Coordination under ambiguity

Stéphane Dupraz*

Columbia University

November 10, 2015

Abstract

How do agents coordinate in a world that they do not fully understand? I consider a dispersed-information coordination game with ambiguity-averse agents who do not trust their models. Because distinguishing models is harder in a noisier economy, the model is one of endogenous ambiguity. Because one agent's noise is another's private information, one agent's reliance on his private information increases how much ambiguity his neighbor faces. I revisit the role of private and public information in this new light. On the positive side, I show that the equilibrium depends less on fundamentals as agents become more ambiguity averse, and not at all in the limit where they become infinitely so. I also show that, because it makes agents trust their model more, the release of public information drives the economy toward fundamentals whenever ambiguity-aversion is high enough, in contrast to the standard result under rational expectations. On the normative side, I show that the equilibrium features too much dependence on fundamentals: agents would rather live in a world that they understand better, even if it means living in a world that is less responsive to changes in fundamentals.

Introduction

The importance of the dispersion and slow diffusion of information has recently received a renewed interest in macroeconomic models of the business cycles. Dispersed information has important consequences in environments where coordination is essential, be it in monetary models—the strategic complementarity in price-setting (Woodford (2003), Mankiw and Reis (2002))—or real models—the strategic complementarity arising from demand spillovers (Angeletos and La'O (2009, 2013); Lorenzoni (2009)). Just as reduced-form models of strategic

*Email: sd2702@columbia.edu. I sincerely thank Ricardo Reis and Michael Woodford for their guidance and support, as well as Jennifer La'O, Emi Nakamura, Jón Steinsson and seminar participants at Columbia University for comments and suggestions. Errors are mine.

complementarities (Cooper and John (1980); Cooper (1999), and references therein) had shed light on the mechanisms at play in macroeconomic models of coordination failures popular in the New Keynesian literature, reduced-form models of strategic complementarities under dispersed information (Morris and Shin (2002), Angeletos and Pavan (2007)) have shed light on the mechanisms at play in these new macroeconomic models. In such models, where an agent's welfare depends partly on how close his action is to others', introducing imperfect and dispersed information lets higher-order beliefs—what I believe you believe and so on—take central stage. Because public information is a better predictor of others' actions, public signals become a coordination device, driving the economy away from fundamentals. Some of the general insights that have emerged from these models are as follows. First, the stronger strategic complementarities are—the stronger the motive to do the same thing as others—the less dependent on fundamentals the equilibrium is. Second, releasing public information decreases the equilibrium dependence on fundamentals. Third, the equilibrium reliance on private and public information, and resulting departure from fundamentals, may however be efficient.

While these models have increased our understanding of the role of dispersed information in business cycles, in such complex interactions as pricing and production decisions which depend on the whole macroeconomy, the uncertainty agents face is likely to consist not only of risk, but also of ambiguity—ignorance of the true model of the economy. The assumption of rational expectations on which beauty contests rely assumes however that agents fully trust their models, ruling out ambiguity as a source of uncertainty. I consider a beauty contest model where agents face both risk and ambiguity, building on the robust control literature (Hansen and Sargent (2008), and references therein) to model the way agents fear model misspecifications. Agents have more trouble uncovering the true model of the economy when the economy features more unforecastable fluctuations—when there is more noise in the regressions they run to learn the model. As a consequence, ambiguity is endogenously determined in equilibrium: each agent's action depends on the ambiguity he faces; but in turn, ambiguity depends on agents' actions. Ultimately, the noise one agent faces is another's private information. Section 1 and 2 detail the set-up and how ambiguity is endogenously determined in equilibrium.

Section 3 analyses the positive properties of the model. The two main results revisit the role of public signals through their interaction with ambiguity-aversion. First, in addition to increasing with strategic complementarities, the dependence on public signals increases with ambiguity aversion. When agents face ambiguity in the actions of others, they are less willing to take the risk to respond to fundamentals and seek to act more like others. To do so, they rely more on public information, driving the economy away from fundamentals. For any degree of fundamental strategic complementarity, at the limit where the concern for model misspecification increases, the equilibrium depends on public signals only, and not at all on fundamentals. Second, releases of public information do not necessarily drive the economy away from fundamentals, as happens under rational expectations, and as some authors have worried could be detrimental to welfare. Whenever ambiguity-aversion is high enough, releasing public information increases agents' understanding of the world, decreases ambiguity, and makes them more willing to take the risk to respond to fundamentals: the economy tracks fundamentals more closely.

