these interest rates, I solve for the optimal borrowing size. The entrepreneur will want to borrow the amount that maximizes his expected profits from running the project given by:

$$\mathbb{E}[\pi_i|\mu_{t|t-1}] \quad = \quad \theta_{t|t-1}(A_i K_i^\alpha - R(\theta_{t|t-1})K_i) \tag{1.13}$$

$$K_i^*(\mu_{t|t-1}) \quad = \quad \left[\frac{\alpha A_i}{R(\theta_{t|t-1})}\right]^{\frac{1}{1-\alpha}} \tag{1.14}$$

Given the optimal investment[25] size, expected profits from the no-screening contracts is given by:

$$\begin{aligned}\mathbb{E}[\pi_i|\mu_{t|t-1}, I] \quad &= \quad \left[\theta_{t|t-1}\frac{\alpha\, A_i}{(1+r_f)^\alpha}\right]^{\frac{1}{1-\alpha}} \left[\frac{1-\alpha}{\alpha}\right] \\ &= \quad (\theta_{t|t-1}A_i)^{\frac{1}{1-\alpha}}\Lambda \end{aligned} \tag{1.15}$$

where $\Lambda \equiv \left[\frac{\alpha}{(1+r_f)^\alpha}\right]^{\frac{1}{1-\alpha}} \left[\frac{1-\alpha}{\alpha}\right]$. Similarly, we can solve for expected profits from a screening contract:

$$\mathbb{E}[\pi_i|\mu_{t|t-1}, S] \quad = \quad (\mu_{t|t-1}\theta_G^{\frac{1}{1-\alpha}} + (1-\mu_{t|t-1})\theta_B^{\frac{1}{1-\alpha}})A_i^{\frac{1}{1-\alpha}}\Lambda - \gamma(1+r_f) \tag{1.16}$$

[25]In the case of the screening contract, expected profits are not always positive given the screening cost. In these cases, optimal loan sizes are zero.

1.A.3 Screening threshold

A given borrower with productivity A_i will choose the screening contract whenever expected screening profits are higher. Then screening is chosen when:

$$Z(\mu_{t|t-1}) \geq \left[\frac{\gamma(1+r_f)}{\Lambda}\right] A_i^{\frac{-1}{1-\alpha}} \tag{1.17}$$

where

$$Z(\mu_{t|t-1}) \equiv \mu_{t|t-1}\theta_G^{\frac{1}{1-\alpha}} + (1-\mu_{t|t-1})\theta_B^{\frac{1}{1-\alpha}} - (\mu_{t|t-1}\theta_G + (1-\mu_{t|t-1})\theta_B)^{\frac{1}{1-\alpha}}$$

The right-hand side of the above equation is decreasing in A_i and the left-hand side is concave in $\mu_{t|t-1}$ whose minimum is zero when $\mu_{t|t-1} \in \{0,1\}$. $Z(\mu_{t|t-1})$ captures the value of information to the borrower arising from prior beliefs about the state of the world.

Further, when γ is zero, the inequality always holds, and screening is always preferred. We can then define a cutoff level of productivity $\tilde{A}(\mu_{t|t-1})$ given expectations about the state of the world such that an entrepreneur with such a level of productivity is indifferent between either contract. For any A_i greater than this cutoff, the equation will hold with strict inequality and an entrepreneur would choose the screening contract. In equilibrium, financiers will also offer the screening contract to these agents and the no-screening contract to the rest. This guarantees that entrepreneurs who asked to be screened and turn out to be the Bad type have no incentive to go to a different financier and ask for the unscreened contract. For this reason, financiers offer a particular contract to each entrepreneur given their observed productivity.

$$\tilde{A}(\mu_{t|t-1}) = \left[\frac{\gamma(1+r_f)}{Z(\mu_{t|t-1})\Lambda}\right]^{1-\alpha} \tag{1.18}$$

1.A.4 Participation threshold

Given the optimal loan contract, an entrepreneur would like to borrow and invest whenever profits from this activity, exceed the outside option w. Since expected profits from production are increasing in A_i, define $A^*(\mu_{t|t-1})$ as the level of productivity such that an entrepreneur would be indifferent:

$$\begin{aligned} A^*(\mu_{t|t-1}) &= min\{A^*(\mu_{t|t-1}, S), A^*(\mu_{t|t-1}, I)\} \\ &\text{where:} \\ A^*(\mu_{t|t-1}, S) &= \left[\frac{w + \gamma(1+r_f)}{(\mu_{t|t-1}\theta_G^{\frac{1}{1-\alpha}} + (1-\mu_{t|t-1})\theta_B^{\frac{1}{1-\alpha}})\Lambda}\right]^{1-\alpha} \qquad (1.19) \\ A^*(\mu_{t|t-1}, I) &= \left[\frac{w}{(\mu_{t|t-1}\theta_G + (1-\mu_{t|t-1})\theta_B)^{\frac{1}{1-\alpha}}\Lambda}\right]^{1-\alpha} \qquad (1.20) \end{aligned}$$

The equations above simply equate expected profits from the screening and no-screening contract with the outside opportunity compensation w. Thus, a borrower finds it optimal to invest in her project $iff \quad A_i > A^*(\mu_{t|t-1})$.

Note that when $\gamma = 0$ then $A^*(\mu_{t|t-1}) = A^*(\mu_{t|t-1}, S)$ and the screening threshold is below the participation threshold whereas when $\gamma = \infty$ then $A^*(\mu_{t|t-1}) = A^*(\mu_{t|t-1}, I)$. In particular, this is the case whenever

$$\gamma(1+r_f) \geq \zeta(\mu_{t|t-1})\Lambda \qquad (1.21)$$

where $\zeta(\mu_{t|t-1}) \equiv \left[\frac{\mu_{t|t-1}\theta_G^{\frac{1}{1-\alpha}} + (1-\mu_{t|t-1})\theta_B^{\frac{1}{1-\alpha}}}{(\mu_{t|t-1}\theta_G + (1-\mu_{t|t-1})\theta_B)^{\frac{1}{1-\alpha}}}\right] \geq 1$.

1.A.5 Optimal updating

Beliefs about the current state today are formed given the full history of past information. This is a summarized in a recursive setup in which prior beliefs at any time t encapsulates all prior information and is updated with information at time t (Σ_t) to form a posterior belief given the new information. Denote the aggregate number of projects run and their composition into screened (S) and unscreened projects (I) as $n_t = n_t^S + n_t^I$. Define the number of successful projects in n_t^I as s_t^I and the number of screened projects (in n_t^S) with success probability θ_G as s_t^S. Define $\Sigma_t \equiv \{\frac{s_t^S}{n_t^S}; \frac{s_t^I}{n_t^I}\}$ as new information generated at time t. Thus, agents born at time t have the information set $\mathcal{I}_{t-1}$ which is recursively updated with outcomes (Σ_t) in each period $\mathcal{I}_t = \mathcal{I}_{t-1} \cup \Sigma_t$.

The proportion of successes among the number of projects run that were not screened is informative about the state of the world since they map into aggregate success rates.

$$\begin{aligned}\theta_H &= \theta_B + \mu_H(\theta_G - \theta_B)\\ \theta_L &= \theta_B + \mu_L(\theta_G - \theta_B)\end{aligned}$$

On the other hand, the proportion of screened projects with success probability equal to θ_G is also a direct statistic that may be used to infer μ_t.

Given some prior probability that $\mu_t = \mu_H$ denoted by $p_{t|t-1}$, we can use the set of signals from the unscreened outcomes to form posterior beliefs using Bayes' Law:

$$p^I_{t|t} = \frac{\theta_H^{s_t^I}(1-\theta_H)^{n_t^I - s_t^I} p_{t|t-1}}{\theta_H^{s_t^I}(1-\theta_H)^{n_t^I - s_t^I} p_{t|t-1} + \theta_L^{s_t^I}(1-\theta_L)^{n_t^I - s_t^I}(1-p_{t|t-1})}$$

The numerator is simply the likelihood of being in the high state given th signal $\frac{s_t^I}{n_t^I}$ weighted by the prior probability. The denominator normalizes the likelihood with the sum over the likelihood of both states. Similarly, we can use the second set of signals from the screened projects using the posterior from the first set of signals as a prior:

$$p_{t|t} = \frac{\mu_H^{s_t^S}(1-\mu_H)^{n_t^S - s_t^S} p^I_{t|t}}{\mu_H^{s_t^S}(1-\mu_H)^{n_t^S - s_t^S} p^I_{t|t} + \mu_L^{s_t^S}(1-\mu_L)^{n_t^S - s_t^S}(1-p^I_{t|t})}$$

Combining the two steps produce the following updating equation:

$$\begin{aligned}p_{t|t} &= \left[1 + \left(\frac{\mu_L(1-\mu_H)}{\mu_H(1-\mu_L)}\right)^{s_t^S}\left(\frac{(1-\mu_L)}{(1-\mu_H)}\right)^{n_t^S}\left(\frac{\theta_L(1-\theta_H)}{\theta_H(1-\theta_L)}\right)^{s_t^I}\left(\frac{(1-\theta_L)}{(1-\theta_H)}\right)^{n_t^I}\left(\frac{1-p_{t|t-1}}{p_{t|t-1}}\right)\right]^{-1}\\ &= \left[1 + LR^{-1}(\Theta, \Sigma_t)\frac{1-p_{t|t-1}}{p_{t|t-1}}\right]^{-1} \qquad (1.22)\end{aligned}$$

and

$$\mu_{t|t} = \mu_L + p_{t|t}(\mu_H - \mu_L) \qquad (1.23)$$

$LR(\Theta, \Sigma_t)$ is the likelihood ratio between the two states given parameters $\Theta = \{\mu_L, \mu_H, \theta_L, \theta_H\}$ and the signals $\{n_t^S, s_t^S, n_t^I, s_t^I\}$. Finally the optimal forecast for the aggregate state of the world is given by the posterior belief and the persistence parameter:

$$\begin{aligned} p_{t+1|t} &= \lambda\, p_{t|t} + (1-\lambda)(1-p_{t|t}) \\ &= (1-\lambda) + (2\lambda - 1)p_{t|t} \qquad (1.24) \\ \mu_{t+1|t} &= \mu_L + p_{t+1|t}(\mu_H - \mu_L) \qquad (1.25) \end{aligned}$$

Information production

The sensitivity of posterior beliefs to period t information depends on the informativeness of period t signals. These are characterized by their quantity and quality. As a measure for how much information is produced consider the likelihood of being in the high or low state given signals Σ_t:

$$\begin{aligned} \mathcal{L}(\mu_H|\Sigma_t) &= \mu_H^{s_t^S}(1-\mu_H)^{n_t^S - s_t^S}\theta_H^{s_t^I}(1-\theta_H)^{n_t^I - s_t^I} \\ \mathcal{L}(\mu_L|\Sigma_t) &= \mu_L^{s_t^S}(1-\mu_L)^{n_t^S - s_t^S}\theta_L^{s_t^I}(1-\theta_L)^{n_t^I - s_t^I} \end{aligned}$$

The above equations are the conditional likelihood of either states given period t information. The likelihood ratio $LR(\Theta, \Sigma_t)$ is just the ratio of these two equations. As a measure of information, the Kullback-Leibler divergence D_{KL} between these two likelihoods is[26]:

$$\begin{aligned} D_{KL}(\mu_H||\mu_L) &= s_t^S log(\frac{\mu_H}{\mu_L}) + (n_t^S - s_t^S)log(\frac{1-\mu_H}{1-\mu_L}) + \ldots \qquad (1.26) \\ & \quad s_t^I log(\frac{\theta_H}{\theta_L}) + (n_t^I - s_t^I)log(\frac{1-\theta_H}{1-\theta_L}) \qquad (1.27) \\ &= n_t\left[\frac{n_t^S}{n_t}\left(\mu_H log(\frac{\mu_H}{\mu_L}) + (1-\mu_H)log(\frac{1-\mu_H}{1-\mu_L})\right) + \ldots \right. \qquad (1.28) \\ & \quad \left. \frac{n_t^I}{n_t}\left(\theta_H log(\frac{\theta_H}{\theta_L}) + (1-\theta_H)log(\frac{1-\theta_H}{1-\theta_L})\right)\right] \qquad (1.29) \\ D_{KL}(\mu_L||\mu_H) &= n_t\left[\frac{n_t^S}{n_t}\left(\mu_L log(\frac{\mu_L}{\mu_H}) + (1-\mu_L)log(\frac{1-\mu_L}{1-\mu_H})\right) + \ldots \right. \qquad (1.30) \\ & \quad \left. \frac{n_t^I}{n_t}\left(\theta_L log(\frac{\theta_L}{\theta_H}) + (1-\theta_L)log(\frac{1-\theta_L}{1-\theta_H})\right)\right] \qquad (1.31) \end{aligned}$$

[26]The Kullback-Leibler divergence is a weighted average *distance* between two probability density functions evaluated at one of the measures taken to be the correct one. Note that this implies that this implies that this measure of divergence between two distributions is not symmetric.

The first equation is conditional on μ_H being the true state, and thus $\mathbb{E}[s_t^S] = \mu_H n_t^S$ and $\mathbb{E}[s_t^I] = \theta_H n_t^I$, whereas the second is with respect to the low state μ_L.

1.A.6 The extensive margin to information production

To construct the proof first we show that the quantity of signals is increasing in beliefs. Denote participation with the variable $\mathbb{1}_{w,i}$ if entrepreneur i decides to participate in credit markets and invest.

Lemma 1. *The number of signals $n_t \equiv \sum_{i \in N} \mathbb{1}_{w,i}$ is non-decreasing in beliefs about the state of the world $\mu_{t|t-1}$.*

Proof.

$$\begin{aligned}
\frac{n_t}{N} &= 1 - F(A^*(\mu_{t|t-1})) \\
& \quad \frac{\partial A^*(\mu_{t|t-1})}{\partial \mu_{t|t-1}} < 0 \\
\text{since} & \\
& \quad \frac{\partial A^*(\mu_{t|t-1}, S)}{\partial \mu_{t|t-1}} < 0 \\
& \quad \frac{\partial A^*(\mu_{t|t-1}, I)}{\partial \mu_{t|t-1}} < 0 \\
\text{Thus} & \\
\frac{\partial n_t}{\partial \mu_{t|t-1}} &\geq 0
\end{aligned}$$

The last equation holds with strict inequality whenever $F'(A^*) \neq 0$. $\square$

Proof of Proposition 1: The extensive margin to information production. The amount of information in Σ_t about the state of the world is characterized by $D_{KL}(\mu_L || \mu_H)$ and $D_{KL}(\mu_H || \mu_L)$. Assuming a constant proportion of screened projects, more optimistic beliefs yield a lower A^* and, by the previous lemma, increases n_t. Then, given the constant screening proportion:

$$\begin{aligned}
\frac{\partial D_{KL}(\mu_L||\mu_H)}{\partial n_t} &= \left[\frac{n_t^S}{n_t}\left(\mu_L log(\frac{\mu_L}{\mu_H}) + (1-\mu_L) log(\frac{1-\mu_L}{1-\mu_H})\right) + \dots\right. \\
&\quad \left.\frac{n_t^I}{n_t}\left(\theta_L log(\frac{\theta_L}{\theta_H}) + (1-\theta_L) log(\frac{1-\theta_L}{1-\theta_H})\right)\right] \\
&> 0 \\
\frac{\partial D_{KL}(\mu_H||\mu_L)}{\partial n_t} &= \left[\frac{n_t^S}{n_t}\left(\mu_H log(\frac{\mu_H}{\mu_L}) + (1-\mu_H) log(\frac{1-\mu_H}{1-\mu_L})\right) + \dots\right. \\
&\quad \left.\frac{n_t^I}{n_t}\left(\theta_H log(\frac{\theta_H}{\theta_L}) + (1-\theta_H) log(\frac{1-\theta_H}{1-\theta_L})\right)\right] \\
&> 0
\end{aligned}$$

The inequality arises since $\mu_H \neq \mu_L$ and $\theta_H \neq \theta_L$. If this were the case, then the two likelihoods are the same and the Kullback-Leibler divergence is always zero. Thus information production, or the information contained in period t signals, is non-decreasing in beliefs through the extensive margin. □

1.A.7 The intensive margin to information production

As before, we first show that the proportion of screened projects is weakly concave (inverse-U shaped) in beliefs.

Lemma 2. *The proportion of screening $n_t^S \equiv \sum_{i \in N} \mathbb{1}_{S,i}$ to total n_t signals in each period is weakly concave in beliefs $\mu_{t|t-1}$. In particular, it is inverse-U shaped whenever the screening cost γ is within a particular range.*

Proof. The measure of entrepreneurs who choose to be screened are given by:

$$n_t^S \propto 1 - \max\{F(A^*(\mu_{t|t-1})), F(\tilde{A}(\mu_{t|t-1}))\}$$

On the other hand, the measure of unscreened entrepreneurs is:

$$n_t^I \propto [F(\tilde{A}(\mu_{t|t-1})) - F(A^*(\mu_{t|t-1}))]^+$$

Then:

$$\frac{n_t^S}{n_t} \equiv \frac{n_t^S}{n_t^S + n_t^I} \quad \begin{cases} = 1 & \text{if } A^*(\mu_{t|t-1}) \geq \tilde{A}(\mu_{t|t-1}) \\ \propto \frac{1-F(\tilde{A}(\mu_{t|t-1}))}{1-F(A^*(\mu_{t|t-1}))} & \text{if } A^*(\mu_{t|t-1}) < \tilde{A}(\mu_{t|t-1}) \end{cases}$$

Thus, it is sufficient to show that the ratio $\frac{A^*(\mu_{t|t-1})}{\tilde{A}(\mu_{t|t-1})}$ is inverse-U shaped. We use equations 1.6 and 1.7 to get:

$$\begin{aligned} \frac{A^*(\mu_{t|t-1})}{\tilde{A}(\mu_{t|t-1})} &= \min\Bigg\{(1-\zeta(\mu_{t|t-1})^{-1})^{1-\alpha}(1+\frac{w}{\gamma(1+r_f)})^{1-\alpha}, \ \ldots \\ &\quad (\zeta(\mu_{t|t-1})-1)^{1-\alpha}(\frac{w}{\gamma(1+r_f)})^{1-\alpha}\Bigg\} \end{aligned}$$

where $\zeta(\mu_{t|t-1}) > 1$ and is concave in $\mu_{t|t-1} \in (0,1)$. In turn, both arguments in the minimization function are concave (inverse-U shaped) in $\mu_{t|t-1}$. □

Proof of Proposition 2: The intensive margin to information production. Keeping the total number of signals constant, and under the previous lemma, the change in informativeness of period t signals given an increase in the proportion of screened projects is given by:

$$\begin{aligned} \frac{\partial D_{KL}(\mu_L||\mu_H)}{\partial n_t^S} - \frac{\partial D_{KL}(\mu_L||\mu_H)}{\partial n_t^I} &= \left(\mu_L log(\frac{\mu_L}{\mu_H}) + (1-\mu_L)log(\frac{1-\mu_L}{1-\mu_H})\right) - \ldots \\ &\quad \left(\theta_L log(\frac{\theta_L}{\theta_H}) + (1-\theta_L)log(\frac{1-\theta_L}{1-\theta_H})\right) \\ &> 0 \\ \frac{\partial D_{KL}(\mu_H||\mu_L)}{\partial n_t^S} - \frac{\partial D_{KL}(\mu_H||\mu_L)}{\partial n_t^I} &= \left(\mu_H log(\frac{\mu_H}{\mu_L}) + (1-\mu_H)log(\frac{1-\mu_H}{1-\mu_L})\right) - \ldots \\ &\quad \left(\theta_H log(\frac{\theta_H}{\theta_L}) + (1-\theta_H)log(\frac{1-\theta_H}{1-\theta_L})\right) \\ &> 0 \end{aligned}$$

Both differences are positive in that $\theta_H - \theta_L = (\mu_H - \mu_L)(\theta_G - \theta_B) < \mu_H - \mu_L$ which guarantees that the first term in both equations are larger than the second. Intuitively, the larger difference in parameters of the binomial likelihoods given the screening signals relative to the no-screening signals makes the screening signals more informative as these signals would be sampled from relatively more different distributions. $\square$

The following corollary identifies special cases of the proposition.

Corollary. *When $\gamma \geq \bar{\gamma}$, no screening ever takes place and $n_t = n_t^I \quad \forall \mu_{t|t-1}$. When $\gamma \leq \underline{\gamma}$ then all credit market participants choose to be screened and $n_t = n_t^S \quad \forall \mu_{t|t-1}$.*

$$
\begin{aligned}
\bar{\gamma} &\equiv Z(\mu^*)\bar{A}^{\frac{1}{1-\alpha}} \frac{\Lambda}{1+r_f} \\
\underline{\gamma} &\equiv [\zeta(\mu^{**}) - 1]\left(\frac{w}{1+r_f}\right)
\end{aligned}
$$

The screening cost $\bar{\gamma}$ is the smallest value such that $min(\tilde{A}(\mu_{t|t-1})) = \bar{A}$. This is $\tilde{A}(\mu^*) = \bar{A}$ where μ^* is the level of beliefs which maximizes the gains from screening:

$$
\mu^* \equiv \frac{1}{\theta_G - \theta_B}\left[\frac{(1-\alpha)(\theta_G^{\frac{1}{1-\alpha}} - \theta_B^{\frac{1}{1-\alpha}})}{\theta_G - \theta_B}\right]^{\frac{1-\alpha}{\alpha}} - \theta_B
$$

On the other hand, $\underline{\gamma}$ is the largest screening cost such that $\tilde{A} \leq A^* \quad \forall \mu_{t|t-1} \in [\mu_L, \mu_H]$ where $\mu^{**} \in \{\mu_L, \mu_H\}$ is the level of beliefs which minimizes the gains from screening[27].

1.A.8 Persistence

Proof of Proposition 3: Dynamics of belief persistence. Define a change in beliefs given market outcomes as $\Delta p_{t|t} \equiv log(\frac{p_{t|t}}{p_{t|t-1}})$ and endogenous persistence as the degree by which this change in beliefs is small arising from the information content of date t signals. Then,

$$
\Delta p_{t|t} = -log(\big[p_{t|t-1} + LR^{-1}(\Theta, \Sigma_t)(1 - p_{t|t-1})\big])
$$

Clearly $\Delta p_{t|t} \neq 0$ when $LR^{-1}(\Theta, \Sigma_t) \neq 1$ and the extent by which posterior beliefs deviate from priors is proportional to the extent by which the likelihood ratio differs from one - the amount of information contained in signals at date t. To illustrate the two

[27]That is, μ^{**} minimizes $Z(\mu_{t|t-1})$ which will be at one of the boundaries, depending on parameter values.

channels clearly, we rewrite the likelihood ratio as follows:

$$LR^{-1}(\Theta,\Sigma_t) = \left(\left[\left(\frac{\mu_L(1-\mu_H)}{\mu_H(1-\mu_L)}\right)^{\frac{s_t^S}{n_t^S}}\left(\frac{1-\mu_L}{1-\mu_H}\right)\right]^{\frac{n_t^S}{n_t}}\left[\left(\frac{\theta_L(1-\theta_H)}{\theta_H(1-\theta_L)}\right)^{\frac{s_t^I}{n_t^I}}\left(\frac{1-\theta_L}{1-\theta_H}\right)\right]^{\frac{n_t^I}{n_t}}\right)^{n_t}$$

$\frac{s_t^S}{n_t^S}$ and $\frac{s_t^I}{n_t^I}$ are date t shocks, $\frac{n_t^S}{n_t}$ and $\frac{n_t^I}{n_t}$ are the relative shares of each type of signal, and n_t is the aggregate quantity of signals. The deviation of $LR^{-1}(\Theta,\Sigma_t)$ from one represents the total information content of all signals generated at time t. Thus, the object of interest is how priors affect the sensitivity of the likelihood ratio to date t shocks.