Section 4 considers the normative properties of the equilibrium. I show that, although the equilibrium depends more on public signals than it does under rational expectations, it does not depend on public signals enough. Under rational expectations, the equilibrium of the game I consider is efficient; no longer so when agents fear model misspecifications. In equilibrium, each agent fails to internalize that when he reacts to his private information and idiosyncratic shock, he creates noise for everyone else, decreasing everyone's understanding of the world, and increasing ambiguity. The equilibrium generally features an overproduction of ambiguity.

Although I stress the interpretation of robust control as a model of ambiguity-aversion, the multiplier preferences I use are consistent with expected utility, leading to an interpretation of the model that involves risk-aversion only. Seen under this light, the model describes how the positive and normative results of coordination games change when one moves away from the quadratic preferences on which much of the reduced-form literature on dispersed information has focused. When the certainty equivalence no longer holds, agents' aversion to risk moves the equilibrium toward public signals, but not as much as efficiency requires: the equilibrium is too risky.

The paper builds on both the literature on dispersed information in coordination games, and the literature on decision-theory under ambiguity, especially robust control. Models of dispersed information have wide applications in macroeconomics. Early applications include the possibility of multiple equilibria in currency attacks (Morris and Shin (1998), Angeletos and Werning (2006)) and the persistence of inflation in models of price-setting (Woodford (2003), Mankiw and Reis (2002)). More recently, several articles have stressed how these models also apply to a central strategic complementarity in decentralized market economies: demand linkages. Angeletos and La'O (2009, 2013) and Lorenzoni (2009) show how common errors in expectations can drive business cycles, while Benhabib et al. (2015, 2013) show how dispersed information can give rise to multiple equilibria. Schaal and Taschereau-Dumouchel (2015) consider investment decisions and show how an economy can fall into a coordination trap, remaining in a recession for a very long time. The present paper is closest to the reduced-form branch of the literature, such as Morris and Shin (2002) and Angeletos and Pavan (2007).

On the ambiguity-aversion side, this paper explicitly relies on the robustness literature as reviewed in Hansen and Sargent (2008). Many of the early models on robustness are decision-theoretical and not equilibrium models. Closer to the spirit of this model are papers that consider ambiguity-averse agents within an equilibrium model. Ilut and Schneider (2014) consider the effect of ambiguity shocks in a real business cycles model where ambiguity—how large the set of models agents consider is—is exogenous. Backus et al. (2015) also consider a real business cycles model with ambiguity-averse agents, this time using Klibanoff et al. (2009)'s recursive smooth ambiguity to model ambiguity-aversion. Again, ambiguity is exogenous in their model. In contrast, Bidder and Smith (2012) use the multiplier preferences of robust control in a similar real business cycles model, and stress the dependence on risk of the worst-case scenario agents fear. Part of my contribution is to explain why robust control generates such a dependence of ambiguity on risk. In addition, none of these models considers strategic complementarities between heterogenous agents or dispersed information, which is at the heart of this paper.

1 Set-up

Each agent i in a continuum $i \in [0,1]$ seeks to pick an action y_i so as to minimize a weighted average of the expected distances to two targets. First, an exogenous, idiosyncratic and observed target a_i. This fundamental is the sum of an aggregate component $\bar{a}$, with mean m and variance $1/\kappa_a$, and an idiosyncratic component b_i, independently and identically distributed with zero mean (so that the average fundamental is $\bar{a}$) and variance $1/\kappa_b$: $a_i = \bar{a} + b_i$. I assume all fundamentals, as well as the signals to be defined, to be jointly gaussian. Second, the average action of others:

$$\bar{y} \equiv \int y_i di. \tag{1}$$

Given an information set ω_i, agent i seeks to minimize its expectation, with respect to his information set ω_i, of the quadratic loss:

$$L(y_i, \bar{y}, a_i) = (1-\alpha)\left(\frac{(y_i - a_i)^2}{2}\right) + \alpha\left(\frac{(y_i - \bar{y})^2}{2}\right),\ \alpha \in (0,1). \tag{2}$$

The weight α on the average action parameterizes the degree of strategic complementarities. I assume it to be positive, and strictly so: for $\alpha = 0$, there is no tradeoff between the two targets and no uncertainty in targeting a_i, so that the solution is trivially $y_i = a_i$.

This reduced-form model is emblematic of the literature on beauty contests. One feature worth stressing is that I assume an idiosyncratic and observed fundamental target a_i.[1] Agents know their own fundamentals and the only variable they need to expect is the aggregate action $\bar{y}$: uncertainty arises only from the need to

[1]The early literature focused on a common and imperfectly observed fundamental (e.g. Morris and Shin (2002); Angeletos and Pavan (2007)). Allowing for idiosyncratic fundamentals has become common. For instance, Bergemann et al. (2015) analyze how the information structure shapes aggregate volatility in a framework that allows for idiosyncratic shocks.

coordinate with others. My reason for doing so is that in many applications, this is the central information problem agents face. For instance, in the application to demand linkages developed in Angeletos and La'O (2009, 2013), the fundamental a_i is a firm's idiosyncratic technology shock which we can think of as being relatively well known to the firm; the firm's uncertainty bears on aggregate demand $\bar{y}$.[2] Technically, it is also easier to treat the case where agents need to form expectations on a single variable.

The information set ω_i of agent i consists in private and public information. The idiosyncratic signal a_i is in itself a private signal of $\bar{a}$. I let there be additional private information by considering a second private signal $\tilde{x}_i = \bar{a} + \tilde{\varepsilon}_i^x$, with precision $\tilde{\kappa}_x$, such that a_i and ε_i^x are independent, and ε_i^x are independent across agents. Private information can be summed up by the single signal $x_i = \frac{\kappa_b}{\kappa_x} a_i + \frac{\tilde{\kappa}_x}{\kappa_x} \tilde{x}_i$, with precision $\kappa_x \equiv \tilde{\kappa}_x + \kappa_b$. There exists a public signal z which also takes the form of average fundamental plus noise: $z \equiv \bar{a} + \varepsilon^z$, with precision κ_z, such that ε^z is independent of ε_i^x for all i. When this prior is not flat ($\kappa_a \neq 0$), the unconditional mean m plays the role of a second public signal. The weighted average $z' = \frac{\kappa_z}{\kappa_a + \kappa_z} z + \frac{\kappa_a}{\kappa_a + \kappa_z} m$ of the two public signals z and m, with precision $\kappa_{z'} = \kappa_a + \kappa_z$, can be treated as the unique public signal.

1.1 Rational Expectations non robust solution

As a benchmark, I derive the rational expectations solution of the game.

Definition 1. *A RE equilibrium is strategies for all i such that all agents minimize their expectations of their loss* (2) *where the aggregate action $\bar{y}$ is defined by* (1) *and beliefs are the equilibrium ones.*

The unique equilibrium can be derived along the same lines as in Morris and Shin (2002). The objective of agent i is concave, so that the first-order condition selects the unique best response of agent i. It is a weighted average of its two targets a_i and $\bar{y}$:

$$y_i = (1 - \alpha) a_i + \alpha E_i(\bar{y}). \tag{3}$$

So that the average action is:

$$\bar{y} = (1 - \alpha)\bar{a} + \alpha \bar{E}(\bar{y}), \tag{4}$$

where $\bar{E}(.) \equiv \int E_i(.)di$ designates the average expectation. It is practical to rely on the apparatus of formal power series that is standard in time-series. Note H the average higher-order belief operator; the equation the

[2] Demand spillovers were also one of the major application of games of strategic complementarities in the early new keynesian literature. See e.g. Hart (1982); Heller (1986); Cooper (1994); Bryant (1983); Roberts (1987); Kiyotaki (1988); Shleifer (1986); Murphy et al. (1989).

average action solves can be written:

$$(I - \alpha H)\bar{y} = (1 - \alpha)\bar{a}. \tag{5}$$

Restricting to solutions such that $\alpha^k \bar{E}^{(k)}(\bar{y})$ tends to zero with k, that is excluding that agents believe that agents believe ... that the aggregate action is infinite, the polynomial can be inverted to give:

$$\bar{y} = (1 - \alpha) \sum_{k=0}^{\infty} \alpha^k \bar{E}^{(k)}(\bar{a}), \tag{6}$$

where $\bar{E}^{(k)}(.) = H^k$ is the k^{th}-order average belief operator, defined recursively by $\bar{E}^{(k)}(.) = \bar{E}(\bar{E}^{(k-1)}(.))$. The higher-order beliefs of $\bar{a}$ are easy to solve for by induction:

$$\bar{E}^{(k)}(\bar{a}) = \mu^k \bar{a} + (1 - \mu^k) z', \tag{7}$$

where $\mu = \kappa_x / \kappa$ is the bayesian weight on private information. Beliefs of higher-orders depend more and more on the public signal z' at the expense of the private signal x_i. Plugging in this expression of higher-order beliefs (7) in (6) gives the unique solution for $\bar{y}$. This solution for $\bar{y}$ can be plugged into (3) to give the corresponding unique best-response function.