Suppose $\frac{s_t^S}{n_t^S}$ and $\frac{s_t^I}{n_t^I}$ are equal to ϵ. We evaluate the sensitivity of posteriors to a small change to such realizations $\Delta\epsilon$.

$$\begin{aligned}\frac{\partial\Delta p_{t|t}}{\partial\Delta\epsilon} &= \frac{\partial\Delta p_{t|t}}{\partial LR^{-1}(\Theta,\Sigma_t)}\frac{\partial LR^{-1}(\Theta,\Sigma_t)}{\partial\Delta\epsilon}\\ &= -\left[\frac{LR^{-1}(\Theta,\Sigma_t)}{\frac{p_{t|t-1}}{1-p_{t|t-1}}+LR^{-1}(\Theta,\Sigma_t)}\right]\left[n_t\left(\frac{n_t^S}{n_t}log(\frac{\mu_L(1-\mu_H)}{\mu_H(1-\mu_L)})+\frac{n_t^I}{n_t}log(\frac{\theta_L(1-\theta_H)}{\theta_H(1-\theta_L)})\right)\right]\end{aligned}$$

The first term reflects the sensitivity of posterior beliefs to current information while the second reflects the information content of shocks to date t signals. We are interested in the second component of the above equation[28].

First, the extensive margin reduces persistence in that sensitivity is increasing in n_t which in turn is non-decreasing in priors $\mu_{t|t-1}$ as shown in proposition 1. Second, sensitivity is increasing in $\frac{n_t^S}{n_t}$ given that signals from the screening outcomes are more precise: $log(\frac{\mu_L(1-\mu_H)}{\mu_H(1-\mu_L)}) < log(\frac{\theta_L(1-\theta_H)}{\theta_H(1-\theta_L)}) < 0$. Thus, the intensive margin generates an inverse-U shaped component to persistence as shown in proposition 2.

Depending on the relative values of the parameters $\{\mu_H,\mu_L\}$ and $\{\theta_H,\theta_L\}$, the contribution of the intensive margin to persistence in beliefs may either be amplified or attenuated. These, together with the sensitivity of the cutoffs A^* and $\tilde{A}$ to $\mu_{t|t-1}$ (also a

[28]The first component reflects persistence generated from how extreme priors are. Further, its functional form is sensitive to the assumption of a bivariate state. In this two-state economy, the ratio reflects both how extreme beliefs are (mean of the prior) and how precise these beliefs are thought to be (variance of the prior).

function of the other parameters r_f, α, w) determine whether there exists a threshold μ^* as defined in the proposition for a given level of γ. Finally case one in the proposition refer to the special cases such that $n_t^S = n_t$ or $n_t^I = n_t$ corresponding to threshold levels of the screening cost identified in the corollary to the previous proposition. In this case, the intensive margin does not apply and through the extensive margin, persistence of beliefs is monotonically decreasing in $\mu_{t|t-1}$ $\square$

1.B Socially Optimal Information Production

In this section I compare the competitive equilibrium allocation in Section 1.2 to that chosen by a constrained social planner. The production of private information generates spillovers to public signals (market outcomes) that become available to succeeding generations. In the economy with a hidden Markov chain for the state, efficiency is in part[29] constrained by how precisely market outcomes reveal the true state of the world. In the competitive equilibrium, agents take this precision as given not internalizing the fact that more participation and screening activity generates more precise prior beliefs for the next period. Instead, the cutoffs that arise in the competitive equilibrium are the ones that maximize only the current period's expected profits. A free-rider problem is present and, even in the presence of a market for information, inter-temporal trade is unlikely to resolve this inefficiency.

To facilitate our analysis, define aggregate expected profits in period t as given by the sum of all expected end-of-period returns to entrepreneurs:

$$\begin{aligned}\mathbb{E}[\pi_t|p_{t|t-1}] &= Nw + \sum_{A_i \in (A^*, \bar{A})} (\mathbb{E}[\pi_i|p_{t|t-1,I}] - w) f(A_i) + ... \\ & \sum_{A_i \in (max\{A^*, \tilde{A}\}, \bar{A})} (Z(p_{t|t-1}) A_i^{\frac{1}{1-\alpha}} \Lambda - \gamma(1+r_f)) f(A_i) \quad (1.32)\end{aligned}$$

Aggregate expected profits is the sum of the outside opportunity return for all entrepreneurs plus the gains from participation given by the no-screening contract for all entrepreneurs above the participation threshold and finally the gains from screening for all entrepreneurs above the screening threshold. This expression is increasing and convex in beliefs about the state ($p_{t|t-1}$) and is concave in the two thresholds A^* and $\tilde{A}$.

The cutoffs that maximize aggregate expected profits are given by the lowest values of A_i that give positive values for the second and third terms. This is exactly the conditions satisfied by the competitive equilibrium (equations 1.6 and 1.7). Thus, the competitive equilibrium yields the same allocation as a constrained social planner with this objective. This is what is done in Van Nieuwerburgh and Veldkamp (2006).

[29] Uncertainty and costly screening at the individual level also reduces the efficiency of the competitive allocation.

I consider an alternative objective for my social planner who internalizes the impact of these thresholds on the distribution of next period beliefs. Define a social welfare function as the discounted sum of all expected profits conditional on prior beliefs.

$$\max_{\{\tilde{A}_t, A_t^*\}_{t=1}^{\infty}} \mathcal{W} \;=\; \sum_{t=1}^{\infty} \beta^{t-1} \mathbb{E}[\pi_t | p_{t|t-1}]$$

The constrained social planner chooses a sequence of cutoffs that maximizes social welfare taking into account the intertemporal information spillover externality. I do not allow our social planner to deviate from the loan contracts specified in Section 1.2. Instead, the social planner only chooses the thresholds A_t^{*CSP} and $\tilde{A}_t^{CSP}$ which determine the set of participating and screened entrepreneurs. As in the competitive equilibrium, these cutoffs must be measurable with respect to period t information.

The constrained social planners problem yields the following Bellman equation and first order conditions

$$\begin{aligned}
V(p_{t|t-1}) &= \max_{\{\tilde{A}_t, A_t^*\}} \mathbb{E}[\pi_t | p_{t|t-1}] + \beta \mathbb{E}[V(p_{t+1|t}) | p_{t|t-1}] \\
\frac{\partial \mathbb{E}[\pi_t | p_{t|t-1}]}{\partial A_t^*} &= -\beta \frac{\partial \mathbb{E}[V(p_{t+1|t}) | p_{t|t-1}]}{\partial A_t^*} \\
\frac{\partial \mathbb{E}[\pi_t | p_{t|t-1}]}{\partial \tilde{A}_t} &= -\beta \frac{\partial \mathbb{E}[V(p_{t+1|t}) | p_{t|t-1}]}{\partial \tilde{A}_t}
\end{aligned}$$

Proposition 4 (Cutoffs under optimal learning)**.** *The solution to the constrained social planner's problem yields cutoffs $\tilde{A}^{CSP}(p_{t|t-1})$ and $A^{*CSP}(p_{t|t-1})$ which are lower than the cutoffs in the competitive equilibrium.*

The proof rests on two features of the model. First, the expected value function tomorrow is increasing and convex in the expected state of the economy $p_{t+1|t}$. As in the individual profit functions, the average of aggregate profits in the high and low states are higher than aggregate profits at the average state. Second, increasing the precision of information, by reducing either cutoff, induces a mean-preserving spread in the distribution of $p_{t+1|t}$. The next two lemmas derive these properties and the proof of the proposition follows.

Lemma 3. *The value function $V(p_{t+1|t})$ is increasing and convex in $p_{t+1|t}$*

Proof. Expected period $t+1$ profits $\mathbb{E}[\pi_{t+1}|t]$ is increasing and convex in $p_{t+1|t}$. This is straightforward from equation 1.32 which is just the aggregation of increasing and convex individual profit functions (equations 1.15 and 1.16). Further, the sequence of expected aggregate states of the economy $\{p_{t+1+s|t}\}_{s=1}^{\infty}$ are a linear and increasing function of $p_{t+1|t}$. Recall that

$$\begin{aligned} p_{t+s|t} - \frac{1}{2} &= (2\lambda - 1)(p_{t+s-1|t} - \frac{1}{2}) \\ &= (2\lambda - 1)^{s-1}(p_{t+1|t} - \frac{1}{2}) \end{aligned}$$

This then implies that $V(p_{t+1|t})$ is just the discounted sum of expected future profits which are increasing and convex in $p_{t+1|t}$ thus completing the proof.[30] □

Lemma 4. *A reduction in the cutoffs A_t^* and $\tilde{A}_t$ induce a mean-preserving spread in the distribution of $p_{t+1|t}$.*

Proof. This lemma is implied by proposition 3. For any given period t shocks, lowering the cutoffs A_t^* and $\tilde{A}_t$ increases the precision of information but preserves the mean as the ratios $\frac{s_t^S}{n_t^S}$ and $\frac{s_t^I}{n_t^I}$ are independent random variables which only depend on μ_t. On the other hand, as persistence falls with the precision of period t information, the expected variance of posterior beliefs increases.

To show that a reduction in the cutoffs preserve the mean forecast, note that

$$\begin{aligned} \mathbb{E}[p_{t+1|t}|p_{t|t-1}] &\equiv p_{t+1|t-1} \\ &= \mathbb{E}[(2\lambda - 1)(p_{t|t} - \frac{1}{2})|p_{t|t-1}] \\ &= (2\lambda - 1)(p_{t|t-1} - \frac{1}{2}) \end{aligned}$$

To show that the forecast variance increases I use the result from proposition 3 which shows a negative relationship between the participation and screening thresholds and

[30] Note that as $s \to \infty$ the sensitivity of expected $t+s$ profits to current priors goes to zero.

the resulting distance between the likelihood ratio of period t signals from one.

$$\begin{aligned}\mathbb{E}[(p_{t+1|t} - p_{t+1|t-1})^2 | p_{t|t-1}] &= (2\lambda - 1)^2 \mathbb{E}[(p_{t|t} - p_{t|t-1})^2 | p_{t|t-1}] \\ &= \left[(2\lambda - 1)(\frac{1}{4} - (p_{t|t-1})^2)\right]^2 \dots \\ &\quad \mathbb{E}[\left(\frac{LR(\Sigma_t) - 1}{1 + p_{t|t-1}(LR(\Sigma_t) - 1)}\right)^2 | p_{t|t-1}]\end{aligned}$$

The first term in the above equation captures the persistence of the state and thus the usefulness of past information. The second term captures the strength of prior beliefs (i.e., if priors are close to one or zero then posteriors are unlikely to move far away from priors). The last term reflects the variation in posteriors due to the precision of period t information as captured by the expected deviation of the likelihood ratio between the High and Low states ($LR(\Sigma_t)$) from one.

By proposition 3 we have that for any realization of the random variables s_t^S and s_t^I

$$\begin{aligned}\frac{\partial |LR(\Sigma_t) - 1|}{\partial A_t^*} &< 0 \\ \frac{\partial |LR(\Sigma_t) - 1|}{\partial \tilde{A}_t} &< 0 \\ &\text{then} \\ \frac{\partial \left(\frac{LR(\Sigma_t)-1}{1+p_{t|t-1}(LR(\Sigma_t)-1)}\right)^2}{\partial A_t^*} &< 0 \\ \frac{\partial \left(\frac{LR(\Sigma_t)-1}{1+p_{t|t-1}(LR(\Sigma_t)-1)}\right)^2}{\partial \tilde{A}_t} &< 0\end{aligned}$$

Thus a reduction in either cutoff will also yield an increase in the expected squared deviation of the likelihood ratio from one which implies that

$$\begin{aligned}\frac{\partial \mathbb{E}[(p_{t+1|t} - p_{t+1|t-1})^2 | p_{t|t-1}]}{\partial A_t^*} &< 0 \\ \frac{\partial \mathbb{E}[(p_{t+1|t} - p_{t+1|t-1})^2 | p_{t|t-1}]}{\partial \tilde{A}_t} &< 0\end{aligned}$$

Thus we have that a fall in either cutoff A_t^* or $\tilde{A}_t$ induces a mean-preserving spread on the distribution of posterior forecasts and next-period priors. □

We can now complete the proof of the proposition.

Proof of Proposition 4: Cutoffs under optimal learning. For the social planner cutoffs $\tilde{A}_t^{CSP}$ and A_t^{*CSP} to be lower than the cutoffs chosen in the competitive equilibrium, it is sufficient to show that:

$$\beta \frac{\partial \mathbb{E}[V(p_{t+1})|p_{t|t-1}]}{\partial A_t^*} < 0$$
$$\beta \frac{\partial \mathbb{E}[V(p_{t+1})|p_{t|t-1}]}{\partial \tilde{A}_t} < 0$$

Using the previous two lemmas, a reduction in the thresholds induce a mean-preserving spread in the distribution of next-period priors $p_{t+1|t}$. In turn, since the value function is increasing and convex in next-period priors, such a mean-preserving spread implies an increase in the expected value of the value function and completes the proof. □

Intuitively, increasing information production leads to a more dispersed distribution of next-period priors. Proposition 3 shows that by lowering the cutoffs, signals in period t are more precise and induce more variation in posterior beliefs towards the true state. This increases the expected future value function since it is increasing and convex over beliefs.

In the model, there are no avenues by which resources can be transferred through time and an inter-temporal market for information does not exist. Nevertheless, since information is a public good, even in the presence of such a market a free-rider problem is present and there are no individual incentives to participate in such an exchange. This creates space for a policy to improve on the competitive allocation by providing subsidies on information production and may be financed by a lump-sum tax on all agents. In particular, a policy which subsidizes investment ($\tau_P(p_{t|t-1})$) and screening ($\tau_S(p_{t|t-1})$) can implement the constrained social planner's solution. It should be noted that these subsidies are a function of prior beliefs insomuch as the social value of information also depend on the same.

1.C Macro Uncertainty

1.C.1 Data description

Uncertainty:

We are interested in measuring macro-uncertainty arising from information flows about aggregate fundamentals over time. Survey forecasts provide a rich dataset for this purpose in that we have available individual forecasts over probability ranges of the same target variable (e.g. annual average real GDP growth for the year 2009) for several forecast horizons. Given the probability range forecasts, we can compute the relative dispersion of each forecast with an entropy measure as detailed below. The variable is defined as the average entropy of individual probability range forecasts.

I make use of probability range forecasts of current and following year *annual average over annual average* change in real GDP. Survey respondents are asked to detail the probability that GDP growth will fall under a certain range. We use survey responses from the first quarter of 1992 to the fourth quarter of 2012 where the GDP growth ranges are binned into < -2,-2 to -1.1, and so on until > 6 percent for a total of ten bins of one percent width each except for the extremes.[31]. Sample responses are provided in the table below:

Table 1.2: Sample Survey Responses

year	quarter	id	industry	Cur > 6	Cur 5 to 5.9	Cur 4 to 4.9	Cur 3 to 3.9	Cur 2 to 2.9	Cur 1 to 1.9	Cur 0 to 0.9	Cur -1 to -0.1	Cur -2 to -1.1	Cur < -2
1992	1	30	3	0	0	10	60	30	0	0	0	0	0
1992	1	35	2	0	0	10	10	20	50	10	0	0	0
year	quarter	id	industry	Fol > 6	Fol 5 to 5.9	Fol 4 to 4.9	Fol 3 to 3.9	Fol 2 to 2.9	Fol 1 to 1.9	Fol 0 to 0.9	Fol -1 to -0.1	Fol -2 to -1.1	Fol < -2
1992	1	30	3	0	0	30	50	20	0	0	0	0	0
1992	1	35	2	0	10	20	20	40	10	0	0	0	0

Table 1.2 reproduces the responses of two participants (identified as 30 and 35). The first two columns report the date of the survey and the fourth column reflects the industry to which the respondent belongs two (financial, non-financial or unknown). Categories *Cur* > 6 to *Cur* < -2 reflect the 10 annual average growth rate bins for real GDP growth

[31] From the second quarter of 2009 onwards, the lowest bin has been split into < -3 and -3 to -2.1. I have aggregated the responses for these two bins. Further, I only use data beginning on the first quarter of 1992 as the survey question from the third quarter of 1981 to the fourth quarter of 1992 corresponds to forecasts of real GNP and only has 6 bins. Prior to this period the survey asks for forecasts of nominal GNP.

for the current year 1992. Given that these responses were obtained in the first quarter of 1992, we have coded these as 4-quarter ahead forecasts. The categories *Fol* > 6 to *Fol* < -2 are the same growth rate bins for the following year (1993) and are coded as 8-quarter ahead forecasts. The values in these categories reflect the probabilities (in percent) that the respondents attach to real GDP growth rate falling within these bins. Our variable of interest is the diffusion of these forecasts. In the table above, forecaster 35 has a more diffuse forecast than forecaster 30 for the current and following year real GDP growth. We may interpret this as higher forecast uncertainty. I do a similar exercise using probability range forecasts for the GDP price index from the same survey (annual average over annual average percent change in the GDP price index) with similar results and are omitted from this paper.

I compute an entropy measure for each individual respondent i in date t defined as[32]:

$$ENT_{i,t,h} \equiv -\sum_{b=1}^{10} p_{i,t,h,b} log(p_{i,t,h,b})$$

where $p_{t,i,b}$ is the probability value for the bin b given by respondent i in date t for real GDP or the GDP Price index growth in year y where each annual average growth rate is forecasted for eight periods (for each period h may refer to either the current year or the following year and t is quarterly). For the sample period (86 quarters), each forecaster makes two predictions per quarter (current and following year) where there are 162 unique forecasters for the real GDP growth series[33]. Within the sample, there is a minimum of 23 (fourth quarter of 2001) and a maximum of 50 respondents (third quarter of 2005) for any given quarter.

This measure is bounded between zero and $log(10)$ and reflects the dispersion of a respondent's forecast. The forecast horizon ranges from 8 quarters ahead (for a forecast of the following year's annual average GDP growth made in the first quarter of the current year) to 1 quarter ahead (a forecast made for the current year annual average GDP growth in the fourth quarter of the current year)[34]. In the following table, I report

[32]Rich and Tracy (2010) conduct a similar exercise. We have also computed a sample standard deviation of the forecast using midpoints of the bins and obtained similar time series.

[33]Note that the FRB Philadelphia's caveat on the identifiers apply. That is sometimes the identifier is associated with a firm rather than an individual.

[34]That is, for the first quarter of each year we have two average entropy measures corresponding to the forecast of the current year with horizon 4 and the forecast of the following year with horizon 8. In the

some summary statistics.

Table 1.3: Entropy measure statistics

	Individuals	Non-zero lowest bin	Non-zero highest bin	Mean	Median	Standard deviation	Min	Max	Obs
GDP Growth	162	326	188	1.0821	1.0627	0.4571	0.0000	2.2804	6120

I observe a total of 6,120 forecasts and about eight percent of these have probabilities on the extreme bins. The following figure is a histogram of the entropy measure for real GDP growth. Note that the entropy measures are biased downwards when the location

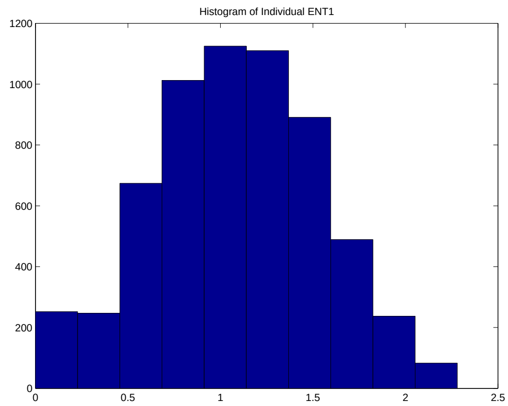

Figure 1.9: Entropy measure histogram

of the reported probability distributions are close to either of the extreme bins as all the probability weights for growth less than two percent (or greater than 6 percent) are lumped into one bin. As an indicator for the severity of this bias, I report the average probability value assigned to the lowest and the highest bins over time in the following figure. The figures suggest that our entropy measure is likely to be biased downwards in 2009 for real GDP growth entropy with the mean probability value assigned to the lowest bin reaching more than thirty percent. To estimate the potential size of this bias in 2009 I conduct the following exercise for real GDP growth forecasts. First I use the responses in the fourth quarter of 2007 to generate a representative distribution. The fourth quarter of 2007 was chosen since it is relatively close to the second quarter of 2009 and the average responses reflect a median probability in the bin two to three

second quarter of each year we would have the same forecast targets but with corresponding horizons of 7 and 3. Similarly, the horizons for measures in the third quarter are 6 and 2 and in the fourth quarter they are 5 and 1

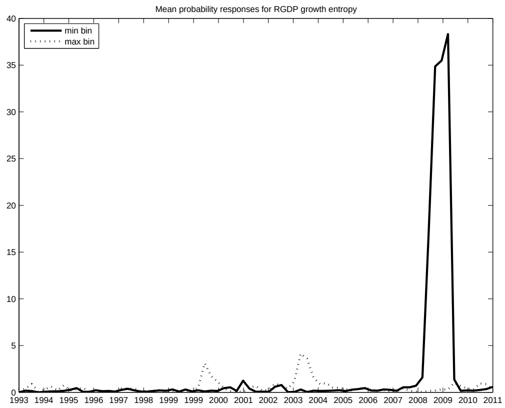

Figure 1.10: Average responses assigned to lowest and highest bins over time

percent which is in the middle of the bins. The following figure plots the averaged responses for the fourth quarter of 2007 and the second quarter of 2009. Note that by averaging responses, the measured entropy from this average distribution will also reflect the disagreement across forecasters[35]. The right panel in Figure 1.11 is quite informative over the source of disagreement and lower uncertainty for the second quarter of 2009. We have that 21 of the 90 current and following year forecasts in this period put all probability mass at the bottom two bins (less than one percent growth) generating a second mode in the average distribution and low individual entropy measures. The rest of the respondents were more optimistic and were forecasting growth to be in the one to three percent bin on average. In estimating average forecast uncertainty over time, I introduce a specification that controls for this potential downward bias at the individual level.

1.C.2 Entropy measure regression

Since this measure may be subject to horizon effects where uncertainty is expected to be increasing in the forecast horizon (see Patton and Timmerman, 2010 for horizon effects on forecast dispersion and Patton and Timmerman, 2011 for mean squared errors), I estimate a panel regression with both individual and forecast horizon fixed effects. I

[35] Suppose for instance that two forecasters put 100 percent probability on the bins less than two percent and greater than 6 percent respectively. Then the average individual entropy is zero whereas the average distribution will have fifty percent mass at both extremes and the entropy of the average distribution is non-zero.