Lemma 1. *There exists a unique RE equilibrium. The aggregate action follows $\bar{y} = \Phi^* \bar{a} + (1 - \Phi^*) z'$, with:*

$$\Phi^* \equiv \frac{1 - \alpha}{1 - \alpha\mu}. \tag{8}$$

Best-responses take the linear form $y_i = \phi^*_a a_i + \phi^*_x x_i + \phi^*_{z'} z'$, with:

$$\phi^*_a \equiv 1 - \alpha, \tag{9}$$

$$\phi^*_x \equiv \Phi^* - (1 - \alpha), \tag{10}$$

$$\phi^*_{z'} \equiv 1 - \Phi^*. \tag{11}$$

With no public information $\mu = 1$, the equilibrium aggregate action follows fundamentals $\bar{y} = \bar{a}$, regardless of the degree of strategic complementarities α. With any positive amount of public information $\mu < 1$, the dependence on fundamentals Φ^* is decreasing from 1 to 0 as α increases from 0 to 1: public signals combined with strategic complementarities move the economy away from fundamentals to common errors in expectations. Besides, the dependence on fundamentals Φ^* is increasing in μ the weight of private information. Releasing public information thus decreases the weight of private information and decreases the dependence on fundamentals.

Is the equilibrium efficient? Define the efficient allocation as the one that agents would pick if they could commit to an action ex ante, before the realization of their idiosyncratic shocks and signals. Under this veil of ignorance, this is akin to minimizing a utilitarian social loss functions, weighted by the probability of occurrence of the idiosyncratic fundamentals and information. For the welfare criterion to be relevant, and as stressed by Angeletos and Pavan (2007), it must incorporate the same information constraints as in equilibrium: the action y_i of agent i needs to be measurable with respect to his information set $\omega_i = \{x_i, z\}$; the aggregate action needs to be measurable with respect to the aggregate variables $\bar{a}$ and z. Formally, an efficient allocation solves the program:

$$\min_{\substack{(y_i(x_i,z))_i \\ \bar{y}(\bar{a},z)}} E\left((1-\alpha)\frac{(y_i-a_i)^2}{2} + \alpha E_i\left(\frac{(y_i-\bar{y})^2}{2}\right)\right), \tag{12}$$

$$\text{st.}\bar{y} = \int y_i d_i.$$

As shown in the appendix, the equilibrium is efficient under rational expectations.

Lemma 2. *The rational expectations equilibrium corresponds to the unique efficient allocation.*

1.2 Concerns for model misspecifications and robust RE equilibrium

I now introduce a concern for model misspecifications in agents' preferences. I follow the literature on robust control (Hansen and Sargent (2008), and references therein) to model how agents value and choose actions in the face of ambiguity. I consider that agents are unsure of their model, not that they worry that others may have a different model (Woodford (2010) is an example of the latter).

Agent i does not know where $\bar{y}$ lies on the real line. Although it is always possible to represent agent i's uncertainty of the location of $\bar{y}$ through a conditional distribution $f(\bar{y}|\omega_i)$ on $\mathbb{R}$, Knight (1921)'s classical distinction between *risk* and *ambiguity* stresses that it is completely legitimate to do so only when uncertainty consists of *risk*, defined as a *"quantity susceptible of measurement"*.[3] A measure is a probability distribution; call it a model.

Definition 2. *A model is any probability distribution on $\bar{y} \in \mathbb{R}$, conditional on available information ω_i, $f(\bar{y}|\omega_i)$.*

Note that the definition of a model does not connote a theoretical model: the definition is agnostic on how the agent came to form his model. It could be from theory—reasoning in a representation of the world in

[3]To be sure, risk and *uncertainty* in Knight's own words. The word *uncertainty* is however commonly used in economics without a precise definition, to refer to several things, including *risk*. I follow the common practice in the literature and use *uncertainty* to refer to any form of not knowing, and *ambiguity* to refer to *Knightian uncertainty*.

which it is possible to deduce relationships by making assumptions and thinking about their implications for equilibrium—or pure atheoretical econometrics—observing the past predictive power of observables contained in ω_i—or any combination of the two. Risk is then the fluctuations in $\bar{y}$ perceived by a model $f(\bar{y}|\omega_i)$. As far as decision-theory is concerned, it matters little whether the distribution $f(\bar{y}|\omega_i)$ is the objective one $f^*(\bar{y}|\omega_i)$, that is whether it corresponds to the actual distribution of equilibrium fluctuations: the agent can be wrong in his perception of risk, without altering the manner he behaves under uncertainty.

Ambiguity is the uncertainty the agent faces on top of the uncertainty embedded in a model because he is unsure of the model: the inability to reduce the set of possible models $\mathcal{M} = \{f(.|\omega_i)\}$ to a singleton. Robust control proposes an approach in two steps to describe how an agent makes decisions under ambiguity. First, given a favor model f_0 and a degree of ambiguity A, the agent restricts the set $\mathcal{M}$ of possible models to the neighborhood of a favored model f^0, $B(f_0, A) \equiv \{f/D(f||f^0) \leq A\}$. The single parameter A parameterizes the lack of confidence in the favored model f_0. When $A = 0$, the agent fully trusts his model, while when A tends to infinity, he makes no prior restrictions on the set of possible models. The metric used is relative entropy (the Kullback-Leibler distance) defined as:

$$D(f||f^0) \equiv E^f \Bigg(\log(f) - \log(f^0) \Bigg), \tag{13}$$

for any distribution f that is absolutely continuous with respect to f^0; distributions that are not are considered infinitely distant from f^0. Relative entropy is the most standard metric between distributions used in information theory.[4] Although there is always some arbitrariness in defining a metric, it is an appealing measure of the distance between models as it considers as close models that would be hard to distinguish empirically with any inference method based on the likelihood function. Indeed, given the unknown model f, picking a model f_0 so as to minimize the Kullback-Leibler distance $D(f||f^0)$ to the true model f is equivalent to maximizing $E^f \log(f^0)$. This is unobserved as f is unknown, but maximizing the sample analogue with an iid sample boils down to maximizing the loglikelihood. Relative entropy selects the models that ML estimation, but also bayesian methods and other estimation methods that rely on the likelihood, would have difficulty distinguishing from the favored model f_0. Although the definition of a model was agnostic about where models come from, restricting the space of possible models through relative entropy highlights the empirical way of coming to a model.

Second, within the restricted class of models $B(f_0, A)$, the agent makes decisions according to the maximin

[4] On information theory, see e.g. Cover and Thomas (2006). Relative entropy is not a symmetric metric, so that there exist in principle two different measures of relative entropy: $E_f(\ln(f/f_0))$ and $E_{f_0}(\ln(f_0/f))$. The robust control literature considers the former: when entertaining the possibility that the true density is f, the agent measures the difference between the two distributions using the distribution f.

principle, which has long been offered as a way to model decisions under ambiguity (e.g. Wald (1945); Bellman (1956)). The agent's desire to make decisions that work well not only under his favored model but also under neighboring models is captured by considering the worst-case model (from the agent's perspective). It is formally as if the agent was playing a zero-sum game against nature, fearing that nature would systematically react to his action y_i by picking the model that hurts him most.

Instead of a constraint problem that fixes the size of the entropy ball A, I consider a multiplier problem: increasing the upper-bound on relative entropy has a constant cost $\lambda > 0$, and the agent chooses the size of the set of models given the exogenous cost λ. I do so merely for analytical convenience. Concerns for model misspecifications are still parameterized by a single free parameter λ representing confidence in the model. To sum up, the robust program of agent i is:

$$\min_{y_i} \max_{f_i(\bar{y}|\omega_i)} (1-\alpha)\frac{(y_i - a_i)^2}{2} + \alpha E_i^f \left(\frac{(y_i - \bar{y})^2}{2} \right) - \lambda D(f_i || f_i^0), \tag{14}$$
$$\text{st.} \int f_i(\bar{y}|\omega_i) d\bar{y} = 1.$$

where $E_i^f(.)$ designates the expected value under the neighboring model with density f_i.[5]

As it is, agent i's program is a function of exogenous beliefs parameterized by the favored model f_i^0, in the logic of a temporary equilibrium. The equilibrium concept disciplines beliefs in much the same way as rational expectations do, by imposing that the favored model be the equilibrium one. Just as with rational expectations, an equilibrium can be seen as a fixed point between beliefs and actual outcomes. I therefore still refer to the equilibrium concept as rational expectations, although *robust* rational expectations.