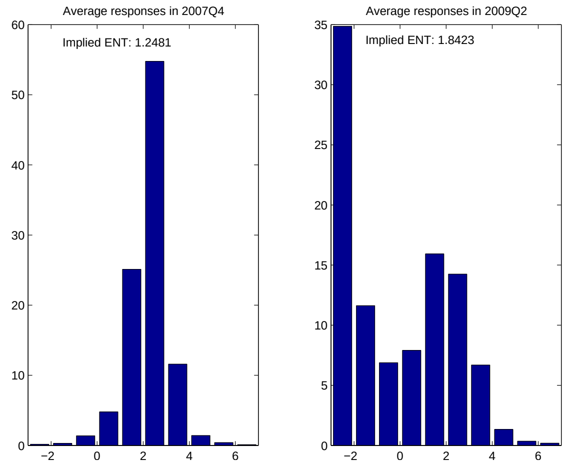

Figure 1.11: Average responses in 2007Q4 and 2009Q2

include time fixed effects to capture the average value of entropy for each period with the following specification:

$$ENT_{i,t}(h) = \alpha_i + \gamma_h + \delta_t + \epsilon_{i,t,h}$$

The time-series of δ_t are plotted in the following figure for each entropy measure along with confidence intervals. These are interpreted as the average variation in time in entropy net of individual fixed effects at an eight quarter forecast horizon and relative to the average for the year 1992. The figure shows three episodes of heightened uncer-

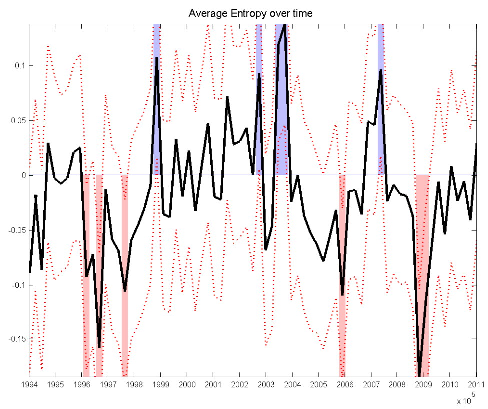

Figure 1.12: Average entropy measures over time

tainty relative to 1992 (omitted time fixed effect coefficients) and four episodes of lower uncertainty for real GDP growth forecasts. The periods of heightened uncertainty are in 1998, 2002-03, and 2007 whereas uncertainty was relatively low in 1993, 1996-1997, 2005, and 2009. Note that we have earlier identified a potential downward bias in entropy for real GDP growth forecasts in 2009. However, I also find a similar patter across these episodes for the entropy measure based on the GDP price index growth forecasts where such a bias does not seem to be severe.

I then derive a bias-adjusted measure of individual entropy for real GDP growth forecasts. First, I estimate the bias by using the average distribution of probability responses when the median bin is at two to three percent. I then move this sample distribution to the left and right of the responses effectively truncating the probability values as the median bin moves closer to either extreme. The next figure plots the transformed average probability responses over the full sample over various median bins. The leftmost panel in the second row is the untransformed average response with the median bin at two to three percent. The next table reports the bias-adjustment factor equivalent to the difference in measured entropy using the transformed responses relative to the original response.

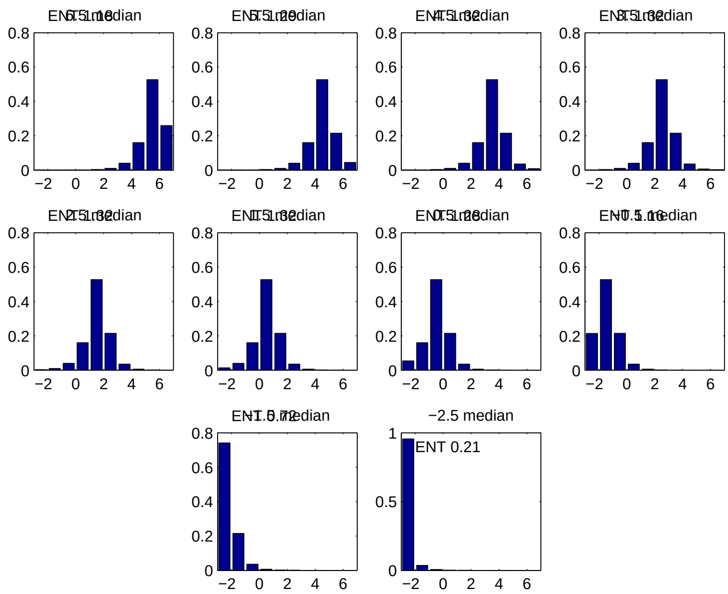

Figure 1.13: Hypothetical responses over various median bins

Table 1.4: Relative reduction in measured entropy

Median Bin	< -2	-2 to -1.1	-1 to -0.1	0 to 0.9	1 to 1.9	2 to 2.9	3 to 3.9	4 to 4.9	5 to 5.9	> 6
Relative Entropy	-1.113	-0.604	-0.159	-0.038	-0.008	0.000	-0.002	-0.007	-0.028	-0.1467

The differences are reported in Table 1.4 where the columns report the difference in entropy corresponding to the median bin of each individual response (e.g. a respondent whose median probability value is at the < -2 percent bin has a lower entropy measure by 1.113). I then run an augmented regression including this bias factor.

$$ENT_{i,t}(h) = \alpha_i + \gamma_h + \delta_t + \beta\, bf_{i,t} + \epsilon_{i,t,h}$$

where $bf_{i,t}$ is the bias factor given by Table 1.4 and the median bin of the forecast for each respondent. The following figure reports the time fixed effects of these regression. The figure suggests that the fall in uncertainty in 2009 is not robust to adjustment for

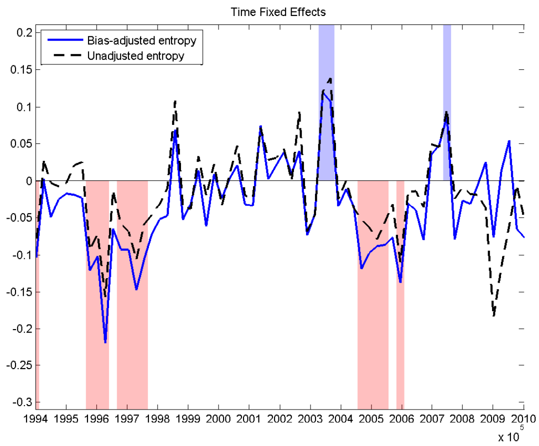

Figure 1.14: Time fixed effects with and without bias-adjustment

potential bias due to truncation of responses.

Does uncertainty precede or proceed from downturns? Up until 2006 the general pattern is that our measure of uncertainty tends to be low during expansions and peaks during or after recessions. On the other hand, our measure of uncertainty also peaked in the fourth quarter of 2007. We compare uncertainty with the median growth forecast (the median probability range across individuals for current and following year forecasts containing the 50th percentile) in the left panel of the next figure. The right panel is a scatter plot of average entropy in the y-axis and the median forecast probability bin for each period in our sample. It seems that the increase in uncertainty prior to the 2001

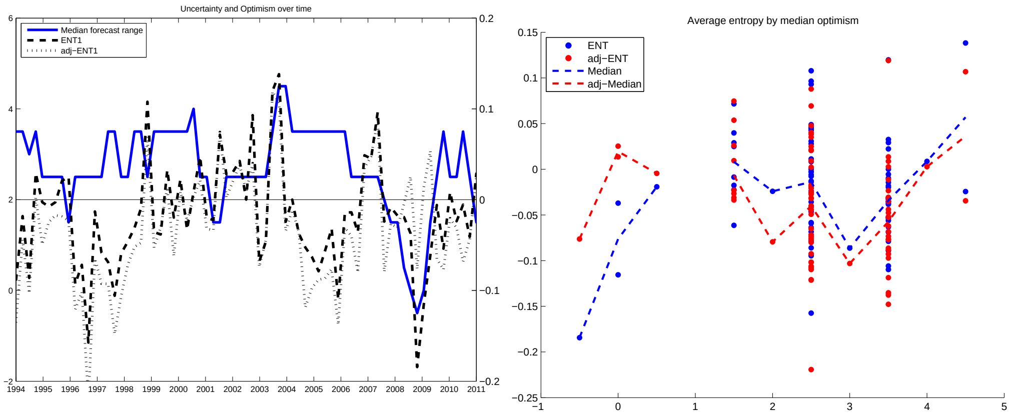

(a) Entropy and Optimism over time

(b) Entropy vs Optimism

Figure 1.15: Entropy and Optimism

recession coincided with a drop in the median probability forecast range, a decline in optimism. However, the short-lived recession exhibited a quick rebound in optimism while uncertainty did not start to fall until the end of 2003. For the 2008-2009 recession however, uncertainty started falling by the first quarter of 2008 and bottoms out at the NBER-dated trough.

This pattern appears to be unique to my measure. In the next figure I compare the entropy measure ($ENT1$) with other common measures of uncertainty, the VIX and the inter-quartile range of quarterly real GDP growth point forecasts (four-quarter ahead) from the SPF. All measures have been standardized for easy comparison. My measure appears to co-move with the other two up until 2006. Stark differences appear around the recent Financial Crisis which perhaps points to another feature that differentiates this episode from ordinary recessions. My measure began increasing since the third quarter of 2006 and peaked in the fourth quarter of 2007. The VIX and inter-quartile forecast dispersion began to increase only around the second quarter of 2007 and peaked in the fourth quarter of 2008 and the third quarter of 2009 respectively. One interpretation is perhaps that as housing prices peaked by early 2006 and the subprime crisis was developing, uncertainty about macroeconomic prospects began to rise but there seems to be a consensus about the direness of the situation. This interpretation of these differences

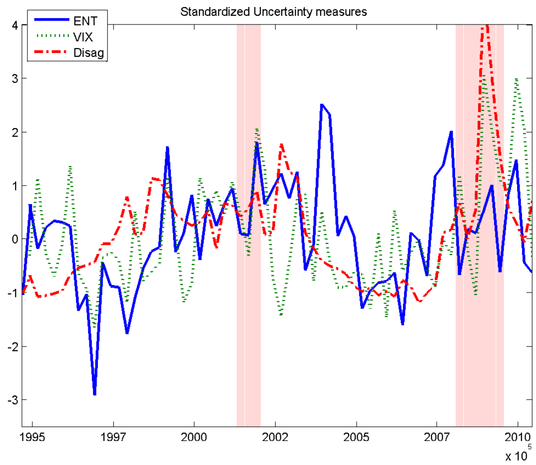

Figure 1.16: Entropy, VIX, and Disagreement

is that the relative magnitudes of private and public information used in these forecasts can explain the coincidence of low forecast disagreement and high uncertainty prior to the crisis and the opposite pattern in the recession that followed.

I also confirm that uncertainty appears to be reduced as the forecast horizon shrinks[36]. Figure 1.17 report the estimated horizon coefficients for real GDP growth forecasts. In the following figures, the estimated horizon coefficients indicate a marked decline in entropy as the forecast horizon shortens.

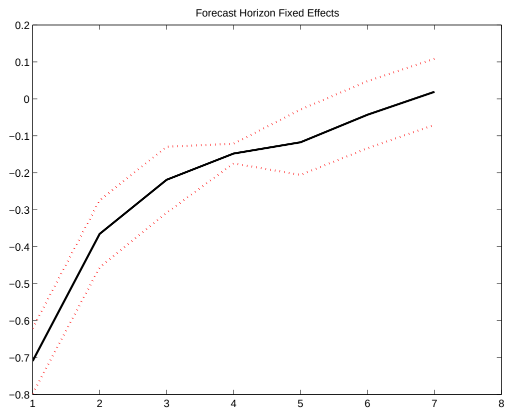

Figure 1.17: Horizon Fixed effects

[36]Patton and Timmerman (2010) and Patton and Timmerman (2011) document these features for the cross-sectional dispersion of forecasts and the mean squared forecast error.

Finally, I also constructed a series of sample standard deviation of probability range forecasts using the midpoints of each bin (e.g. a 40 percent probability assignment to the range 1 to 1.9 percent is coded as 40 percent probability to 1.5 percent) as an alternative measure of forecast dispersion with similar results.

Chapter 2

PRIVATE INFORMATION PRODUCTION IN CREDIT MARKETS

2.1 Introduction

A key ingredient to banking and financial crises is the systemic deterioration of bank loan portfolios that largely go unnoticed in the preceding credit boom. One of the proposed hypothesis behind this deterioration is a fall in borrower screening intensity during the boom.[1] This may occur because the incentives and benefits to private information production, or screening of borrowers, fall when the average borrower quality is either 'too low' or 'too high'.[2] Pro-cyclical competition may also produce the same effect.[3] These may further be exacerbated by the counter-cyclical quality of credit ratings.[4] These hypotheses raise the following empirical questions. How much private information about their borrowers do financial intermediaries produce? To what extent are these reflected in the terms of credit? How has this evolved over time and over the business cycle?

[1]For the purposes of this chapter, lending standards refer to loan contract decisions based on publicly observable or costlessly verifiable information whereas screening refers to the production of private information about borrower fundamentals which may also be costly to communicate (e.g. 'soft' information).

[2]See for example Ruckes (2004); Dell'Ariccia and Marquez (2006), and Gorton and Ordonez (2014).

[3]Examples are Hauswald and Marquez (2006); Gorton and He (2008), and Petriconi (2012).

[4]See Bolton et al. (2012) and Bar-Isaac and Shapiro (2013).

To date the cyclical variation (or lack thereof) in credit market screening intensity has yet to be empirically demonstrated and understandably so. Private information production is an elusive concept and is quite difficult to measure.[5] This chapter fills in this gap. I merge information spanning over two decades and across four datasets to construct a measure of private information production in the term loan credit market in the United States. The measure is constructed by decomposing the variation in term loan spreads (the average loan rate less LIBOR) into that which can be explained by observable borrower, lender, and macroeconomic characteristics and a residual variation which I interpret as the private information that lenders possess about the ability of the borrowers to repay.

In particular, I use data on over eleven thousand term loans between 1987 and 2011 in the *Dealscan* dataset from the Loan Pricing Corporation. I match this information to borrower balance sheet and stock market performance characteristics from the *Compustat* and *CRSP* datasets. Finally, I identify and match the most frequently occurring lenders in the loan dataset to the the Federal Reserve's Bank Consolidated Holdings database to obtain lender-specific information. These combined datasets, along with time fixed effects to control for macroeconomic and financial conditions, allow me to decompose the variation in the terms of credit for these loans into that which can be explained by publicly available information and a residual variance - my measure for private information. I then run a second set of regressions to identify the time-variation in my private information measure controlling for other factors. I also check whether the residuals I estimated in the first exercise indeed constitutes private information. Using an event study approach, I evaluate whether the residual component of the term loan spread can explain some of the variation in the stock market performance of borrowers' equity in the days following the loan agreement.

The main results of the chapter are the following. First, I find that my measure of private information production has a hump-shaped pattern over the last two expansions in the U.S. business cycle. Consistent with the theoretical predictions in the literature,

[5]On the other hand there is survey-based evidence of counter-cyclical variation in lending standards. Note that this does not directly imply the same for screening intensity which results in private information production. For example, in Ruckes (2004) and in the first chapter of this dissertation, both lending standards and screening which results in private information production are used in financial markets. In both, lending standards are counter-cyclical whereas screening has a hump-shaped pattern over the business cycle.

screening intensity appears to initially increase at the beginning of expansions and peak towards the middle before declining some years prior to the onset of the next recession. They are also lower for securitized loans, loans of longer maturities, and loans facilitated by a larger number of lenders. Second, in the days following the loan agreements, I find some evidence of better stock market performance for the borrowers who obtained interest rates that were lower than that predicted by the observable factors. This would suggest that the unpredicted component of loan interest rates estimated in the previous exercise contain new, and private, information about the borrowers that then becomes publicly available and priced into the borrowers' equity in the days following the loan agreement.

My sample consists of term loans taken out by large borrowers in the U.S. and are mostly syndicated loans.[6] Prior evidence would suggest that these are typically the types of borrowers for which private information or discretion has a small role to play (Cerqueiro et al., 2011). On the other hand, as argued in Sufi (2007) in his analysis of syndication structure and the role of lead lenders in resolving information asymmetry, lead lenders play an important role as *informed capital* (a la Holmstrom and Tirole, 1997). That is, private information is also an important factor in this segment of the credit market. Further, aside from the availability of data spanning decades, the private information that is revealed by these term loans are also likely to be more relevant to the macro-economy as a whole.

The apparent cyclical variation in screening intensity provide some evidence to the largely theoretical literature on credit cycles. Further, the results suggest a fall in screening intensity that leads the business cycle (i.e. the onset of a recession). This is consistent with the view that financial fragility and the build-up of systemic risk, through bank risk-taking behavior, varies over the business cycle and may be excessive during booms. As such, the evidence presented in this chapter provide some support for macro-prudential policies that provides dis-incentives to risk-taking by banks and financial intermediaries in a cyclical manner.

Related Literature. I am not the first to employ this strategy to estimate private information production in credit markets. Cerqueiro et al. (2011) use survey data from

[6] An earlier study by Strahan (1999) uses a similar dataset and interested readers may refer to it for a discussion on the average type of firm in the sample. My regression specifications are also patterned after his although I include other variables previously not available.

the 1993, 1998, and 2003 National Survey of Small Business Finances by the Federal Reserve Board in the U.S. They examine the determinants of the variation in loan rates for small businesses as decomposed into rules (variation that can be explained by loan, firm, and relationship characteristics) and discretion (the residual variance). They find that discretion is more important for small and unsecured loans, small and opaque borrowers, and more importantly, a downward trend over the three survey periods.[7] The main contribution of this chapter is that I use term loan data spanning over two decades in the U.S. which allows me to evaluate how private information production has evolved over time and in relation to the business cycle.

There is also some evidence on counter-cyclical lending standards as can be seen in the Loan Officer Survey of the Federal Reserve Survey or the recent work by Dell'Ariccia et al. (2012) but these do not directly imply the same for private information production.[8] Largely due to data constraints, contributions to the literature which try to measure the extent to which private (or soft) information is produced by banks typically focus on the cross-section rather than the time dimension. For example, Agarwal and Hauswald (2010) use loan data over a 16-month period starting in January of 2002 from a large U.S. bank to examine the relationship between distance and soft information production. They regress an internal credit score rating by the bank on publicly available estimates of credit worthiness (credit bureaus) and use the residual variation as a measure of private information. Degryse and Ongena (2005) conduct a similar exercise although they focus on the effect of distance on price discrimination using loan data from a large Belgian bank initiated since 1995 and still on the bank's loan portfolio by August 1997. Although the authors do not highlight the result, in a regression of loan rates on loan, borrower, and relationship characteristics as well as distance, they find that the unexplained variation in loan rates is decreasing in the size of the loan, whether the loan is collateralized, and larger for sole proprietorships. More recently, Cerqueiro et al. (2014) use loan data from a large bank in Sweden around a change in Swedish laws that reduced the value of collateral for some loans to determine the effects of collateralization on loan rates and monitoring. They use loan data of about one year before

[7]A similar strategy has also been employed in the evaluation of qualitative or private information in credit ratings. See Griffin and Tang (2012) and Norden and Roscovan (2013).

[8]Senior loan officer surveys conducted by the Federal Reserve suggests counter-cyclical lending standards. The data is available at `http://www.federalreserve.gov/boarddocs/snloansurvey`

and after the change in law took effect (January 2004) and use a difference in difference strategy to establish that a fall in collateral value leads to an increase in the loan rate, tightened credit limits, and an overall reduction in the intensity of monitoring (defined as the frequency of borrower or collateral reviews made by the bank).

The second exercise I conduct on the effect of residual loan spreads on borrower stock returns follows the literature on the impact of loan announcements.[9] The seminal paper by James (1987) associated bank loan announcements with almost two percent increase in the cumulative abnormal return of a borrower's stock price. Other works have also associated the magnitude of the effect as depending on firm size (Slovin et al., 1992), loan and lender characteristics (e.g. loan initiations or renewals in Lummer and McConnell 1989[10], lender's credit rating in Billet et al. 1995, location relative to borrower in Ongena and Roscovan 2013 or portion of loan retained by arranger in Focarelli et al. 2008). Ongena et al. (2014) show that bank loan announcements elicit responses in both bond and stock markets. Of special interest are results from Best and Zhang (1993) which indicate that banks produce and use more information when other sources of information (analyst forecasts) are more noisy. More recent work by Li and Ongena (2014) estimate stock market response around the 2007-08 Financial crises and their findings suggest that the magnitude of stock market reactions is different across the credit cycle, lower or near-zero during booms and much larger after crashes. In my setting, I use these factors as controls and include the residual variation in term loan spreads as an additional explanatory variable.

The rest of the chapter is organized as follows. The next Section provides the Empirical framework while Section 2.3 provides a brief description of the data. The estimation results on private information production over the business cycle are documented in Section 2.4. I then proceed with analysis of the impact of the estimated private information component to the loan contract on borrower stock returns in Section 2.5. Finally, Section 2.6 concludes.

[9] I refer mainly to works using US data where the time period ranges from the 1980s to the recent Financial Crisis.

[10] In contrast with Lummer and McConnell (1989) who find no statistically significant effect on loan initiations, Billet et al. (1995) still find a statistically significant albeit lower effect whereas Best and Zhang (1993) qualify that initiations do not elicit abnormal returns conditional on financial analyst prediction errors (outside information) being low.

2.2 Empirical Framework

To obtain a measure for private information production in credit markets, I assume the following to be true for the setting of term loan contracts. Let the log of the loan rate spread to a borrower i by lender j at time t ($logSpread_{i,j,t}$) be a linear combination of aggregate credit market conditions at time t (Z_t), borrower credit quality ($X_{i,t}$), and the lender's ability to supply credit ($Y_{j,t}$).[11] Suppose further that borrower credit quality can be decomposed into publicly observable characteristics and lender-acquired private information ($X^O_{i,t}$ and $X^U_{i,t}$). Then, the orthogonal component of the variation in credit terms to the set of observables $\{Y_t, X^O_{i,t}, Z_{j,t}\}$ should be proportional to the variation in privately acquired information embedded into loan spreads.

$$\begin{aligned} logSpread_{i,j,t} &= Z_t\beta^0 + Y_{j,t}\beta^1 + X^O_{i,t}\beta^2 + X^U_{i,t}\beta^3 \\ &= Z_t\beta^0 + Y_{j,t}\beta^1 + X^O_{i,t}\beta^2 + \epsilon_{i,j,t} \end{aligned}$$

That is, the root-mean-squared error (RMSE) of a cross-sectional regression on credit terms on observable macroeconomic, borrower, and lender characteristics is representative of the extent of the average private information acquired and used by lenders to determine the terms of credit for that period in time. To implement this procedure I assume that the unobserved component pertains only to borrower characteristics pertinent to the loan terms. Further, the estimated RMSE would capture only the fraction of the unobserved component $X^U_{i,t}$ orthogonal to the observed factors. This procedure then requires a rich set of observables and sufficient controls for borrower, lender, and aggregate conditions. To fulfill this requirement I use balance sheet and stock market performance information for borrower characteristics, observable loan and lender characteristics as well as regulatory information for lender factors, and finally I use time dummy variables to absorb aggregate macro-economic and financial conditions.