Definition 3. *A robust RE equilibrium with (common) degree of ambiguity λ consists in:*

1. *Best responses such that each agent i, for a favored model f_i^0, solves the robust program* (14)*.*

2. *The aggregate action $\bar{y}$ is given by* (1)*, and agent i's favored model f_i^0 is the equilibrium one.*

Rational expectations stricto sensu correspond to the absence of ambiguity $\lambda = \infty$. This highlights the two components of rational expectations: first the agent's uncertainty consists of risk only—he faces no ambiguity—second he is correct in his assessment of risk. A robust rational expectations equilibrium weakens the requirement on beliefs: an agent's favored assessment of risk is the equilibrium one, but he may face uncertainty stemming

[5]In line with what the agent needs to expect, a model is a distribution on $\bar{y}$, and relative entropy bounds the distance between distributions on $\bar{y}$, not distributions on the exogenous $\bar{a}$. A nice feature of relative entropy is that bounding the distance between distributions on $\bar{a}$ would not affect the relative entropy metric. However, as the agent perceives only $\bar{y}$, it is more natural to phrase the problem this way. Besides, the invariance property is likely to depend on the linear structure of the model, and on the fact that $\bar{y}$ and $\bar{a}$ have the same dimension.

from model ambiguity.[6]

I solve the model using a guess and verify/identify approach. I assume agent i has a favored model $f_i^0(\bar{y}|\omega_i)$ that is Gaussian with variance V_0 independent of his information set ω_i:

$$\bar{y}|\omega_i \sim \mathcal{N}\bigg(E^{f_0}(\bar{y}|\omega_i), V_0\bigg). \tag{15}$$

(V_0 is common to all agents). I then check that these beliefs generate gaussian conditional distributions of this form, and identify the function $E^{f_0}(\bar{y}|.)$ and V_0 to determine the equilibrium.[7] I make no claim to the uniqueness of equilibrium in a broader class of distributions, although the gaussian-linear structure of the model would make the existence of non gaussian-linear equilibria surprising. Note that in the case of rational expectations studied in the previous section, the uniqueness of equilibrium was proven without any prior restriction to a class of models, and the gaussian-linear nature of the unique equilibrium was proved and not assumed.

1.3 Risk interpretation

Although I will mostly focus on the interpretation of the model as one of ambiguity aversion, it is useful to be aware that the multiplier preferences of robust control can equivalently be interpreted as standard (expected-utility) decision under risk, up to increasing the degree of risk aversion embedded in the utility function. The interpretation of the multiplier preferences in (14) as "risk-sensitive" (Hansen and Sargent (1995),Tallarini (2000)) is justified by the possibility to rewrite them as:

$$\min_{y_i} \lambda \log(E_i(e^{\frac{1}{\lambda}L(y_i,\bar{y},a_i)})), \tag{16}$$

where the loss function L is the one defines in (2)—a proof is in the appendix. This representation is of the expected utility form, with the more convex Bernouilli utility function $\exp(\frac{1}{\lambda}L(.))$. The model will shed some light on the reason for this equivalence.

[6]Bayesian decision theory dismisses the practical relevance of the distinction between risk and ambiguity, as defended most famously on axiomatic grounds by Savage (1954). The argument is that agents should form a distribution on the set of possible models, deduce from it a unique probability on $\bar{y}$, and from there rely on standard decision theory—(subjective) expected utility. However, this way of taking ambiguity into account is ruled out in a rational equilibrium model: when equating $f(\bar{y}|\omega_i)$ to the true equilibrium distribution $f^*(\bar{y}|\omega_i)$, rational expectations assume both that the agent perceives the risk correctly, and that the ambiguity he faces is nil. In contrast, robust control allows for an equilibrium concept that imposes that agents have a correct assessment of risk, but do not necessarily face zero ambiguity.