In particular, I combine loan data with balance sheet, stock market, and credit rating information for the borrower characteristics. I use time fixed effects to control for macroeconomic and financial conditions where one period is defined as one calendar year. For lender and loan variables, I include a set of dummy variables indicating the

[11] The loan spread is defined as the average interest rate of the loan considering all fees and charges less some cost of financing measure such as the LIBOR.

lead lender, the loan purpose, as well as the loan amount and maturity. From the set of loan terms, I consider the loan rate (spread over a risk free measure) and collateralization as endogenous variables which contain both private and public information with the former being the dependent variable. For robustness, I include a specification with a dummy variable for collateralized loans. That is, I run an ordinary least squares regression on the following specification.

$$logSpread_{i,t} = \alpha_t + \theta_j + \delta_q + \xi_{fs} + \phi_{lp} + \sum_{k=1}^{K} \beta_t^k X_{i,t}^k + \epsilon_{i,t} \tag{2.1}$$

where α_t are time fixed effects for macro-economic conditions, θ_j are lender fixed effects, δ_q are two-digit SIC industry fixed effects, ξ_{fs} are dummy variables controlling for the auditor's report on the quality of the financial statements in *Compustat*, ϕ_{lp} are dummy variables on the stated purpose for the loan, and $X_{i,t}$ are a host of balance sheet and stock market performance characteristics for the borrower and a host of lender controls. Finally, I also include specifications which allow most coefficients ($\beta_{i,t}^k$) to vary over time.

To reduce the dimensionality of the specification, the regression is done in two stages. First, all other variables were regressed on lender, industry, auditor's opinion and loan purpose fixed effects. Second, I regress equation 2.1 for each year in my sample where the constant reflects the time dummy variable and a set of β_t^k coefficients are estimated for each year. I also consider a pooled specification where I run a regression on the whole sample and constrain the parameters to be constant over time and with time fixed effects.

The regressions give us time-varying estimates of the residual variance. I use the root-mean-squared error of each regression for each year as well as the ratio of the residual sum of squares to the total sum of squares as measures of private information production.[12].

Second, I take the square of the estimated residuals and regress these on a time trend or time fixed effects along with other controls to assess and verify the time-series properties of my proposed measure of private information production. That is, I run a

[12] I also conduct the same exercise on the level of spreads with similar results.

regression with the following specification.

$$\hat{\epsilon}_{i,t} = \gamma_t + \sum_{c=1}^{C} \beta^c X_{i,t}^c + \mu_{i,t} \tag{2.2}$$

Here, γ_t may represent a time trend or a set of dummy variables across time and $X_{i,t}^c$ are a set of controls including lender fixed effects, a dummy variable on whether a loan is collateralized, the number of lenders participating in the syndication, the logs of the size and maturity of the loan, and the log of the size of the borrower in terms of sales.[13] These regressions help isolate the time-variation in the measure for private information prediction that is not driven by time-variation in the control variables (i.e. borrowers tend to get larger over time).

In the next section, I describe in greater detail the data used for the regressions.

2.3 Data description

I use loan pricing data from *DealScan*, borrower balance sheet data from *Compustat*, and borrower equity stock returns from *CRSP*. I narrow the dataset to senior term loans only in the US by US-based non-Financial and non-Utilities (SIC codes 4900 and 6000) borrowers and US-based lenders. In the matching, I aggregate time into quarters beginning the fourth quarter 1987 until the third quarter of 2011. To avoid endogeneiety, I use the *Compustat* and *CRSP* data that would have been available in the quarter prior to the loan date. This yields 11,173 unique loans by 4,190 borrowers over 96 quarters and averaging 116.4 loans per quarter. The mean number of loans per borrower is 2.44 with a maximum of 21 loans. The average loan rate is 290 basis points over LIBOR, the average loan size is 198 million dollars, and an average maturity of 62 months.

From the *Compustat* and *CRSP* datasets, I include the log of total assets, the return on assets, the current ratio, interest coverage, the profit margin, the book to market value of equity, the leverage ratio, the quick ratio, the log of the market value of equity, the sales to total assets ratio, the ratio of total liabilities to the market value of total assets, the ratio of book to market value of total assets, and the ratio of property, plant,

[13]These are based on the specification of Cerqueiro et al. (2011) who conduct a similar exercise for small business loans.

Table 2.1: Data summary

	Dealscan	Compustat	CRSP	FRB BHC
Match Date	Quarter date of loan	Balance sheet as of quarter prior to loan date	Average from previous year	As of quarter of loan date
Match ID	Borrower and Lender ID	Borrower ID†	Borrower ID	Lender ID‡
	Facility amount	Log of total Assets	Mean	EBITDA
	Maturity	Leverage ratio (Total Debt to Assets)	Standard deviation	Net Income
	Spread over libor (fees and rates)	Return on Assets (Operating income over total Assets)	Skewness	Charge offs on C&I loans
	Secured debt dummy	Interest coverage (Operating income over interest payments)	Kurtosis	Recoveries on C&I loans
	Loan primary purpose	Profit margin (Operating income over Sales)		Total assets
	Number of lenders	Current ratio		Total liabilities
		Book to market value of equity and total Assets		Allowance for loan losses**
		Asset tangibility ratio		Tier 1 Capital**
		Log of market value of Equity		Tier 2 Capital
		Sales to total Assets ratio		Total risk-based capital**
		Firm age in years		Average total assets for leverage capital**
		Stock return (y-o-y)		Tier 1 leverage ratio**
		Stock vol: scaled high-low from previous year		Tier 1 risk-based capital ratio
		SIC Industry code of borrower		Total risk-based capital ratio**
		S & P Long term ratings**		
		S & P Short term ratings**		

† Matching based on Chava and Roberts (2008), updated August 2012.

‡ Manually matched lender names and loan dates using the National Information Center Institution database of the Federal Reserve System.

** Data available from 2002 onwards.

Asset tangibility ratio refers to Property, Plant, Equipment and Inventories over total assets.

Stock return moments from CRSP are computed using the daily returns in the preceding four quarters to the loan date.

The Federal Reserve Bank Holding Company dataset is available at the Federal Reserve Bank of Chicago `http://www.chicagofed.org/applications/bhc/bhc-home`.

equipment, and inventories to total assets as a measure of asset tangibility. I also include two stock market performance ratios from *Compustat* - the year-on-year stock market return given by the growth rate of the closing price of the current calendar year over the closing price of the previous calendar year and a proxy for stock price volatility given by the difference between the highest and lowest price of the current year over the average between the closing price of the current and previous calendar year. From 2002 onwards, I also observe short and long term S&P credit ratings for a subset of borrowers (matched to 2,019 out of possible 5,652 loans) which I have truncated into two categories, speculative and investment grade.[14] A speculative grade is defined as lower than a BBB rating for long term and lower than $A3$ for short term ratings. I also include the first four moments of stock returns of borrower equity matched from the *CRSP* database. After merging the datasets, I am left with 6,767 loans with at least some observations from all three datasets.[15]

To control for lender characteristics, I define lender dummy variables for the most frequently occurring lenders in the dataset. Here, I only select the top 100 unique lenders by loans facilitated. I also associate a lead lender and her characteristics only to the loans for which I can identify a unique lead lender.[16] About 75 percent of the loans in the sample are matched to one of the identified lead lenders with the remaining 25 percent of loans were from lead lenders which do not occur frequently enough in the dataset. I then match these lenders to the Federal Reserve's Bank Consolidated Holdings database using institution names and loan dates.[17] The following table provides a summary of the main loan, borrower, and lender characteristics. Table 2.7 in the appendix summarizes the full list of variables used in the regressions.

[14]The rest are categorized as unrated.

[15]Of these, only 5,512 have information for all variables in at least one of the specifications I use.

[16]The lead lenders were identified according to their lender role descriptions with *Administrative agent* or *Syndication agent* as top priority and *Agent* or *Sole lender* in the absence of the previous two. See Sufi (2007) for a similar approach.

[17]Only 51 of the identified lead lenders were completely matched to the Federal Reserve database. Of those not matched, one lender's Federal Reserve identification number could not be found, two could not be uniquely identified, and the rest were matched to the Federal Reserve ID system but could not be matched to a Holding company in the Federal Reserve Bank Consolidated Holdings database.

Table 2.2: Summary statistics

		Mean	Standard deviation	Count
Loan characteristics				
	Loan spread (bps)	290.49	154.10	11,173
	Loan amount (millions)	197.81	509.24	11,173
	Loan maturity (months)	61.99	24.14	11,173
	Number of lenders	3.02	3.20	8,352
	Secured loans indicator	0.94	0.24	8,567
Borrower characteristics				
	Size (log TA)	5.96	1.86	9,326
	Leverage ratio	0.65	1.06	9,324
	CRSP stock retun vol (daily)	0.04	0.02	6,767
	Credit ratings	n.a.	n.a.	5,652
Lender characteristics				
	Total Assets (100 billions)	6.64	6.80	5,203
	Chargeoffs on C&I loans (100 billions)	0.34	0.54	5,203
	Recoveries on C&I loans (100 billions)	0.67	0.81	5,203
	Tier 1 leverage ratio	6.64	1,44	2,693

The dependent variable in the regression is the log of the loan spread.[18] I consider 6 specifications in terms of the explanatory variables. The *Baseline* specification uses all available borrower characteristics except for credit ratings which are only available 2002 onwards and uses lender and time fixed effects to control for lender and aggregate factors. The *Expanded* specification includes lender characteristics from the Federal Reserve's regulatory database. A third specification includes a dummy variable if the loan is collateralized (*Expanded_c*). The fourth specification (*Baseline02*) augments the set of variables in *Baseline* with credit ratings data and limits the regression to only loans from 2002 onwards. Similarly the specification *Expanded02* augments the second specification with both credit ratings data and more lender variables from the regulatory database that become available from 2002 onwards. Finally, the last specification (*Expanded02_c*) adds the collateralization dummy variable to the previous specification.

[18]Strahan (1999) use the same dependent variable in an earlier study using the same dataset. I also consider specifications using the spread in levels with similar results.

2.4 Estimates of time-varying private information

As earlier stated, the regression is done in two stages. I first run a regression of the loan spread and all remaining explanatory variables on the dummy variables for the lead lender, industry classification, loan purpose, and auditor's opinion on the financial statements in *Compustat*. This leaves me with 6,980 loans with non-missing observations across the dummy variables in the first stage regressions.

I then run an ordinary least squares regression on equation 2.1 where all the coefficients are restricted to be the same over time. Table 2.3 reports the second stage regression results. The specifications can explain between 42 to 57 percent of the variation in loan spreads.[19] This is similar to the results in the earlier study by Strahan (1999) who had an (un-adjusted) R-squared of between 41 to 45 percent using a similar dataset and much higher than 25 percent in Cerqueiro et al. (2011) who focus on small business loans. The second and third specifications yield better fits than the baseline in the first set of regressions. The inclusion of additional control variables in the last three specifications, which restricts the sample to observations from 2002 onwards, yield significantly better fits in the third and fourth specifications but can only be done for observations beginning 2002.

Table 2.3: Pooled Regressions

	Baseline b/se	Expanded b/se	Expanded_c b/se	Baseline02 b/se	Expanded02 b/se	Expanded02_c
log loan amount	-0.007	0.017*	0.014	-0.008	0.017	0.024*
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
log loan maturity	0.078***	0.091***	0.022	0.089***	0.087***	0.037
	(0.02)	(0.02)	(0.03)	(0.03)	(0.03)	(0.04)
Borrower baseline Controls	Yes	Yes	Yes	Yes	Yes	Yes
Borrower credit ratings (post 2002)	No	No	No	Yes	Yes	Yes
d_secured			0.539***			0.501***
			(0.04)			(0.05)
Lender baseline Controls	No	Yes	Yes	No	Yes	Yes
Lender additional Controls (post 2002)	No	No	No	No	Yes	Yes
N	5512	3463	2605	3284	1995	1461
r2	0.225	0.267	0.340	0.292	0.366	0.431

Time, Industry, Lender, auditor opinion, and loan purpose fixed effects omitted

[19] The first stage regression explains about 25 percent of the variation in loan spreads.

In Figures 2.6 to 2.8 in the appendix, I plot the average RMSE for all the specifications (full sample and post 2002). There is substantial variation in the variance of the residual over time. Further, although the inclusion of additional explanatory variables improves the model fit for the 2002 and onwards specification, the results suggest that the various specifications yields the same patterns over time. I also run a regression where I allow most of the coefficients in equation 2.1 to vary over time (annually). Estimates of time variation in the RMSE from these regressions are similar to the first set and are also available in the appendix.

Table 2.4: Squared residual regressions

	(1)	(2)	(3)	(4)	(5)	(6)
	squaredlresid	squaredlresid	squaredlresid	squaredlresid	squaredlresid	squaredlresid
	b/se	b/se	b/se	b/se	b/se	b/se
yearstart	-0.002**		-0.000	-0.002		-0.002
	(0.00)		(0.00)	(0.00)		(0.00)
Rec less 3			0.046			0.059
			(0.03)			(0.04)
Rec less 2			0.024			0.007
			(0.02)			(0.03)
Rec less 1			-0.014			-0.023
			(0.02)			(0.03)
Rec year			-0.053***			-0.064***
			(0.02)			(0.02)
Rec plus 1			-0.056***			-0.061**
			(0.02)			(0.03)
Rec plus 2			-0.065***			-0.069***
			(0.02)			(0.03)
Rec plus 3			-0.053***			-0.044
			(0.02)			(0.03)
Secured loan				-0.073***	-0.081***	-0.075***
				(0.02)	(0.02)	(0.02)
Log sales				0.010	0.012*	0.010*
				(0.01)	(0.01)	(0.01)
Log loan maturity				-0.037*	-0.049**	-0.045**
				(0.02)	(0.02)	(0.02)
Log loan amount				-0.010	-0.015	-0.010
				(0.01)	(0.01)	(0.01)
Number of lenders				-0.006***	-0.006***	-0.006***
				(0.00)	(0.00)	(0.00)
Constant	3.741**	0.070***	0.852	4.921	-0.013	3.510
	(1.73)	(0.01)	(2.03)	(3.60)	(0.05)	(3.20)
Time variable	Time trend	Time FE	Time trend and rec dum	Time trend	Time FE	Time trend and rec dum
R2	0.00119	0.0345	0.0166	0.0521	0.0808	0.0672
N	3441	3441	3441	2586	2586	2586

Lender fixed effects omitted.

Following Cerqueiro et al. (2011), I also conduct a heteroskedasticity-type regression on the squared residuals of specifications *Expanded* for observations prior to 2002 and *Expanded02* thereafter and verify how much of the time-variation in the RMSE is associated with time conditional on other potential explanatory variables. This specification was chosen as it includes time-varying lender-specific characteristics but omits the securitization variable which may be endogenous to the private information I am trying to capture with the residual. I consider a time trend, time fixed effects, and dummy variables around recession dates (annual) as measures of variation over time. I also include a dummy variable for collateralized loans, the log size and maturity of the loan, the size of the borrower in terms of log sales, the number of lenders, and lender fixed effects as additional explanatory variables. The results are reported in Table 2.4.

The first column includes only a time trend and here the squared residuals do appear to decline over time although this coefficient is no longer significantly different from zero once I add other control variables (column four). I then include dummy variables for three years before and up to three years after the beginning of a recession in columns three and six. Finally, columns two and and five use time fixed effects with and without additional controls. Figure 2.1 plots the estimated time fixed effects coefficients in column five of Table 2.4.

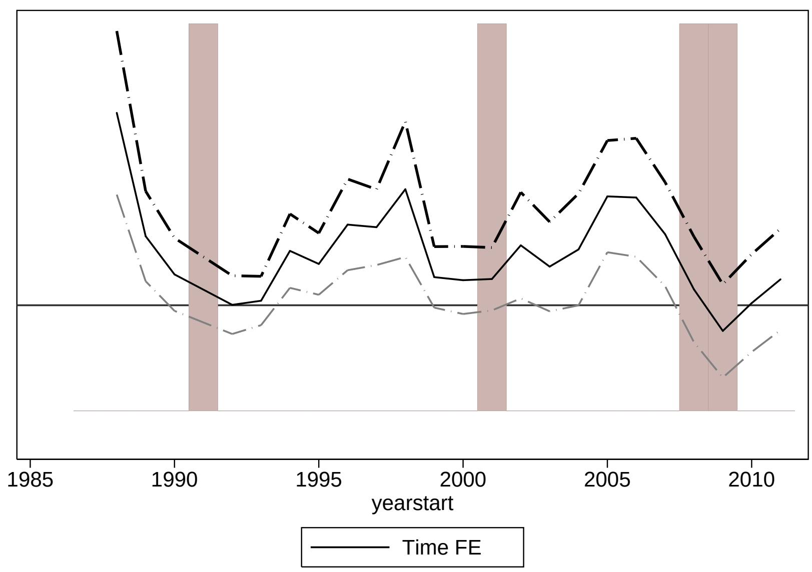

Figure 2.1: Time Fixed effects: squared residuals regression

The estimates suggest a hump-shaped pattern for information production over expansions. In these regressions, there appears to be a downward trend around recession start dates which bottoms out two years after the start of a recession. Interestingly, and consistent with the theoretical literature, private information production appears to decline well before the start of a recession where the estimates of the coefficients suggest that the unexplained variation in loan spreads decline from three years before the start of a recession. With respect to the other explanatory variables, the squared residual is lower for secured loans, loans with longer maturity, and for more lenders in the syndication. These are consistent with the findings of Sufi (2007).[20] Of these, only the negative relationship with loan maturities seems counter-intuitive. Perhaps, longer maturities capture other factors (e.g. deeper relationships between lenders and borrowers) not included in the specifications.

2.5 Stock market reaction to private information

Interpreting the residuals as private information about borrower quality is based on two major assumptions. The first is that we have not omitted publicly available information which is not spanned by our list of regressors. This appears unlikely given the large set of balance sheet and market-based measures we account for. The second concern is whether the estimated residual is simply idiosyncratic variation. To address this concern I check whether the residual I estimate elicits a stock market reaction to the equity of publicly listed borrowers following the loan agreement. Intuitively, the residual represents deviations of the loan spread from that predicted by publicly observable factors. Hence, if stock market participants consider this deviation as informative about the borrowers' conditions, the stock market performance of the borrowers' equity following the loan agreement should partially depend on the estimated residual. To do this, I estimate the sensitivity of the cumulative abnormal returns of the borrower's stock price to the estimated residual in the days following the loan issuance.

I combine the estimated residuals with the daily returns data from *CRSP* and am able to match 1,551 loan events for which I have returns data and do not observe another

[20]He finds that the percentage of the loan held by the lead lender or the Herfindahl concentration index across lenders in syndicated loans are decreasing in the log maturity of the loan and when the loan is secured.

loan in the preceding 125 trading days for each event. I date each loan event using the loan issue date. The literature has typically used loan announcement dates culled from various media sources although several (e.g. Harvey et al. 2004, Focarelli et al. 2008, Li and Ongena 2014) use the loan issue date as I do. The use of loan issue dates may potentially underestimate the effects if the loan details are observed by the market at a loan announcement date significantly different from the loan issue date[21]. The potential underestimation bias may be less severe for the loan spread variable we consider as typical news announcements may not include the full terms of the loan contract. On the other hand the database I use is also accessible in real-time to stock market participants and thus we may reasonably assume that these terms of credit become public around the dates these loans have been issued.

I first generate a predicted series for each equity return by estimating a Fama-French factor model using a 120-day window ending 5 days prior to each loan event[22]. That is I estimate

$$\begin{aligned} \hat{R^e_{i,t}} &= \hat{\alpha_i} + \hat{\beta_{i,1}} Mkt^e_t + \hat{\beta_{i,2}} SMB_t + \hat{\beta_{i,3}} HML_t \\ AR_{i,t} &= R^e_{i,t} - \hat{R^e_{i,t}} \end{aligned}$$

where $R^e_{i,t}$ are excess returns of borrower i at time t and Mkt^e_t, SMB_t, and HML_t are the Fama-French size and book-to-market factors. The estimation window is $(\tau - 125, \tau - 5)$ where τ is the event date. I then compute abnormal returns around the loan event by subtracting the Fama-French model predicted returns from the actual daily returns. I also Winsorize the abnormal returns and compute the cumulative abnormal return for three event windows - from a day before to a day after the loan event, between two days before and two days after the loan event, and between three days before and

[21]This appears to be a small matter in the event study analysis of Harvey et al. (2004) using bond and syndicated debt issuances between 1980 to 1997 for emerging market issuers.

[22]Daily Fama-French factors are obtained from Kenneth French's website on 07 May 2014

three days after.

$$
\begin{aligned}
CAR_{i,-1,1} &= \sum_{t=\tau-1}^{\tau+1} AR_{i,t} \\
CAR_{i,-2,2} &= \sum_{t=\tau-2}^{\tau+2} AR_{i,t} \\
CAR_{i,-3,3} &= \sum_{t=\tau-3}^{\tau+3} AR_{i,t}
\end{aligned}
$$

The following figures plot the average abnormal returns and cumulative abnormal returns from three days before the loan event to three days after.[23] The solid lines depict average abnormal and cumulative abnormal returns for borrowers who obtained a negative residual in the loan spread regressions (a lower interest rate than predicted by observables) and the dashed line represents the averages for borrowers who obtained a positive residual.

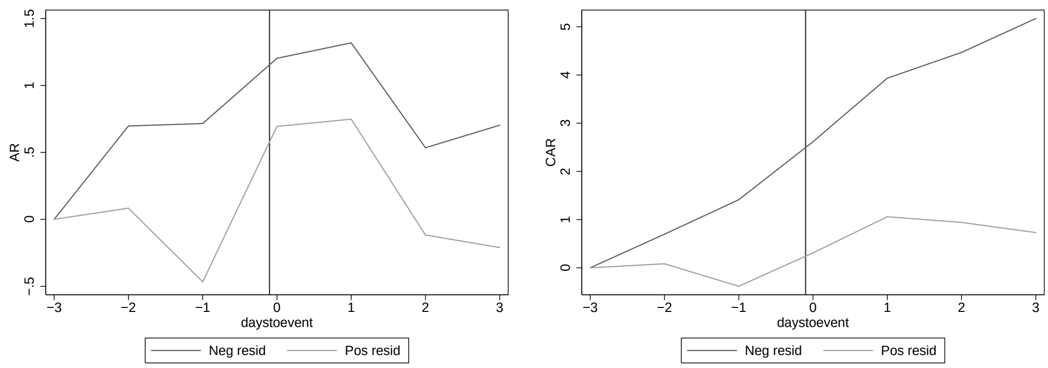

(a) Abnormal returns

(b) Cumulative abnormal returns

Figure 2.2: Averages around loan event

I then check whether the cumulative abnormal return for the series are systematically different for loans with spread higher or lower than predicted by constructing a dummy variable which takes the value of one if the residual from the loan regression is negative.

[23] For this plot, the cumulated abnormal return is computed starting at $\tau - 3$.

The main regression specification is as follows:

$$CAR_{i,t,T} = \alpha + \beta \mathbb{1}_{\hat{\epsilon}_i<0} + \gamma X_i + \mu_i$$

where the left hand side variable is the cumulative abnormal return, $\mathbb{1}_{\epsilon_i<0}$ is a dummy variable for a negative residual, and X_i are a set of controls including the size (in log total assets), book-to-market value, and leverage ratio of the borrower, along with the ratio of the loan amount to total assets of the borrower, the log maturity of the loan, the log size of the loan, and a dummy variable if the loan took place during a recession.

A typical event study would have α as the variable of interest. My main parameter of interest however is β which we expect to be positive on the dummy variable for negative residuals. Table 2.5 reports the regression results. Here I find that the coefficient is only significantly different from zero and positive at the ten percent confidence level for the three day before and after event window.

Table 2.5: CAR regression results: negative residual dummy

	(1)	(2)	(3)	(4)	(5)	(6)
	CAR11	CAR11	CAR22	CAR22	CAR33	CAR33
	b/se	b/se	b/se	b/se	b/se	b/se
Neg residual dummy	-0.597	-0.532	0.098	0.021	0.974*	1.016*
	(0.56)	(0.56)	(0.56)	(0.56)	(0.57)	(0.57)
Constant	1.223***	0.455	-0.242	1.085	-0.238	-0.596
	(0.41)	(1.70)	(0.40)	(2.07)	(0.42)	(2.00)
controls	No	Yes	No	Yes	No	Yes
r2	0.000739	0.00517	0.0000196	0.0111	0.00189	0.00726
N	1554	1554	1551	1551	1550	1550

I also consider a specification with the actual residual spread as an explanatory variable with similar results.

$$CAR_{i,t,T} = \alpha + \beta \hat{\epsilon}_i + \gamma X_i + \mu_i$$

In this case, we expect the coefficient on the residual to be negative. For these regressions, I also include the predicted loan spread as an additional control variable. Table 2.6 reports regression results. Again, we only find significant and negative coefficients for the three day before and after event window.

Table 2.6: CAR regression results: residuals

	(1) CAR11 b/se	(2) CAR11 b/se	(3) CAR22 b/se	(4) CAR22 b/se	(5) CAR33 b/se	(6) CAR33 b/se
Residual	1.216	1.088	0.531	0.606	-1.256*	-1.341**
	(0.75)	(0.76)	(0.73)	(0.74)	(0.68)	(0.68)
Predicted spread	0.040	-0.095	-0.124	-0.284	0.995*	0.974
	(0.53)	(0.56)	(0.63)	(0.67)	(0.60)	(0.62)
Constant	0.727	0.714	0.476	2.686	-5.057	-5.534
	(2.72)	(3.64)	(3.27)	(4.31)	(3.11)	(4.03)
controls	No	Yes	No	Yes	No	Yes
r2	0.00173	0.00595	0.000334	0.0116	0.00334	0.00851
N	1554	1554	1551	1551	1550	1550

The results provide suggestive evidence that, indeed, obtaining a lower interest rate than that which can be predicted by observable factors lead to a positive impact on the stock market return of a borrower. However, this does not directly imply that the residual contains previously private information that the lenders had about the borrowers. For instance, a lower-than-predicted interest rate on the loan implies a lower-than-predicted stream of future interest payments for the borrowers and thus, in itself, lead to an improvement in the borrower's future financial conditions. Disentangling the signalling role of the residual term, in revealing previously private information about the borrower, from its direct effect on the borrowers' equity returns require further investigation and is left for future work.

2.6 Conclusion

Consistent with existing theory, my results suggest that information production in credit markets varies over the business cycle. It appears to initially increase in the beginning of an expansion but starts to decline well before the next recession. Nevertheless, my interpretation of RMSE as private information production comes with several caveats. First, the interpretation requires that the specifications I use span the full set of publicly available information on borrowers used by lenders and relevant for the determination of the terms of the loan. In principle, the RMSE contains all variation in omitted variables. Thus, crucial to my interpretation of the RMSE as a measure for private information production in credit markets is that these omitted variables are only those unobserved borrower characteristics known to the lender and used to set the terms of credit. The extensive set of balance sheet variables, stock market performance and credit rating measures, and lender controls would suggest that this assumption is likely to hold. My specification attempts to encompass a wide set of publicly available information from which an outsider may evaluate the creditworthiness of these borrowers.

Further, my procedure is subject to a potential selection bias in that I only observe loans that have been granted and my sample consists of loans for which I can match observations across all four datasets. For the former, it may be the case that public and private information has been used to screen borrowers (i.e. reject or accept loan applications). Under this scenario, the $RMSE$ is only indicative of the intensive margin to information production in fine-tuning the terms of the loan contract whereas I do not account for the extensive margin or the production and use of private information to reject or accept loans. Thus, my results potentially underestimate actual private information production and use in credit markets. Finally, it should be noted that the borrowers in my sample, being those that are present in both *Compustat* and *Dealscan* are typically large publicly listed corporations for which information is relatively more publicly available.

Limitations aside, this approach is one of the first attempts to estimate time variation in private information production in credit markets. The evidence I show is also consistent with the broader theoretical literature on credit market screening and the business cycle. These findings serve to complement the existing literature and lends support to cyclical macro-prudential policies which aim to reduce the build-up of systemic fragility in credit booms. If the cyclical variations in private information production documented in this chapter lead to variations in the quality of the loan portfolios of financial intermediaries, then it would seem that indeed it is prudent to impose tighter restrictions on financial risk-taking around the phases of the cycle where private information production begins to decline. Given the evidence presented in this chapter, this coincides with about three years prior to the start of the last two recessions in the United States. In light of this, fine-tuning pro-cyclical macro-prudential policies to address the fall in private information production in credit markets require the identification of early warning indicators and is left for future work.

Bibliography

Agarwal, S. and Hauswald, R. (2010). Distance and private information in lending. *Review of Financial Studies*, 23(7):2757–2788.

Bar-Isaac, H. and Shapiro, J. (2013). Ratings quality over the business cycle. *Journal of Financial Economics*, 108:62–78.

Best, R. and Zhang, H. (1993). Alternative information sources and the information content of bank loans. *Journal of Finance*, 48 (4):1507–1522.

Billet, M., Flannery, M., and Garfinkel, J. (1995). The effect of lender identity on a borrowing firm's equity return. *Journal of Finance*, 50 (2):699–718.

Bolton, P., Freixas, X., and Shapiro, J. (2012). The credit ratings game. *Journal of Finance*, 67(1):85–111.

Cerqueiro, G., Degryse, H., and Ongena, S. (2011). Rules versus discretion in loan rate setting. *Journal of Financial Intermediation*, 20:503–529.

Cerqueiro, G., Ongena, S., and Roszbach, K. (2014). Collateralization, bank loan rates, and monitoring. *Journal of Finance*, forhcoming.

Chava, S. and Roberts, M. (2008). How does financing impact iinvestment? the role of debt covenenants. *Journal of Finance*, 63:2085–2121.

Degryse, H. and Ongena, S. (2005). Distance, lending relationship, and competition. *Journal of Finance*, 60 (1):231–266.

Dell'Ariccia, G., Igan, D., and Laeven, L. (2012). Credit booms and lending standards: Evidence from the subprime mortgage market. *Journal of Money, Credit and Banking*, 44(2-3):367–384.

Dell'Ariccia, G. and Marquez, R. (2006). Lending booms and lending standards. *Journal of Finance*, 61 (5):2511–2546.

Focarelli, D., Pozzolo, A. F., and Casolaro, L. (2008). The pricing effect of certification on syndicated loans. *Journal of Monetary Economics*, 55:335–349.

Gorton, G. and He, P. (2008). Bank credit cycles. *Review of Economic Studies*, 75:1181–1214.

Gorton, G. and Ordonez, G. (2014). Collateral crises. *American Economic Review*, 104 (2):343–378.

Griffin, J. and Tang, D. Y. (2012). Did subjectivity play a role in cdo credit ratings? *Journal of Finance*, 67(4):1293–1328.

Harvey, C., Lins, K., and Roper, A. (2004). The effects of capital structure when expected agency costs are extreme. *Journal of Financial Economics*, 74:3–30.

Hauswald, R. and Marquez, R. (2006). Competition and strategic information acquisition in credit markets. *Review of Financial Studies*, 19 (3):967–1000.

Holmstrom, B. and Tirole, J. (1997). Financial intermediation, loanable funds, and the real sector. *Quarterly Journal of Economics*, 112 (3):663–691.

James, C. (1987). Some evidence on the uniqueness of bank loans. *Journal of Financial Economics*, 19:217–235.

Li, C. and Ongena, S. (2014). Bank loan announcements and borrower stock returns before and during the recent financial crisis. *Swiss Finance Institute Research Paper Series*, 14-26.

Lummer, S. and McConnell, J. (1989). Further evidence on the bank lending process and the capital-market response to bank loan agreements. *Journal of Financial Economics*, 25:99–122.

Norden, L. and Roscovan, V. (2013). The dynamics of credit rating disagreement. *Working Paper*.

Ongena, S. and Roscovan, V. (2013). Bank loan announcements and borrower stock returns: does bank origin matter? *International Review of Finance*, 13 (2):137–159.

Ongena, S., Roscovan, V., Song, W.-L., and Werker, B. J. (2014). Banks and bonds: The impact of bank loan announcements on bond and equity prices. *Journal of Financial Management Markets and Institutions*, 2 (2):131–155.

Petriconi, S. (2012). Bank competition, information choice, and inefficient lending booms. *Job Market Paper*.

Ruckes, M. (2004). Bank competition and credit standards. *The Review of Financial Studies*, 17 (4):1073–1102.

Slovin, M., Johnson, S., and Glascock, J. (1992). Firm size and the information content of bank loan announcements. *Journal of Banking and Finance*, 16:1057–1071.

Strahan, P. (1999). Borrower risk and the price and nonprice terms of bank loans. *Federal Reserve Bank of New York Staff Reports*, 90.

Sufi, A. (2007). Information asymmetry and financing arrangements: Evidence from syndicated loans. *Journal of Finance*, 62 (2):629–668.

2.A Data descriptive statistics

Table 2.7 summarizes the general features of the data obtained from various sources and the regression specifications. The dependent variable is the spread in terms of the average interest rate (all fees and charges as a percent of the loan amount) less LIBOR. The second and third terms refer to the size and maturity of the loan. For all three variables, summary statistics are reported on the levels although the regression specification have these variables in logs.

Table 2.7: Descriptive statistics and specifications

	Specification						Statistics		
	Baseline	Expanded	Expanded_c	Baseline02	Expanded02	Expanded02_c	Mean	sd	count
	Dependent variable								
Loan spread* (bps)	Yes	Yes	Yes	Yes	Yes	Yes	290.49	154.10	11,173
	Explanatory variables								
First stage variables									
Lender dummy variables	Yes	Yes	Yes	Yes	Yes	Yes	n.a.	n.a.	8,352
Industry dummy variables	Yes	Yes	Yes	Yes	Yes	Yes	n.a.	n.a.	9,527
Loan purpose dummy variables	Yes	Yes	Yes	Yes	Yes	Yes	n.a.	n.a.	11,173
Auditor's opinion dummy variables	Yes	Yes	Yes	Yes	Yes	Yes	n.a.	n.a.	9,312
Year dummy variables	Yes	Yes	Yes	Yes	Yes	Yes	n.a.	n.a.	11,173
Loan-specific variables									
Loan amount* (millions)	Yes	Yes	Yes	Yes	Yes	Yes	197.81	509.24	11,173
Maturity* (months)	Yes	Yes	Yes	Yes	Yes	Yes	61.99	24.14	11,173
Number of lenders	No	Yes	Yes	No	Yes	Yes	3.02	3.20	8,352
Dummy for secured loans	No	No	Yes	No	No	Yes	0.94	0.24	8,567
Borrower-specific variables									
Size (log TA)	Yes	Yes	Yes	Yes	Yes	Yes	5.96	1.86	9,326
Return on assets	Yes	Yes	Yes	Yes	Yes	Yes	0.10	1.12	0,293
Current ratio	Yes	Yes	Yes	Yes	Yes	Yes	2.12	5.28	9,085
Interest coverage	Yes	Yes	Yes	Yes	Yes	Yes	51.53	921.30	9,152
Profit margin	Yes	Yes	Yes	Yes	Yes	Yes	-0.18	12.54	9,268
Book-to-market ratio	Yes	Yes	Yes	Yes	Yes	Yes	125.73	3,972.26	8,330
Leverage ratio	Yes	Yes	Yes	Yes	Yes	Yes	0.65	1.06	9,324
Tangible assets ratio	Yes	Yes	Yes	Yes	Yes	Yes	0.44	0.24	9,237
Quick ratio	Yes	Yes	Yes	Yes	Yes	Yes	1.55	5.26	9,059
Asset book to market	Yes	Yes	Yes	Yes	Yes	Yes	0.77	0.30	8,342
Stock return mean (y-o-y)	Yes	Yes	Yes	Yes	Yes	Yes	5.43	398.79	7,717
Stock price range (hi-lo)	Yes	Yes	Yes	Yes	Yes	Yes	0.84	0.93	7,717
CRSP stock return mean (daily)	Yes	Yes	Yes	Yes	Yes	Yes	0.00	0.00	6,767
CRSP stock retun vol (daily)	Yes	Yes	Yes	Yes	Yes	Yes	0.04	0.02	6,767
CRSP stock retun skew (daily)	Yes	Yes	Yes	Yes	Yes	Yes	-0.12	2.15	6,767
CRSP stock retun kurtosis (daily)	Yes	Yes	Yes	Yes	Yes	Yes	12.89	22.06	6,767
Credit ratings	No	No	No	Yes	Yes	Yes	n.a.	n.a.	5,652
Lender variables									
EBITDA	No	Yes	Yes	No	Yes	Yes	3.61	4.63	5,203
Net Income	No	Yes	Yes	No	Yes	Yes	3.64	4.66	5,203
Chargeoffs on C&I loans	No	Yes	Yes	No	Yes	Yes	0.34	0.54	5,203
Recoveries on C&I loans	No	Yes	Yes	No	Yes	Yes	0.67	0.81	5,203
Total Assets	No	Yes	Yes	No	Yes	Yes	6.64	6.80	5,203
Total Liabilities	No	Yes	Yes	No	Yes	Yes	6.09	6.20	5,203
Allowance for loan losses (beg.)	No	No	No	No	Yes	Yes	9.78	10.25	2,719
Tier 1 capital	No	No	No	No	Yes	Yes	4.95	4.34	4,192
Tier 2 capital	No	No	No	No	Yes	Yes	2.68	1.62	2,690
Total risk-based capital	No	No	No	No	Yes	Yes	6.97	5.89	4,192
Average TA for leverage capital	No	No	No	No	Yes	Yes	7.81	6.57	4,192
Tier 1 leverage ratio	No	No	No	No	Yes	Yes	6.64	1,44	2,693
Tier 1 risk-based capital ratio	No	No	No	No	Yes	Yes	9.07	1.80	2,693
Total risk-based capital ratio	No	No	No	No	Yes	Yes	12.67	1.93	2,693

* Reported summary statistics are on levels although regression specifications use logs.

The following plots the evolution of the quarterly average values of various terms of the loan over our sample period.

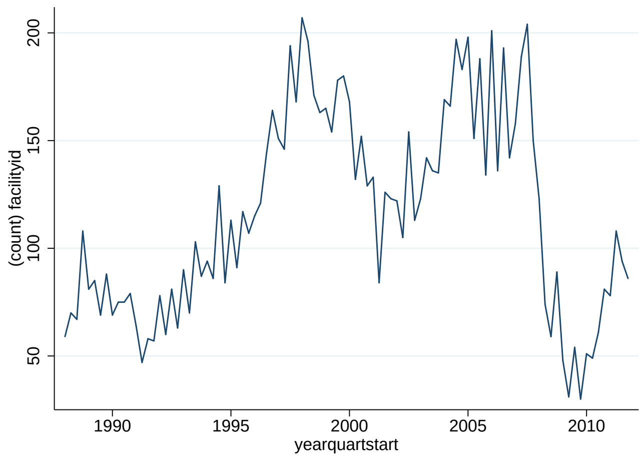

Figure 2.3: Number of loans

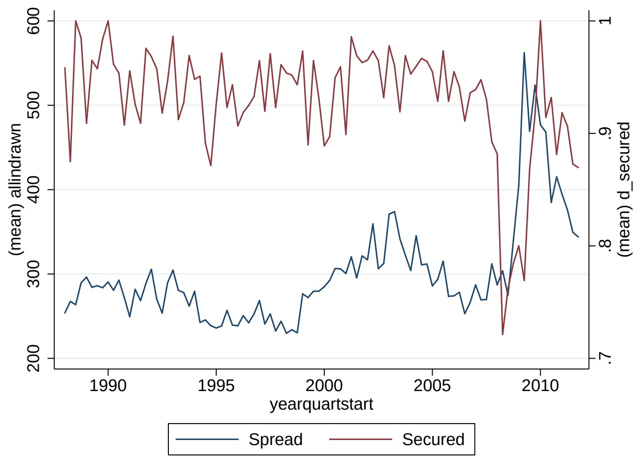

Figure 2.4: Loan spread and collateralization rates

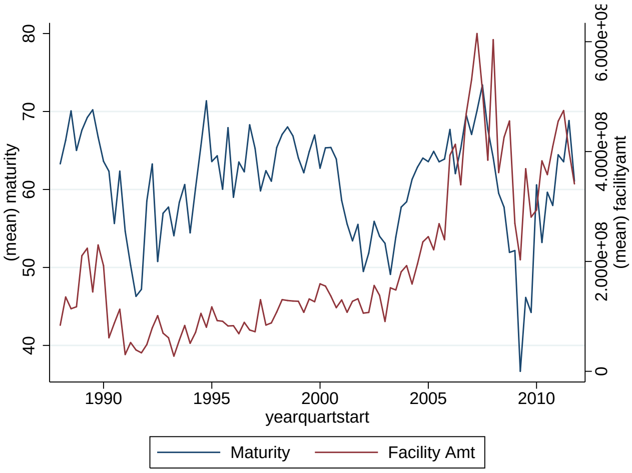

Figure 2.5: Loan maturity (months) and amounts

2.B Loan spread regression figures

The following figures plot the root-mean-squared-errors for the regression specifications without time-varying coefficients and reported in Table 2.3. In the first figure, the specification *Baseline* appears as the solid line while *Baseline02* with variables available from 2002 onwards is represented as the dashed line. The next two figures plot the results from the remaining four specifications in the same manner.

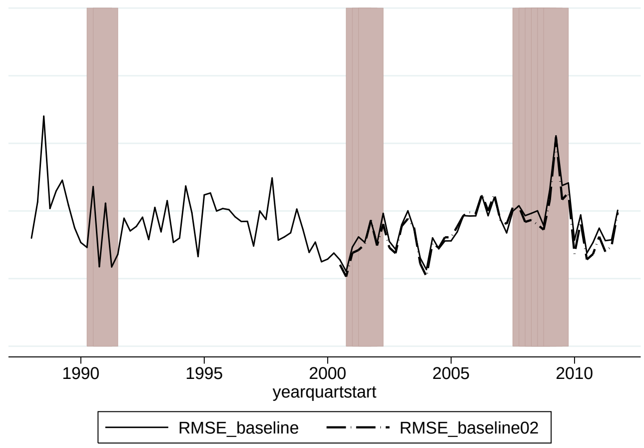

Figure 2.6: Baseline specification estimated RMSE

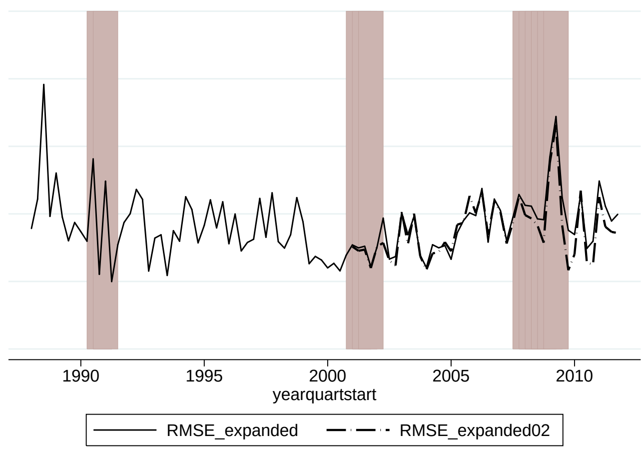

Figure 2.7: Expanded specification

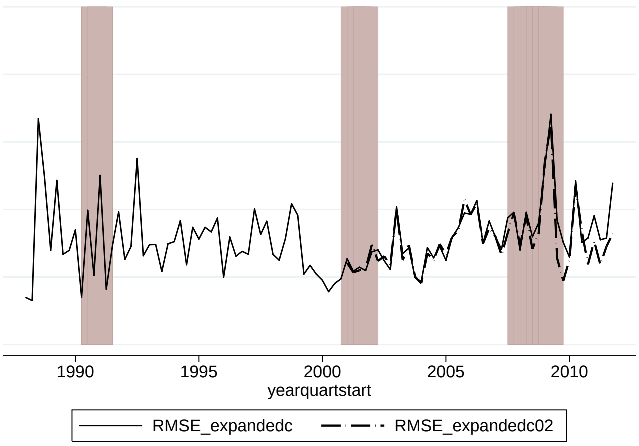

Figure 2.8: Expanded specification with collateral dummy

The next set of figures report estimates of private information with time-varying coefficients. For each figure, I plot the RMSE and proportion of unexplained variation (*RSS2TSS*) for each year in the sample using the *Baseline* and *Baseline02* specifications for the periods 1987 to 2001 and 2002 onwards respectively in the first plot and do the same for the remaining specifications in the next two plots. Note that in the figures the specification changes from the year 2002 onwards and relative comparison of the two measures should strictly only be done within the periods 1990 and 2001 and 2002 to 2011. However, as the preliminary results suggest, this is not a major concern. There were an insufficient number of observations to conduct the second-stage regression for the first few years of the sample.

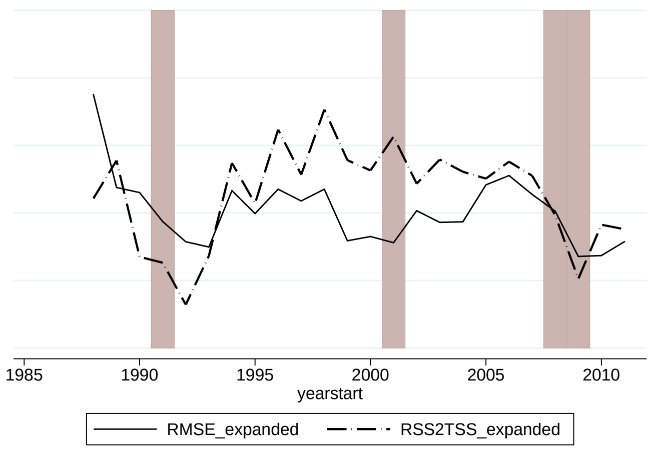

Figure 2.9: Baseline specification estimates of private information

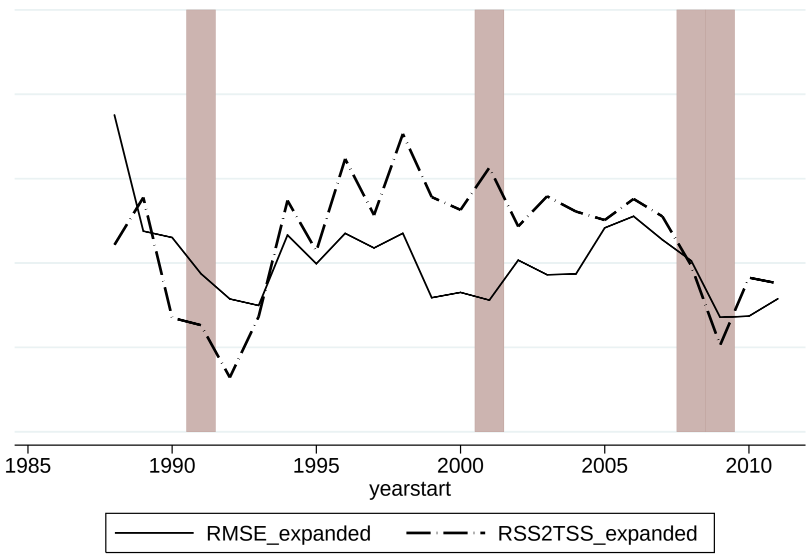

Figure 2.10: Expanded specification estimates of private information

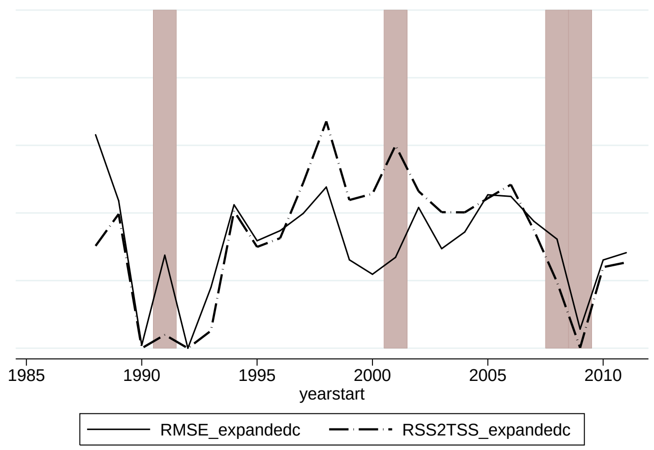

Figure 2.11: Expanded specification with collateralization estimates of private information

Chapter 3

EVALUATING FORECAST OPTIMISM AND CONFIDENCE IN THE SURVEY OF PROFESSIONAL FORECASTERS

3.1 Introduction

Excessive optimism and overconfidence are thought to be significant factors in driving recent boom-bust episodes and may have played a significant role in triggering financial crises. Although recent contributions to the macroeconomic and financial literature provide theories and mechanisms which take these into account, there is little in terms of empirical support and guidance with respect to the business cycle features of over-optimism and overconfidence. In this chapter, I use over two decades of survey forecast data from professional forecasters in the U.S. to help fill this gap. I use density forecasts of real GDP growth from the Survey of Professional Forecasters in the United States across multiple forecasters and forecast horizons and employ existing methodologies to identify biases in over-optimism and overconfidence.

Timely and reliable forecasts of future economic conditions are a crucial input to economic agents and policy makers alike. Surveys of professional forecasts provide such information. These surveys include density forecasts which offer users not only with a mean prediction of future realizations of macroeconomic variables but also their variance. Existing tests in the forecast evaluation literature typically evaluate the overall predictive accuracy of density forecasts. However, perhaps equally important to policy makers and other users of such forecasts is to understand the nature of these variations in predictive accuracy. That is, it is quite important to discern along what dimension these forecasts systematically deviate from realizations of the forecast target. In this chapter, I decompose the tests proposed in the literature into tests for over-optimism and overconfidence to help guide users of survey forecasts. I define over-optimism as the case when the mean of a density forecast is consistently higher than the actual realization of real GDP growth. Similarly, over-confidence is defined as the case when the standard deviation of the density forecast is lower than the standard deviation of the forecast error.

The procedure I employ relies on the assumption that real GDP growth is conditionally Normally distributed given information at the time a forecast is made. Under this assumption the standardized forecast error, defined as the difference between the realized GDP growth and the mean of the forecast density divided by the implied standard deviation from the forecast density, is also Normally distributed with mean equal to the optimism bias and variance equal to the confidence bias. I estimate these parameters at the individual and forecast horizon level using the responses from the quarterly Survey of Professional Forecasters (SPF) from 1992Q1 to 2013Q4. I also provide estimates which take into account a particular form of the correlation structure given by the three dimensional nature of the survey data proposed by Davies and Lahiri (1995).

I find the following. First, the results indicate that, in general, the survey density forecasts are overconfident. The variance of forecast errors are about 2.88 times larger than the implied variance of forecast densities on average. Univariate tests confirm the statistical significance of the overconfidence bias for a substantial proportion of forecasters and forecast horizons with about 65 percent of the tests rejecting the null hypothesis of no bias. On the other hand, the results confirm the existence of over-optimism biases to a lesser degree with only 22 percent of the tests rejecting the null hypothesis of no bias. Second, the estimated biases appear to be increasing in the forecast horizon. I also estimate the optimism and confidence bias on 5-year rolling sub-samples of the survey responses and find some variation over time. It appears that the most recent U.S. expansion and recession featured the highest levels of overconfidence bias.

Related Literature. The evaluation of survey forecasts has a long history in the economic literature stemming from evaluating the predictability of forecast errors of point forecasts as in Mincer and Zarnowitz (1969) to more recent contributions such as Coibon and Gorodnichenko (2012); Patton and Timmerman (2012); Coibon and Gorodnichenko (2015) which take advantage of the availability of forecasts at multiple horizons. Since the introduction of survey density forecasts, which can potentially shed more light behind forecasters' expectation formation process, new tools and methodologies to assess the overall predictive accuracy of such forecasts have also began to appear in the literature. Of particular relevance are Diebold et al. (1998) and later on Berkowitz (2001) who respectively propose tests on the distribution of probability integral transforms (PITs) of the forecasts and inverse-normal transformations of the PITs.[1]

Diebold et al. (1998) focus on density forecasts and evaluates the predictive accuracy of the density forecast as a whole by looking into probability integral transforms (PITs). Intuitively, if we transform a random variable by feeding its realization into its cumulative density function (cdf) we would have generated an independent and identically distributed standard uniform random variable. The procedure yields several benefits. First, the ranking it induces is applicable to any loss function that the forecaster may have in generating their forecasts (i.e. a forecaster would prefer forecast densities which are closer to the true density). Second, it does not require knowledge of or estimation of the true densities. The PITs measure relative differences in the forecasted vis-a-vis true densities. We may then use a variety of tests for uniformity and independence.

This paper builds on the inverse-normal transformation of PITS first suggested by Berkowitz (2001). Such a transformation generates a random variable which, under the null hypothesis of predictive accuracy, is distributed as a standard Normal. The second transformation provides us with a variable that is easier to work with and allow us to employ a variety of existing tests of independence and distribution fit. In his paper, Berkowitz focus on testing for independence of the transformations and whether they follow a standard Normal distribution (a joint test on zero mean and unit variance) where he proposes likelihood ratio tests on the estimated parameters.[2]

[1] Comprehensive surveys of the nature, use and evaluation of density forecasts may be found in Corradi and Swanson (2006); Pesaran and Weale (2006); Mitchell and Wallis (2011); Boero et al. (2014) and Rossi (2014).

[2] More recent contributions along this line are Kalliovirta (2012) and Gonzales-Rivera and Yoldas

However, perhaps equally important to policy makers and other users of such forecasts is to understand the nature of these variations in predictive accuracy. That is, it is quite important to discern along what dimension these forecasts systematically deviate from realizations of the forecast target. Diebold et al. (1998) for example propose a visual analysis of quantile plots of PITs to assess the nature of deviations in predictive accuracy. Two dimensions of forecast accuracy of particular relevance are on the accuracy of the mean of the forecast density and its variance. By and large the literature has focused on the accuracy of the forecast mean - whether forecasts errors are consistently positive or negative. A few have also looked at whether the implied variance of forecast densities accurately reflect the conditional variance of the forecast target.[3] A seminal contribution along this avenue is the coverage test of Christoffersen (1998) which looks at how often the forecast error is within a given confidence interval implied by forecast densities.

One of the few contributions to the literature which look at both the mean and the variance is Giordani and Soderlind (2006). They look at pessimism (optimism) and doubt (over-confidence) in both point and density survey forecast data and focus on its implications on the equity premium. To test for optimism (pessimism), they use point forecast from the SPF and Livingston surveys of real GDP and consumption (1982-2003). To test for doubt, they apply the approach of Christoffersen (1998) and test whether confidence intervals derived from the SPF survey of real GDP growth, either by fitting a normal distribution to the responses or uniform-within-bin (categorical distribution), accurately reflect the occurrence of the forecast target at a given confidence level. The main result is that forecasters tend to underestimate uncertainty, even after adjusting for the pessimism/optimism bias obtained using the point forecasts. I use an alternative methodology which jointly estimates both biases from survey density forecasts.

More recently, Kenny et al. (2014) and Kenny et al. (2015) are two related works which use the ECB's survey of professional forecasters density forecasts on real GDP growth and inflation. In the former, they evaluate how the predictive accuracy of pro-

(2012) which extends the Berkowitz (2001) framework to the multivariate case (i.e. multiple forecast targets).

[3] A related stream of the literature focuses on the measurement of forecaster uncertainty starting from Zarnowitz and Lambros (1987) and more recently Boero et al. (2008b); Rich and Tracy (2010); Clements (2012); Boero et al. (2014); Abel et al. (2015); Jurado et al. (2013); Scotti (2013), and Bloom (2014).

fessional forecasters compare with statistical benchmarks and in the latter they look at how forecast accuracy correlates with several features of the density forecasts such as the mean, variance, skewness, and kurtosis.[4] In particular, in Kenny et al. (2015), the authors conduct cross-sectional and panel regressions of the predictive accuracy of density forecasts on moments of the densities themselves. They find that density forecasts with low variance tend to perform worse. They separately consider two forecast horizons (current year and following year) and find this effect to be stronger at the shorter horizon. They consider this as evidence of over-confidence for the ECB survey of professional forecasters. In contrast, the two-step transformation used in this chapter directly compares moments of the density forecast with moments of the forecast error which allows me to directly asses the relative performance of density forecasts along these dimensions.

Finally, to make full use of the survey data, I also explicitly take into account the three-dimensional nature of survey forecasts and the correlation structure it implies by following the framework proposed in Davies and Lahiri (1995) and extended in Davies (2006).[5] The survey data provides us with, at any given survey date, forecasts by several participants at multiple forecast horizons. The framework proposed by Davies and Lahiri (1995) provides a covariance structure of forecast errors from point forecasts in light of covariates in time, forecast horizons, and individuals. Boero et al. (2008a) adopted the same framework in evaluating forecast errors using the Bank of England's Survey of External Forecasters. I use their framework to enable me to conduct joint hypothesis tests across forecasters and forecast horizons as available in the survey forecast data.

The next section outlines the empirical framework and Section 3.3 provides a description of the survey data. In Section 3.4 I report the main findings using univariate and joint tests of optimism and confidence biases. Finally, section 3.5 concludes.

[4] Their measure of predictive accuracy is the Ranked Probability Score which is the preferred measure of overall density forecast performance in the study of Boero et al. (2011).

[5] Note that the added dimensionality in the survey forecast data does not directly tranlate to the correlation structure generally accounted for in the multivariate density forecast evaluation literature such as in Diebold et al. (1999); Kalliovirta (2012) and Gonzales-Rivera and Yoldas (2012).

3.2 Empirical framework

The framework that I use implements the testing framework of Berkowitz (2001) on the quarterly density forecasts of annual real GDP growth taken from the survey of professional forecasters (SPF) conducted by the Federal Reserve Bank of Philadelphia. In the survey, respondents are asked to assign probabilities corresponding to realizations of real GDP growth for the current (and also for the following year) in one percent intervals. Beginning the first quarter of 1992, respondents are asked to assign probabilities to ten bins with the lowermost being less than negative two percent growth, followed by between negative two and negative one percent and so forth in one percent intervals with the uppermost bin being real GDP growth greater than six percent. In the rest of this section, I briefly describe the methodology of Berkowitz as applied to the survey data I use.

Denote with t the year being forecasted and with i a particular forecaster. Each target year may potentially be forecasted up to eight times in that the survey is quarterly and in each survey a respondent is asked to provide a forecast of the current and following year annual real GDP growth. Define the forecast horizon h as the number of quarters until the end of the year being forecasted (e.g. $h = 1$ for a forecast of the current year annual real GDP growth made in the fourth quarter of the year and $h = 8$ for a forecast of the following year real GDP growth made in the first quarter of the current year).

Let X_t be the real GDP growth of the U.S. in year t with a conditional distribution $f(X_t|\theta_{t,h})$ at time h quarters prior to the end of year t. where $\theta_{t,h}$ represents the underlying parameters of the density. Denote with $f(X_t|\hat{\theta}_{i,t,h})$ the density forecast of respondent i made h quarters before the end of year t.

Diebold et al. (1998) propose the construction of a probability integral transform (PIT) of the forecast target using the forecast density defined as $y_{i,t,h} = F(X_t|\hat{\theta}_{i,t,h})$. Then, if the forecast density coincides with the true density of the forecast target the PIT has a standard uniform distribution $y_{i,t,h} \sim U(0,1)$. Berkowitz (2001) propose yet another transformation where we feed the PITs into the inverse normal cdf $z_{i,t,h} = \Phi^{-1}(y_{i,t,h}, 0, 1)$ and the resulting random variable would then have the standard normal distribution under the same presumption.

Under the special case that both $f(X_t|\theta_{t,h})$ and $f(X_t|\hat{\theta}_{i,t,h})$ are Normally distributed then the two-step transformation reduces to a standardized forecast error, the difference between the actual realization of real GDP growth and mean of the forecast divided by the standard deviation of the density forecast. Further it is straightforward to show that the distribution of this transformed random variable is given by:

$$z_{i,t,h} \equiv \frac{X_t - \hat{\mu}_{i,t,h}}{\hat{\sigma}_{i,t,h}} \tag{3.1}$$

$$z_{i,t,h} \sim \mathbf{N}\left(\frac{\mu_{t,h} - \hat{\mu}_{i,t,h}}{\hat{\sigma}_{i,t,h}}, \frac{\sigma^2_{t,h}}{\hat{\sigma}^2_{i,t,h}}\right) \tag{3.2}$$

That is, the mean and variance of the transformed random variable maps to the standardized difference in the predicted-against-actual means (optimism) and the ratio of the actual and predicted variances (confidence). Under the hypothesis that forecast densities are unbiased and accurately capture the conditional distribution of real GDP growth, then $z_{i,t,h}$ is an independent and identically distributed standard Normal random variable. The literature has so far not exploited this result into separate tests of optimism and confidence.

Define $\alpha_{i,h}$ as the conditional mean of the standardized forecast error for each forecaster-forecast horizon pair. Second, define $\beta_{i,h}$ as the conditional standard deviation of the same variable. Further, as in Berkowitz (2001) I include a persistence parameter ($\rho_{i,h}$) intended to capture dependencies in the forecast errors and re-express equation 3.2 as:

$$z_{i,t,h} - \alpha_{i,h} = \rho_{i,h}(z_{i,t-1,h} - \alpha_{i,h}) + \beta_{i,h}\epsilon_{i,t,h} \tag{3.3}$$

where

$$\bar{\epsilon} \sim \mathbf{MVN}(\mathbf{0}, \mathbf{P}) \tag{3.4}$$

The parameter $\alpha_{i,h}$ represents the optimism bias with a positive value indicating a pessimistic forecast (i.e. the realization of annual real GDP growth is larger than the mean of the forecast density) and a negative value for overly-optimistic forecasts. The scale of this parameter is in terms of standard deviations from the forecast density. On the other hand, the $\beta_{i,h}$ parameter represents the over-confidence bias. A value of $\beta_{i,h}$ greater than one implies that the forecast error is $\beta_{i,h}$ times larger than implied standard deviation of the survey forecast which we define as over-confidence. Alternatively, when $\beta_{i,h}$ is less than one then the implied standard deviation of the density forecast is larger than the forecast error and would suggest doubt or under-confidence in the forecasters' forecast.

To implement the Berkowitz (2001) procedure to the survey data, I treat each (i, h) pair as a univariate series. For each series, the parameters may then be estimated by Ordinary Least Squares and I perform likelihood ratio tests on whether $\rho = 0$ indicating independence, $\alpha = 0$ for no optimism bias and $\beta = 1$ for no confidence bias.[6]

However, this exercise ignores the three dimensional nature of the survey data and in a second exercise I jointly test the hypothesis of biases across forecasters and forecast horizons by implementing the framework proposed by Davies and Lahiri (1995) and allow the correlation across forecasters and forecast horizons to be non-zero for the same forecast target and up to one year lag. In the modified framework, the parameters of the correlation matrix map to N individual-level variance estimates, a quarterly shock variance for real GDP growth, and an i.i.d. forecast error variance.

Their framework yields a block diagonal correlation matrix. That is if we define the vector z_t as the vector of standardized forecast errors of the same forecast target year across all forecast horizons and forecasters, and ϵ_t the corresponding vector of residuals from equation 3.3 we have that:

$$z = \begin{bmatrix} z_1 \\ \vdots \\ z_t \\ \vdots \\ z_T \end{bmatrix} \quad , \quad z_t = \begin{bmatrix} z_{1,1} \\ \vdots z_{1,N} \\ \vdots \\ z_{h,n} \\ \vdots \\ z_{H,N} \end{bmatrix}$$

$$\mathbb{E}\left[\begin{bmatrix} \epsilon_t \\ \epsilon_{t+s} \end{bmatrix} \begin{bmatrix} \epsilon_t & \epsilon_{t+s} \end{bmatrix}\right] = \begin{cases} \begin{bmatrix} C & D' \\ D & C \end{bmatrix} & \text{if } s == 1 \\ \begin{bmatrix} C & \mathbf{0} \\ \mathbf{0} & C \end{bmatrix} & \text{otherwise} \end{cases}$$

Davies and Lahiri (1995) maps the matrices C and D into a set of parameters re-

[6]In Berkowitz (2001) and subsequent contributions to the literature, joint tests on $\alpha = 0$ and $\beta = 1$ are typically conducted.

flecting individual predictive ability, the variance of unpredictable shocks to real GDP growth, and the forecast horizon.[7]

Conditional on estimates of the matrix P, I then verify the univariate regression estimates by re-estimating the parameter vectors for ρ, α, and β using maximum likelihood across all observations. As in the first exercise with univariate regressions, I then conduct likelihood ratio tests to check for independence, optimism, and confidence biases.

3.3 Data description

The forecasting data are taken from the Survey of Professional Forecasters managed by the Federal Reserve Bank of Philadelphia (SPF).[8] In particular, I focus on survey responses on forecasts of current and following year annual average Real GDP growth for the surveys beginning the first quarter of 1992 and ending the fourth quarter of 2013. The advantage of using this forecast is that respondents provide probabilities on certain ranges (bins) of Real GDP growth from less than negative two percent to more than six percent.[9] I also limit the survey respondents to those who have participated in at least fifteen consecutive surveys and have participated in at least half of the surveys over the period.[10] This leaves me with 2,863 forecasts from 31 forecasters over 88 quarters (or 22 forecast targets) and at eight different forecast horizons.

The following figure plots a typical response to which I fit a Normal distribution. The gray bars represent the probabilities that the respondent has assigned to each range of possible outcomes and the blue line represents the fitted density of a Normal distribution. The mean and variance of the Normal distribution are chosen so as to minimize the squared area in between the two densities where the areas are categorized into intervals as in the survey question and weighted with the probability given in the response.[11] The

[7] See the appendix for estimation details.

[8] The survey forecast and real-time GDP growth data are available at the Real-Time Data Research Center of the Federal Reserve Bank of Philadelphia. `http://www.philadelphiafed.org`

[9] These are at one percent intervals for a total of ten bins. Starting from the second quarter of 2009, an eleventh bin was added splitting the lowest bin to between negative two to negative 2.9 percent and less than three percent.

[10] This still results in an unbalanced panel of forecast data. However, these restrictions imply that each of the remaining forecasters in the sample would have made forecast for at least two consecutive target years across all forecast horizons.

[11] Giordani and Soderlind (2003) use the same approach in their analysis of inflation forecast uncertainty in the SPF.

red dot represents the actual realization of annual real GDP growth.[12]

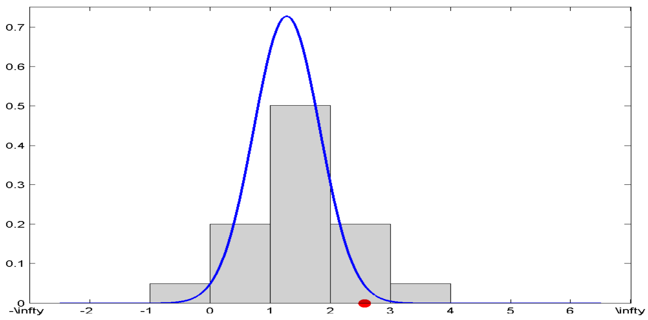

Given the implied mean and standard deviations from the survey forecast densities and the realized annual real GDP growth, I then construct standardized forecast errors as described in the previous section. The following table reports summary statistics for the sample. The sample third and fourth moments would suggest that the Normality

Table 3.1: Standardized Forecast Error Statistics

	Mean	**Var**	**Skew**	**Kurt**	**Obs**
Std FE	-0.443	7.456	-1.376	10.831	2863.000

assumption may not be appropriate for the survey density forecast. In the appendix, I provide robustness check where I compare estimates of the mean and variance of the survey responses when using a Normal distribution as against a location-scale transformed skewed Students' t distribution. I find no significant difference in the mean estimate although I may be underestimating the variance by about 30 percent under the Normality assumption. Nevertheless, closer inspection of the data reveals that this is largely true only for a subset of forecasters and forecast targets.[13]

Further, note that the survey data may suffer from a truncation bias in that the possible ranges for real GDP growth at the lower end is trimmed to less than negative two percent until the survey on the first quarter of 2009 and thereafter reduced to less than three percent. Around this period, a substantial number of forecasters have placed a

[12]The actual realizations are taken as of one quarter after the end of the year being forecasted except for 1995 real GDP growth in which case the real-time value 5-quarters afterwards is used. See the appendix for a comparison of forecast errors across different vintages of real-time data.

[13]In particular, only three forecasters in the sample exhibited significantly large skewness estimates and across responses, only the surveys 2006:Q3 to 2008:Q1 exhibited significantly large (and negative) skewness.

significant amount of probabilities on this lowest bin. This suggests that around this period, had there been more categories below this threshold, survey respondents may have distributed this probability in the lowest bin to more categories. Hence, we may be underestimating the mean and variance of survey forecasts around this period and may be introducing artificial skewness in the data.

3.4 Tests for optimism and confidence

3.4.1 Summary statistics

I start with a simple exercise in which I estimate the optimism and confidence bias parameters in terms of sample averages of means and standard deviations on subsets of the data.[14] First, I report the estimated mean and standard deviation of standardized forecast errors across forecast horizons in the following table. The results indicate no

Table 3.2: Estimated biases across forecast horizon

	H1	H2	H3	H4	H5	H6	H7	H8	All
alpha	0.064	0.233	-0.074	0.006	-0.335	-0.646	-0.774	-0.694	-0.254
t-stat	0.035	0.198	-0.049	0.003	-0.081	-0.182	-0.218	-0.177	-0.088
beta	1.837	1.177	1.518	1.940	4.141	3.551	3.548	3.923	2.880
χ^2-stat	1261.759	515.514	850.660	1362.935	5676.233	4123.537	4091.525	4815.863	23062.353

bias in optimism overall and forecast variances to be about 80 percent to four times smaller than the true conditional variances of the forecast target or about three times smaller overall.[15] This bias appears to increase in the forecast horizon as well in contrast with the result in Kenny et al. (2015) using the ECB survey data.

I then conduct a second exercise where I estimate the same parameters for each forecaster in the sample. In the following table I report the number of individual-level parameter estimates (over 31) deviating from $\alpha = 0$ and $\beta = 1$. The table suggests that none of the forecasters appear to be systematically pessimistic (optimistic). However, and consistent with the earlier result, all of the forecasters appear to be overconfident.

As a third preliminary exercise I estimate the optimism and confidence bias parameters using a rolling sample over the forecast target where I consider a 5-year rolling

[14]These sample means and variances are potentially biased in that they do not take into account the correlation structure across forecast errors.

[15]Note that, as shown in the appendix, estimates of the variances using a Normal distribution for the survey forecasts could be about 30 percent underestimated relative to estimates under a more general location-scale transformed skewed Student's t distribution.

Table 3.3: Individual bias tests

	90	**95**	**99**
$\alpha > 0$	0	0	0
$\beta > 1$	31	31	31
$\alpha < 0$	0	0	0
$\beta < 1$	0	0	0

1 sided t and χ^2 tests with column headers indicating confidence levels

window of forecasts (e.g. forecasts for 1992 to 1996 real GDP). The following figure plots the parameter estimates with a 90 percent confidence intervals.[16]

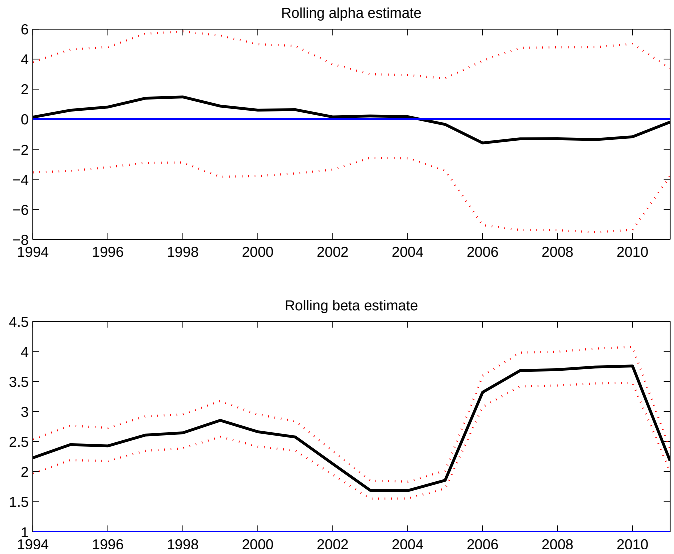

Figure 3.1: 5-year rolling window estimates

Consistent with the findings of Giordani and Soderlind (2006) on SPF GDP forecast data between 1982 to 2003, there seems to be excessive pessimism (an $\alpha > 0$ in the upper panel) in rolling estimates of the optimism bias from 1993 to 2002 using our sample as well although not significantly different from zero. Thereafter, it seems that the average forecast is overly optimistic but still not statistically significantly different

[16]The confidence intervals are constructed from the t and chi-squared distributions respectively with degrees of freedom equal to the number of observations less one.

from zero. With regard to the overconfidence bias, estimates suggest forecasts are consistently overconfident throughout the sample which seems to be more so around recessions.[17] Interestingly, over-confidence appears to be more pronounced in the 2006-2010 period.

3.4.2 Univariate series tests

Now I estimate the set of parameters for each pair of forecaster and forecast horizon treating each as a univariate series. That is I run a regression on the specification given by equation 3.3 for each individual and forecast horizon pair. The next figure plots estimates of the persistence parameter $\rho_{i,h}$ with the values on the vertical axis and the estimates sorted according to the forecast horizon on the horizontal axis. The circles represent estimates for each forecaster and the solid black line represents the average of the optimism biases for each forecast horizon.

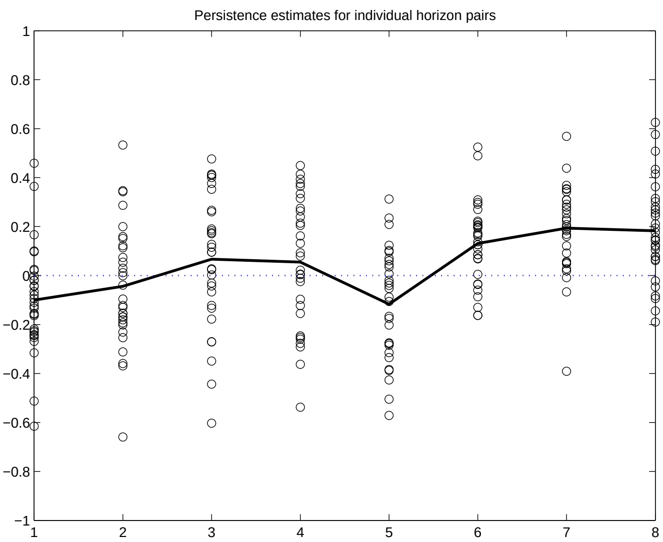

Figure 3.2: Individual and average persistence biases

[17]Estimates of overconfidence by Giordani and Soderlind (2006) which uses the confidence interval coverage test by Christoffersen (1998) similarly suggests overconfidence for the SPF forecasts in between 1982 to 2003.

Only 42 of the 248 likelihood ratio tests reject the null hypothesis of zero persistence at the ten percent significance level with 25 of these at forecast horizons 5 quarters ahead and larger (i.e. following year forecasts). For ten of the forecasters, I could not reject the null hypothesis for any forecast horizon. Similarly Figure 3.3 plots point estimates of the optimism bias.

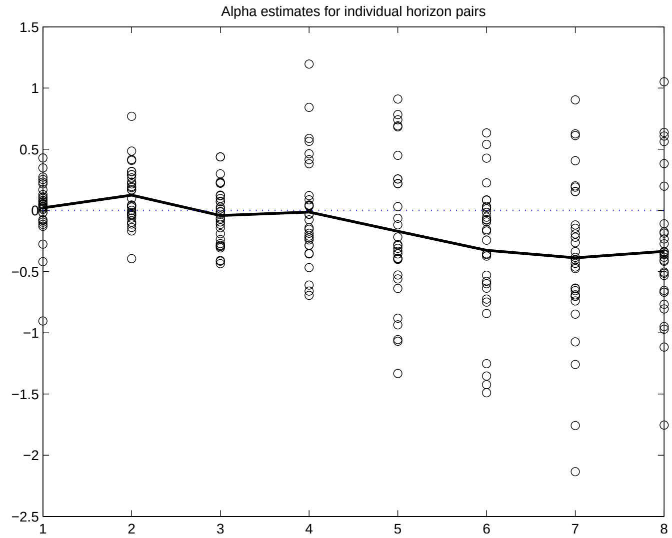

Figure 3.3: Individual and average optimism biases

The estimates confirm the earlier results that the optimism bias appears to be increasing over the forecast horizon although we also note that the α estimates for longer horizons are also more dispersed. Likelihood ratio tests on whether the optimism bias is zero is rejected for 55 estimates at the ten percent significance level. We failed to reject the null hypothesis for eight of the forecasters across any forecast horizon. Finally, 27 of the 55 estimates for which I reject no optimism bias are at the forecast horizons 5 quarters and higher.

In the next figure I plot estimates of the confidence bias. Here I reject the null hypothesis of no confidence bias for 161 out of 248 estimates. There were more or less an equal amount of rejections across forecast horizons and I am able to reject the null hypothesis for at least two forecast horizons for all forecasters with one forecaster having a β estimate significantly different from one for all eight forecast horizons.

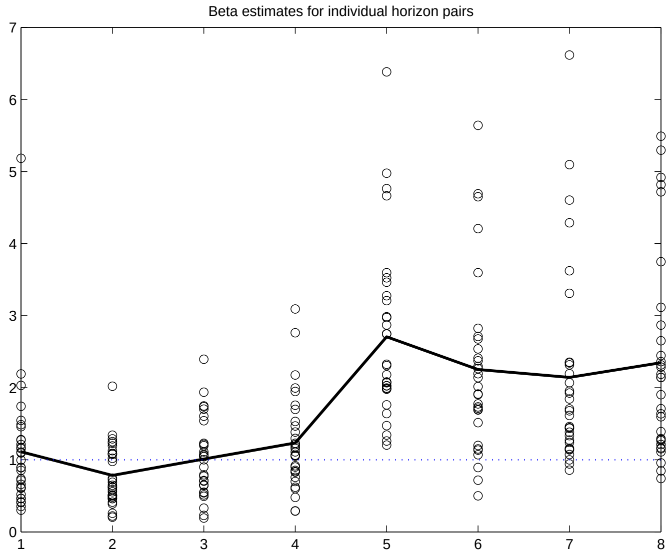

Figure 3.4: Individual and average confidence biases

3.4.3 Joint hypothesis tests

In this section, I jointly test the null hypothesis of zero biases across estimates for forecast horizons or individuals. I first verify that the joint maximum likelihood estimate of the parameters conditional on an estimate of the correlation matrix P are equivalent to the univariate estimates. To perform joint tests on the persistence parameters ρ and the optimism bias parameters α I conduct Likelihood ratio tests using a multivariate normal likelihood where the covariance matrix is constructed from the confidence bias parameters β and an estimate of the correlation matrix P.

As I noted earlier, I construct the correlation matrix P using a modified version of the Davies and Lahiri (1995) framework.[18] First, when I re-estimated the parameters using the multivariate normal likelihood, all but one set of parameter estimates (out of 248 forecaster-forecast horizon pairs) were unchanged. Second, when I conducted joint tests of zero persistence, zero optimism bias, and no confidence bias across forecasters and forecast horizons, I failed to reject the null hypothesis in all cases and parameters.

[18]Davies and Lahiri (1995) estimates the covariance matrix as a whole whereas I decompose the covariance matrix into a correlation matrix and estimates of the variances $\beta_{i,h}$. See the appendix for details.

Although I could reject individual parameter estimates for some horizons (forecasters) for a given forecaster (horizon) in the previous exercise, the joint tests fail to reject that all optimism or confidence biases are zero for a given forecaster or forecast horizon.

In light of the test results in the univariate cases, the failure to reject any of the hypotheses in the joint tests merit further investigation. This may be due to the small sample size (we have $T = 22$ and $N * H = 248$). This is left for future work where alternative frameworks suitable for multi-dimensional panel data where the time dimension is considerably smaller than the others would also be considered.

3.5 Conclusion

Timely and reliable forecasts of future economic conditions are a crucial input to economic agents and policy makers alike. Surveys of professional forecasts provide such information and in this regard it is also quite important to recognize the limitations of such forecasts. Density forecasts in these surveys provide users with both mean and variance predictions of the future realizations of macroeconomic variables such as real GDP growth. It is not surprising then that such density forecasts have gained considerable attention given the added information they contain over point forecasts. It is then important to estimate potential biases in both the mean and variance of survey forecasts. This chapter uses existing methodologies in the literature and decomposes the proposed tests therein to evaluate the predictive accuracy of survey density forecasts along these two dimensions - what I refer to as optimism and confidence biases. I provide estimates of these biases across forecasters and forecast horizons using the Survey of Professional Forecasters responses to annual real GDP growth forecasts in the U.S. for the years 1992-2013.

I find some evidence for optimism biases for about 22 percent (55 out of 248 tests) of the forecaster-forecast horizon pairs and a substantially larger proportion of forecaster-forecast horizon pairs to have an over-confidence bias at 65 percent of the tests. Further, the sample variance estimates suggest that this bias seems to increase with the forecast horizon. However, I fail to reject the null hypothesis of no biases in the joint tests across forecasters or forecast horizons for both optimism and confidence. Given the large dimensionality of the number of forecasters and forecast horizons relative to forecast targets, future work in this direction could focus on improving estimates of the correlation structure across forecast errors.

Nevertheless, the results in the chapter provide suggestive evidence of overconfidence in survey forecasts which could potentially be time-varying. This corroborates the evidence presented in Kenny et al. (2014) and Kenny et al. (2015) on the ECB's survey of professional forecasters and provide suggestive evidence in support of the theoretical literature on behavioral biases as key ingredients to booms and crashes (e.g. Gennaioli and Shleifer, 2010; Gennaioli et al., 2013).

Bibliography

Abel, J., Rich, R., Song, J., and Tracy, J. (2015). The measurement and behaviour of uncertainty: evidence from the ecb survey of professional forecasters. *Journal of Applied Econometrics*, forhcoming.

Berkowitz, J. (2001). Testing density forecasts, with applications to risk management. *Journal of Business and Economic Statistics*, 19 (4):465–474.

Bloom, N. (2014). Fluctuations in uncertainty. *Journal of Economic Perspectives*, 28 (2):153–176.

Boero, G., Smith, J., and Wallis, K. (2008a). Evaluating a three-dimensional panel of point forecasts: The bank of england survey of external forecasters. *International Journal of Forecasting*, 24 (3):354–367.

Boero, G., Smith, J., and Wallis, K. (2008b). Uncertainty and disagreement in economic prediction: the bank of england survey of external forecasters. *The Economic Journal*, 118:1107–1127.

Boero, G., Smith, J., and Wallis, K. (2011). Scoring rules and survey density forecasts. *International Journal of Forecasting*, 27 (2).

Boero, G., Smith, J., and Wallis, K. (2014). The measurement and characteristics of professional forecasters' uncertainty. *Journal of Applied Econometrics*, forthcoming.

Christoffersen, P. (1998). Evaluating interval forecasts. *International Economic Review*, 39 (4):841–862.

Clements, M. (2012). Subjective and ex post forecast uncertainty: Us inflation and output growth. *Warwick Economic Research Papers*.

Coibon, O. and Gorodnichenko, Y. (2012). What can survey forecasts tell us about information rigidities? *Journal of Political Economy*, 120 (1):116–159.

Coibon, O. and Gorodnichenko, Y. (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review (forthcoming)*.

Corradi, V. and Swanson, N. (2006). *Handbook of Economic Forecasting*, volume 1, chapter Predictive Density Evaluation, pages 197–284. Elsevier.

Davies, A. (2006). A framework for decomposing shocks and measuring volatilities derived from multi-dimensional panel data of survey forecasts. *International Journal of Forecasting.*, 22:373–393.

Davies, A. and Lahiri, K. (1995). A new framework for analyzing survey forecasts using three-dimensional panel data. *Journal of Econometrics*, 68:205–227.

Diebold, F., Gunther, T., and Tay, A. (1998). Evaluating density forecasts. *International Economic Review*, 39:863–883.

Diebold, F., Hahn, J., and Tay, A. (1999). Real-time multivariate density forecast evaluation and calibration: monitoring the risk of high frequency returns on foreign exchange. *Review of Economics and Statistics*, 81:661–673.

Gennaioli, N. and Shleifer, A. (2010). What comes to mind. *Quarterly Journal of Economics*, 125 (4):1399–1433.

Gennaioli, N., Shleifer, A., and Vishny, R. (2013). A model of shadow banking. *Journal of Finance*, 68 (4):1331–1363.

Giordani, P. and Soderlind, P. (2003). Inflation forecast uncertainty. *European Economic Review*, 47:1037–1059.

Giordani, P. and Soderlind, P. (2006). Is there evidence of pessimism and doubt in subjective distributions? implications for the equity premium puzzle. *Journal of Economic Dynamics & Control*, 30:1027–1043.

Gonzales-Rivera, G. and Yoldas, E. (2012). Autocontour-based evaluation of multivariate predictive densities. *International Journal of Forecasting*, 28:328–342.

Jurado, K., Ludvigson, S., and Ng, S. (2013). Measuring uncertainty. *NBER Working Paper*, 19456.

Kalliovirta, L. (2012). Misspecification tests based on quantile residuals. *Econometrics Journal*, 15:358–393.

Kenny, G., Kostka, T., and Masera, F. (2014). How informative are the subjective density forecasts of macroeconomists? *Journal of Forecasting*, 33:163–185.

Kenny, G., Kostka, T., and Masera, F. (2015). Density characteristics and density forecast performance: a panel analysis. *Empirical Economics*, 48:1203–1231.

Mincer, J. and Zarnowitz, V. (1969). *Economic Forecasts and Expectations.*, chapter The evaluation of economic forecasts. National Bureau of Economic Research, New York.

Mitchell, J. and Wallis, K. (2011). Evaluating density forecasts: forecast combinations, model mixtures, calibration and sharpness. *Journal of Applied Econometrics.*, 26:1023–1040.

Patton, A. and Timmerman, A. (2010). Why do forecasters disagree? lessons from the term structure of cross-sectional dispersion. *Journal of Monetary Economics*, 57 (7):803–820.

Patton, A. and Timmerman, A. (2012). Forecast rationality tests based on multi-horizon bounds. *Journal of Business and Economic Statistics*, 30 (1):1–17.

Pesaran, M. H. and Weale, M. (2006). *Handbook of Economic Forecasting*, volume 1, chapter Survey Expectations, pages 715–776. Elsevier.

Rich, R. and Tracy, J. (2010). The relationships among expected inflation, disagreement, and uncertainty: Evidence from matched point and density forecasts. *The Review of Economics and Statistics*, 92 (1):200–207.

Rossi, B. (2014). Density forecasts in economics, forecasting, and policymaking. *Els Opuscles del CREI.*

Scotti, C. (2013). Surprise and uncertainty indexes: Real-time aggregation of real-activity macro surprises. *FRB International Finance Discussion Paper*, 1093.

Zarnowitz, V. and Lambros, L. (1987). Consensus and uncertainty in economic prediction. *Journal of Political Economy*, 95 (3):591–621.

3.A The SPF survey

See the Chapter 1 Appendix on Macro-uncertainty for a more detailed description of the data. For this chapter, I restrict the sample to include only those forecasters who have made at least fifteen consecutive surveys and at least half of all the surveys in the sample period. This leaves me with 2,863 forecasts of annual real GDP growth made by 31 forecasters over 8 quarterly forecast horizons and 22 forecast targets (annual real GDP growth from 1992 to 2013).

To estimate the mean and variance of each forecast, I fit an (unbounded) Normal distribution by minimizing the squared difference between the cumulative probability associated with each bin by a Normal distribution against the probability given in the survey response. Each of these squared differences are then weighted by the probabilities in the survey response and the mean and variance are chosen so as to minimize the sum of these weighted squared differences. That is:

$$\{\hat{\mu}, \hat{\sigma}\} = \arg\min_{\mu,\sigma} \sum_{k=1}^{K} p_k \Big(\int_{k_l ow}^{k_u p} \phi(x|\mu, \sigma) - p_k \Big)^2$$

where p_k is the probability assigned to bin k in the survey with lower and upper edges $k_l ow$ and $k_u p$ and $\phi(x|\mu, \sigma)$ is the Normal pdf.[19]

3.B Choice of forecast target vintage

To construct forecast errors I use the real time data set on quarterly real GDP growth levels from the Federal Reserve Bank of Philadelphia. In particular, I construct four vintages of the data releases, one quarter after the year of the forecast (e.g. the real GDP growth rate for 2008 as of the first quarter of 2009), five quarters after, nine quarters after, and the 'final' release (the real GDP growth rate for all years as of the first quarter of 2015).

First, I look at the mean squared standardized forecast errors constructed from various vintages of the outcome variable to see which vintage of real GDP growth data the forecasters are most closely able to forecast. The following table reports the overall average squared forecast error by vintage data release. It seems that the survey forecasts

[19]This is similar to minimizing the Kullback-Leibler divergence between the two densities.

Table 3.4: Average MSSE by data release vintage

	1Q	5Q	9Q	FN
MSSE	6.9906	9.1434	11.2838	12.7173
Obs	2,700	2,781	2,621	2,453

perform best relative to initial releases of the data when compared to data revisions in succeeding years. I also check whether average forecast errors across vintages differ over forecast horizons. The following figure plots average forecast errors across data release vintages and forecast horizons. Again, it seems that forecast accuracy is consistently better for initial release data across all horizons.

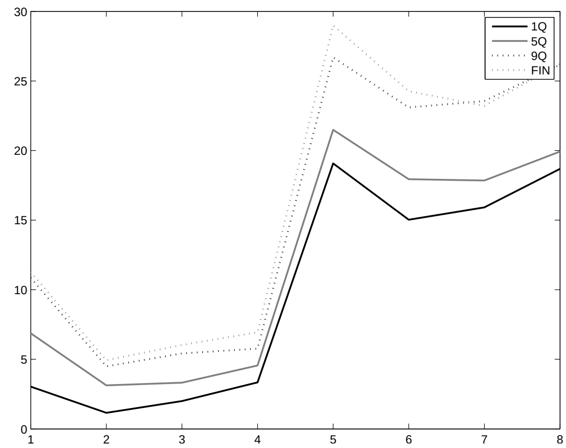

Figure 3.5: Average MSE across forecast horizons

A forecast horizon of one is associated with a forecast of current year annual real GDP growth made in the fourth quarter whereas horizon 8 is for a forecast of following year annual real GDP growth made in the first quarter.

In the next figure, I plot the average mean-squared standardized forecast error by forecast target from 1993 to 2013 annual real GDP growth.

However, to consistently estimate the correlation structure in the forecast errors, missing observations in the initial release data prove to be problematic given that the real time dataset has missing observations to construct the one-quarter afterwards 1995 annual real GDP growth. To resolve this issue, I use forecast errors constructed from one-quarter afterwards for all of the series except for the 1995 forecasts for which I use the 5-quarter afterwards release. I have also run all the exercises using only 5-quarter afterwards realized real GDP with similar results.

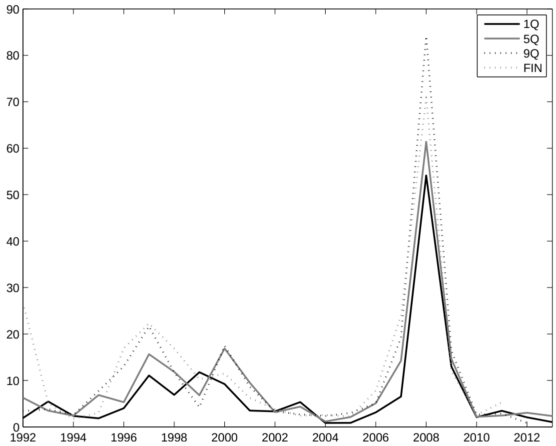

Figure 3.6: Average MSE across forecast targets

3.C Robustness of mean and variance estimates

The implied mean and variance of the survey forecasts are derived by finding the mean and variance of a Normal distribution that minimizes the weighted squared integral of the area between the two densities. First, I check whether the assumption of Normality affects the estimates of the mean and variance from the survey responses. To do so, for a sub-sample of the survey responses I estimate a four-parameter density corresponding to a location-scale transformed skewed Student's t distribution. I then check whether estimates of the mean and variance are preserved when I fit this four parameter distribution to the survey responses taking into account the possibility of skewness and excess kurtosis being different from zero. Note that in order to estimate four parameters I would also need the survey responses to have non-zero probabilities to at least four bins hence the procedure may only be done on a subset of the data.

Second, I also construct the two-step transformation as proposed in Berkowitz (2001) directly. That is, first I construct Probability Integral Transforms defined as $y = F_{i,t,h}(x_t)$ where F is the forecast by forecaster i for real GDP growth in year t at h quarters ahead and x_t is real GDP growth for year t as recorded in the first quarter of the following year. I then generate $z_{i,t,h}$ by feeding the PIT into the inverse cdf of a standard Normal distribution $z_{i,t,h} = \Phi^{-1}(y_{i,t,h})$. Under the null hypothesis the transformation still yields a zero-mean and unit-variance Normally distributed random variable. Differences in the location of the survey forecast and the true conditional density of real GDP growth will be reflected in the mean of the transformed variable and similarly, differences in the respective second moments reflect differences in the scale of the distributions.

However, since we do not assume Normality in the forecast distributions, we lose the closed form interpretation of the mean being the difference in terms standard deviations and the variance being the ratio of the two variances. Further, we also lose some information in the survey responses in the cases when the forecast target realization is in a survey response bin with zero probability (e.g. real GDP growth turned out to be 5 percent but the survey forecast only placed non-zero probabilities between two and four percent.). Direct calculation of the two-step transformation will yield negative or positive infinity (the zeroth and 100th percentile of a Normal random variable respectively). In order to facilitate the analysis, I truncate these values to the first and 99th percentiles of a standard Normal distribution instead. Of course, this procedure does not differentiate from a forecast realization one bin away from the non-zero probabilities in a forecast as against another forecast where the realization is two or three bins away.

The aforementioned truncation turns out to be a significant issue as 160 forecasted distributions were to the right of the forecast realization (a zeroth percentile transformation with real GDP growth much lower than predicted) and 800 forecast distributions were to the left. Given that over a third of the transformations had to be truncated, I expect that differences in the estimates using the standardized forecast error and the direct two-step transformation would largely be due to how these occurrences are handled.

3.C.1 Comparison of estimates against four-parameter Student's t

In the following table, I report the frequency of responses containing one to all bins with non-zero probabilities.

Table 3.5: Number of bins with non-zero responses

	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**All**	**Total**
Count	71	556	946	442	365	188	155	99	51	69	2942
Prop	2	19	32	15	12	6	5	3	2	2	100

All may be 10 or 11 bins

As the table indicates, this exercise is only feasible for about half of the responses. For these I estimate the location scale, degrees of freedom and skewness of a location-scale transformed skewed t distribution.[20]. In the following table, I report the average difference in estimated means under Normality and the location-scale transformed skewed Students' t distribution as well the ratio of the variance estimates along with other sample statistics. There does not seem to be a significant difference in the mean

Table 3.6: Comparison of mean and variance estimates

	Mean	**StDev**	**5th-tile**	**95th-tile**
Diff in mean	-0.059	0.754	-0.071	0.248
Ratio of variance	1.293	1.820	1.052	1.391

The first row reports the estimated mean from the location-scale transformed skewed Student's t less the estimated mean from a Normal density. The second row reports the ratio between the estimated variances using the same densities.

estimates across the two distributions. However, assuming Normality seems to imply smaller variances as can be seen in the second row in the table above. Even though the median ratio between the two variance estimates is one, the average seems to indicate that assuming Normality yields variance estimates 30 percent lower than the location-scale transformed skewed Student's t distribution. In particular, only ten percent of the estimates yielded a larger variance estimate under Normality and for about 40 percent of the forecasts, the variance is at least ten percent smaller.

In the next table I report statistics on the estimated skewness and kurtosis of the location-scale transformed skewed Students' t distribution.

Table 3.7: Estimated Skewness and degrees of freedom

	Mean	**5th-tile**	**10th-tile**	**25th-tile**	**50th-tile**	**75th-tile**	**90th-tile**	**95th-tile**
Skewness	-63.197	-1.000	-1.000	-1.000	-0.255	0.000	0.902	1.000
df	299.022	299.000	299.000	299.000	299.000	299.001	299.002	299.002

The starting value for the degrees of freedom parameter was set to 299. Note that the degrees of freedom parameter may not be reliably estimated given that the survey responses are not very informative about the tail probabilities.

The results suggest that it is quite difficult to reliably estimate excess kurtosis from the survey responses. Nevertheless, it seems that skewness is an issue and ignoring this appears to affect the variance estimates. However, the unusually high skewness seems to be driven by only several forecasters. Only five forecasters have an average skewness larger than 1 and of these only three remain in the sample.

[20] I use a modified version of Andrew Patton's code available at `http://public.econ.duke.edu/~ap172/code.html`

3.C.2 Two-step transformation without Normality

As earlier described, I construct the two-step transformation by first feeding the forecast target realization into the survey forecast to generate the PIT,

$$y_{i,t,h} = \sum_{k=1}^{K(x_t)} p_{i,t,h}$$

where $p_{i,t,h}$ are the probabilities for each bin and $K(x_t)$ is the bin (in ascending order) that contains the realization of the forecast target. I then construct our variable of interest by using the inverse-Normal transform on the computed PIT $z_{i,t,h} = \Phi^{-1}(y_{i,t,h})$. As earlier mentioned, I replace PITs of zero and one with 0.01 and 0.99 to avoid having infinite values in the sample.

I first report sample means and standard deviations across forecasters and forecast horizons. As in the main text, a sample mean greater than zero suggests that the forecast mean is lower than realization, and a standard deviation greater than one implies that the forecast variance is also lower than the variance of the forecast target. In the next figures, I first plot the estimated mean and standard deviation across forecasters and forecast horizons. Each circle represents a sample mean or standard deviation for a forecaster-forecast horizon pair. The solid line represents the average for a given forecast horizon across forecasters.

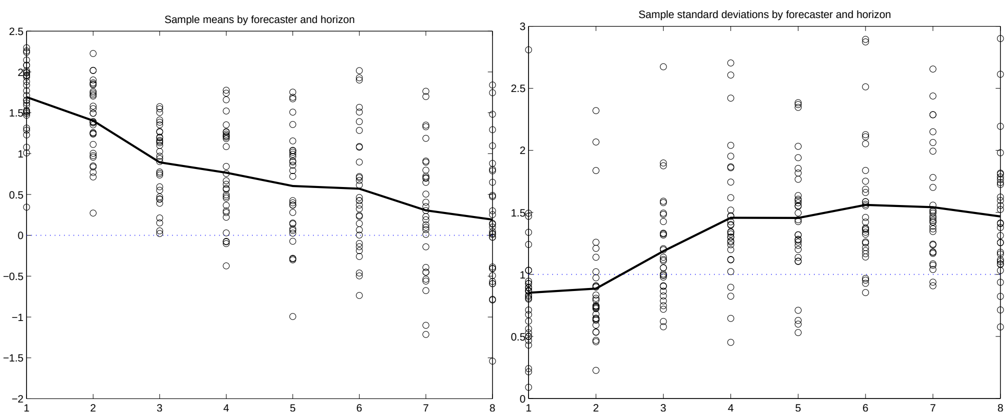

(a) Sample means (b) Sample standard deviations

Figure 3.7: Sample means and standard deviations across forecaster-horizon pairs

Although a similar pattern across forecast horizons in the sample means appear in the figure above relative to that with the standardized forecast errors, the sample means are now mostly above zero which suggests pessimism in the forecasts across forecast horizons. On the other hand, the sample standard deviation estimates are qualitatively the same. I then estimate sample means and standard deviations using a 5-year rolling sample of survey forecast responses. Again, the results suggests qualitatively similar patterns over time, albeit the sample means appear to be consistently pessimistic until the rolling window estimates centered around the 2009 forecasts.

(a) Sample means

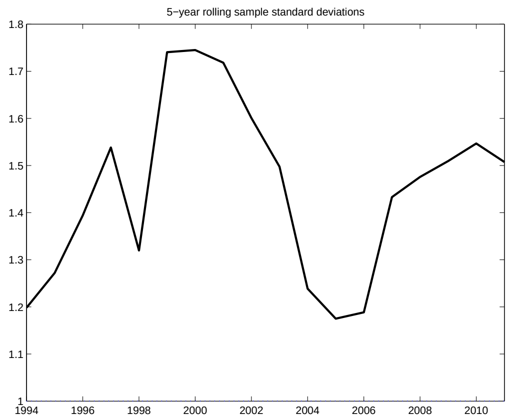

(b) Sample standard deviations

Figure 3.8: Sample means and standard deviations over 5-year rolling windows

I then proceed with estimating the persistence, optimism, and confidence bias parameters for each forecaster-forecast horizon pair as was done in the main text. The following table reports the number of likelihood ratio tests for which I reject the null hypothesis for each of the parameters.

Table 3.8: Individual bias likelihood ratio tests

	All (248)	Individual (31)	H1 (31)	H2 (31)	H3 (31)	H4 (31)	H5 (31)	H6 (31)	H7 (31)	H8 (31)
$\rho = 0$	64	3	9	13	5	5	6	6	9	11
$\alpha = 0$	121	9	23	26	17	12	12	11	12	8
$\beta = 1$	131	14	14	23	14	19	12	15	20	14

Rejection is at the ten percent significance level. The column labels report the relevant category over which the tests belong. The value in parentheses report the number of tests for each category. The Individual column reports the number of individuals for which more than half of the tests across forecast horizons reject the null.

The results are qualitatively similar although the number of rejections are fewer than when I use the standardized forecast error.

3.D Structural estimation of the correlation matrix

To obtain estimates of the covariance matrix of the parameters α and β we have to take into account the correlation structure across the forecast errors over time, forecasters, and forecast horizons. To do this, I provide an estimate of the correlation across standardized forecast errors using a modified version of the covariance framework proposed by Davies and Lahiri (1995).

To illustrate the nature of the correlation structure across forecast errors, first I specify a simple model for the data generating process behind the forecast target. For ease of notation, define a period as one quarter and let x_t be quarter-on-quarter real GDP growth while X_t is annual real GDP growth as computed if the year ended in quarter t. Assume that quarterly real GDP growth is a persistent process, then we have:

$$\begin{aligned}
x_t &\equiv log(GDP_t) - log(GDP_{t-1}) \\
x_t &= \rho x_{t-1} + \epsilon_t \\
\epsilon_t &\sim i.i.d.\mathbf{N}(0, \sigma_\epsilon^2) \\
X_t &\equiv \frac{\sum_{\tau=0}^{3} x_{t-\tau}}{\sum_{\tau=4}^{7} x_{t-\tau}} - 1
\end{aligned}$$

To facilitate the construction of forecasts of X_t, we expand the representation of quarterly growth to include six lags of x_t:

$$\bar{x}_t \equiv \begin{bmatrix} x_t \\ x_{t-1} \\ \vdots \\ x_{t-6} \end{bmatrix} = F\bar{x}_{t-1} + C\epsilon_t \tag{3.5}$$

where $F = \begin{bmatrix} \rho & \mathbf{0}_{1x6} \\ \mathbf{I}_6 & \mathbf{0}_{6x1} \end{bmatrix}$, and $C = \begin{bmatrix} 1 \\ \mathbf{0}_{6x1} \end{bmatrix}$. Patton and Timmerman (2010) show that the annual real GDP growth rate may be approximated by a linear combination of the current and up to six lags of the quarter-on-quarter growth rate $X_t = \bar{w}'\bar{x}_t$ where the ith element of $\bar{w}$ is given by $w_i = \frac{4-|4-i|}{4}$. Combining this with equation 3.5 we then have that the conditional distribution of X_t at time $t-h$ is given by:

$$X_{t|t-h} \sim \mathbf{N}\left(\mu_{t,h}, \sigma^2_{t,h}\right) \tag{3.6}$$

$$\mu_{t,h} = \bar{w}'F^h\bar{x}_{t-h} \tag{3.7}$$

$$\sigma^2_{t,h} = \sum_{i=1}^{h} \bar{w}'F^{h-i}CC'F^{h-i'}\bar{w}\sigma^2_\epsilon \tag{3.8}$$

Further, we can decompose $X_{t|t-h}$ into its predictable component $\mu_{t,h}$ and the sum of unforecast-able shocks between $t-h$ and t: $\nu_{t,h} \equiv \sum_{i=1}^{h} \bar{w}'F^{h-i}C\epsilon_{t-h+i}$. Similarly, the forecast error of a particular forecast may be decomposed into a prediction error on the conditional mean and the unpredictable component which is common to all forecasters' errors.

$$X_{t|t-h} - \hat{\mu}_{i,t,h} = \mu_{t,h} - \hat{\mu}_{i,t,h} + \nu_{t,h} \tag{3.9}$$

The expected value of the mean error component represent the extent of optimism bias which I decompose into a forecaster-specific and forecast horizon-specific terms:

$$\frac{1}{\hat{\sigma}_{i,t,h}}\mathbb{E}[\mu_{t,h} - \hat{\mu}_{i,t,h}] = \alpha_{i,h} \tag{3.10}$$

Further, I assume that the covariance in the mean error component across forecasters, time and horizon to be zero unless the forecast is made by the same forecaster with the same forecast target:

$$Cov(\mu_{t,h} - \hat{\mu}_{i,t,h}, \mu_{s,k} - \hat{\mu}_{j,s,k}) = \sigma^2_{i,h,k} = \begin{cases} \sigma^2_i & \forall \quad i=j, t=s \\ 0 & \text{else} \end{cases} \tag{3.11}$$

This covariance may be interpreted as some measure of the forecasting ability of a forecaster and earlier estimates by Davies and Lahiri (1995); Boero et al. (2008a) for the Blue Chip Survey and Bank of England Survey of External Forecasters respectively suggest substantial heterogeneiety of forecasting ability across forecasters.

Then, following the framework developed by Davies and Lahiri (1995) and extended by Davies (2006), the covariance structure across two (un-standardized) forecast errors is:

$$\mathbb{E}[(X_{t,h} - \hat{\mu}_{i,t,h})(X_{s,k} - \hat{\mu}_{j,s,k})] = \begin{cases} \sigma_i^2 + \sigma_{t,min(h,k)}^2 & \forall i = j, t = s \\ \sigma_{t,min(h,k)}^2 & \forall i \neq j, t = s \\ \sigma_{t,min(h,k-4)}^2 & \forall t + 1 = s, k > 4 \\ 0 & else \end{cases}$$

$$\begin{aligned} \mathbb{E}[\mu_{t,h} - \hat{\mu}_{i,t,h}, \nu_{t,h}] &= 0 \\ \sigma_{t,h}^2 &= \sigma_\epsilon^2 \sum_{i=1}^{h} (\rho^{h-i} w_{h+1-i})^2 \end{aligned}$$

where the last result is derived from equation 3.8. My covariance structure differs from Davies and Lahiri (1995) and Davies (2006) in that I standardize the forecast errors and consequently, the covariance between two standardized forecast errors $z_{i,t,h} \equiv \frac{X_t - \hat{\mu}_{i,t,h}}{\hat{\sigma}_{i,t,h}}$ is better interpreted as a correlation coefficient. This is only true however, if indeed the variance estimate in the forecast equals the variance of the forecast error. In my framework I separately estimate the bias estimate with the potential correlation across forecast errors. First, I decompose the standard deviation of a density forecast into the product of a confidence bias component and the true standard deviation of the forecast:

$$\hat{\sigma}_{i,t,h} = (\beta_{i,h})^{-1} (\sigma_i^2 + \sigma_{t,h}^2)^{\frac{1}{2}} \tag{3.12}$$

Further, use equation 3.8 to pin down the relationship between the variance of forecasts with different horizons $\sigma_{t,h}^2 = \tilde{w}_h \sigma_\epsilon^2$ and for simplicity let $min(h,k) = h$.[21] We can then decompose the covariance between standardized forecast errors into the product of

[21] $\tilde{w}_h \equiv \sum_{i=1}^{h} (\rho^{h-i} w_{h+1-i})^2$ and $\tilde{w}_h = 0 \forall h \leq 0$.

overconfidence biases and a correlation coefficient:

$$\mathbb{E}[z_{i,t,h}, z_{j,s,k}] = \begin{cases} \beta_{i,h}\beta_{i,k}p_{i,t,h,k} & \forall i = j, t = s \\ \beta_{i,h}\beta_{j,k}p'_{i,j,t,h,k} & \forall i \neq j, t = s \\ \beta_{i,h}\beta_{j,k}p'_{i,j,t,h,k-4} & \forall t = s-1, k > 4 \\ 0 & \text{else} \end{cases}$$

where

$$p_{i,t,h,k} = \left[\frac{\sigma_i^2 + \sigma_\epsilon^2 min(\tilde{w}_h, \tilde{w}_k)}{\sigma_i^2 + \sigma_\epsilon^2 max(\tilde{w}_h, \tilde{w}_k)}\right]^{\frac{1}{2}}$$

$$p'_{i,j,t,h,k} = \left[\frac{\sigma_\epsilon^4 (min(\tilde{w}_h, \tilde{w}_k))^2}{(\sigma_i^2 + \sigma_\epsilon^2 \tilde{w}_h)(\sigma_j^2 + \sigma_\epsilon^2 \tilde{w}_k)}\right]^{\frac{1}{2}}$$

Finally, to complete the description of the standardized forecast errors, define the vector of standardized forecast errors z such that:

$$\underset{THN\times 1}{z} = \begin{bmatrix} z_1 \\ \vdots \\ z_t \\ \vdots \\ z_T \end{bmatrix}, \quad \underset{HN\times 1}{z_t} = \begin{bmatrix} z_{1,t} \\ \vdots \\ z_{h,t} \\ \vdots \\ z_{H,t} \end{bmatrix}, \quad \underset{N\times 1}{z_{h,t}} = \begin{bmatrix} z_{1,h,t} \\ \vdots \\ z_{n,h,t} \\ \vdots \\ z_{N,h,t} \end{bmatrix}$$

Then $z \sim \mathbf{MVN}(\alpha, \beta P \beta)$ where:

$$\alpha = \begin{bmatrix} \alpha_t \\ \vdots \alpha_t \end{bmatrix} \quad , \alpha_t = \begin{bmatrix} \alpha_{1N} \\ \vdots \\ \alpha_{hN} \\ \vdots \alpha_{HN} \end{bmatrix} \quad , \alpha_{hN} = \begin{bmatrix} \alpha_{h,1} \\ \vdots \\ \alpha_{h,n} \\ \vdots \\ \alpha_{h,N} \end{bmatrix}$$

$$\beta = \begin{bmatrix} \beta t \\ \vdots \beta_t \end{bmatrix} \quad , \beta_t = \begin{bmatrix} \beta_{1N} \\ \vdots \\ \beta_{hN} \\ \vdots \beta_{HN} \end{bmatrix} \quad , \beta_{hN} = \begin{bmatrix} \beta_{h,1} \\ \vdots \\ \beta_{h,n} \\ \vdots \\ \beta_{h,N} \end{bmatrix}$$

$$P = \begin{bmatrix} C & D^{'} & \mathbf{0} & \dots & \mathbf{0} \\ D & C & \ddots & \dots & \vdots \\ \mathbf{0} & \ddots & \ddots & \ddots & \vdots \\ \vdots & & \ddots & & D^{'} \\ \mathbf{0} & \dots & \dots & D & C \end{bmatrix} , C = \begin{bmatrix} A_1 & B_{1,2} & \dots & B_{1,H} \\ B_{2,1} & A_2 & & \vdots \\ \vdots & & \ddots & \vdots \\ B_{H,1} & & & A_H \end{bmatrix} , D = \begin{bmatrix} \tilde{B}_{1,1} & \dots & \tilde{B}_{1,H} \\ \vdots & \ddots & \vdots \\ \tilde{B}_{H,1} & & \tilde{B}_{H,H} \end{bmatrix}$$

where

$$\begin{aligned} A_h(n,m) &= \begin{cases} 1 & \text{if } n = m \\ \left[\frac{(\tilde{w}_h \sigma_\epsilon^2)^2}{(\sigma_n^2 + \tilde{w}_h \sigma_\epsilon^2)(\sigma_m^2 + \tilde{w}_h \sigma_\epsilon^2)}\right]^{\frac{1}{2}} & \text{if } n \neq m \end{cases} \\ B_{h,k}(n,m) &= \begin{cases} \left[\frac{\sigma_n^2 + min(\tilde{w}_h, \tilde{w}_k)\sigma_\epsilon^2}{(\sigma_n^2 + max(\tilde{w}_h, \tilde{w}_k)\sigma_\epsilon^2)}\right]^{\frac{1}{2}} & \text{if } n = m \\ \left[\frac{(min(\tilde{w}_h, \tilde{w}_k)\sigma_\epsilon^2)^2}{(\sigma_n^2 + \tilde{w}_h \sigma_\epsilon^2)(\sigma_m^2 + \tilde{w}_k \sigma_\epsilon^2)}\right]^{\frac{1}{2}} & \text{if } n \neq m \end{cases} \\ \tilde{B}_{h,k}(n,m) &= \left[\frac{(min(\tilde{w}_h, \tilde{w}_k)\sigma_\epsilon^2)^2}{(\sigma_n^2 + \tilde{w}_h \sigma_\epsilon^2)(\sigma_m^2 + \tilde{w}_k \sigma_\epsilon^2)}\right]^{\frac{1}{2}} \end{aligned}$$

where $\tilde{w}_h \equiv \sum_{i=1}^{h} \rho^{h-i} w_{h+1-i}$ and $\tilde{w}_h = 0 \forall h \leq 0$ while $w_i \equiv \frac{4-|4-i|}{4}$ and $w_i = 0 \forall i \notin [1,7]$.

The vectors α and β yield $2 \times (N + H)$ optimism and bias parameters to estimate while the correlation matrix P requires N variance parameters plus the quarterly shock variance σ_ϵ^2 and the quarterly growth persistence parameter ρ. Note that the last two parameters may be estimated from quarterly GDP growth data directly. The parameters are estimated by Maximum Likelihood and we use the compacted version of the vectors α and β as well as the matrix P by deleting the appropriate rows and columns given the unbalanced panel data from the survey forecasts. Finally note that in the estimation exercise of this chapter, the mean vector α includes an AR(1) term.