[7]The fixed point problem goes: from beliefs on the aggregate action to the best response; from the best response to the aggregate action; from the aggregate action to beliefs on the aggregate action. Formally, it is possible to start with a guess at any of these three points. Starting at beliefs stresses the view of an equilibrium as a fixed point in beliefs, and stresses that agents do not need to have any knowledge of the underlying structure of the game. He does not even need to be aware that the distribution of $\bar{y}$ is the result of some other people's actions.

2 Equilibrium characterization

Although it is formally possible to solve for the best-response by taking first-order conditions with respect to beliefs and action simultaneously, solving the program sequentially—first deriving worst-case beliefs, plugging them into the program, then maximizing in the action y_i—is insightful.

2.1 Worst-case beliefs

The maximization in the distribution $f(\bar{y}|\omega_i)$ is a concave program, so that a solution is characterized by the first-order conditions. Noting ζ the Lagrange multiplier on the $\int f = 1$ constraint, the Lagrangian corresponding to agent i's robust program (14) is:

$$\min_{y_i} \max_{f(\bar{y}|\omega_i)} (1-\alpha)\left(\frac{(y_i - a_i)^2}{2}\right) + \int \left(\alpha \frac{(y_i - \bar{y})^2}{2} - \lambda\Big(\log(f(\bar{y}|\omega_i)) - \log(f_0(\bar{y}|\omega_i))\Big) - \zeta\right) f(\bar{y}|\omega_i) d\bar{y}. \quad (17)$$

The first-order condition on $f(\bar{y}|\omega_i)$ is:

$$\alpha \frac{(y_i - \bar{y})^2}{2} - \lambda(\log(f(\bar{y}|\omega_i)) - \log(f_0(\bar{y}|\omega_i)) + 1) - \zeta = 0. \quad (18)$$

This implies that the worst-case belief is the following normal distribution:

$$\bar{y}|\omega_i \sim \mathcal{N}\left(\frac{\frac{E^{f_0}(\bar{y}|\omega_i)}{V_0} - \frac{\alpha}{\lambda} y_i}{\frac{1}{V_0} - \frac{\alpha}{\lambda}}, \frac{1}{\frac{1}{V_0} - \frac{\alpha}{\lambda}}\right), \quad (19)$$

provided the following condition to guarantee that the variance is positive:

$$\lambda > \alpha V_0. \quad (20)$$

When condition (20) is not satisfied, the supremum is infinite and worst-case beliefs are not defined. For any action y_i, agent i fears that any of his actions will be followed by a worst-case scenario that hurts him infinitely. There is no meaningful equilibrium in this case. Condition (20) imposes that considering alternative models is costly enough so that this does not happen. To see the effect of a concern for model misspecification under the maximin principle, consider the distortion in the mean:

$$E_i^f(\bar{y}|\omega_i) - E_i^{f_0}(\bar{y}|\omega_i) = -\left(\frac{1}{\frac{\lambda}{\alpha V} - 1}\right)(y_i - E^{f_0}(\bar{y}|\omega_i)). \quad (21)$$

Worst-case beliefs systematically act to exaggerates the distance of agent i's action y_i to its expectation of the uncertain target $\bar{y}$. First, consider the sign of the distortion. If agent i is at the right of the expectation of the target given by his favored model, he will fear that the true model is one that predicts a target even further on the left. Second, consider the amplitude of the distortion. It is determined by the factor $\left(\frac{\lambda}{\alpha V_0} - 1\right)^{-1}$. Quite expectedly, the factor is decreasing in λ—increasing in the concern for model misspecification. It is increasing in α as model misspecifications are more of a concern when the agent initially cares more about $\bar{y}$.

But—and it is key—the inflating factor is also increasing in V_0, the conditional variance. This essential feature tells that in an economy that is more risky—an economy with a higher conditional variance V_0 of the endogenous $\bar{y}$—agent i will also be facing more ambiguity—as measured by a higher mean distortion between the favored and worst-case models. In the vocabulary of econometrics, parameter and specification uncertainty (ambiguity) increase with innovation uncertainty (risk). This is because more volatility makes it more difficult to detect statistical regularities, something the relative entropy metrics takes into account. If $\bar{y}$ depends on variables the agent does not observe—noise from his perspectives—not only does the forecasting power of the variables in his information set (x_i, z) diminishes: the greater noise also reduces the confidence he puts in the relationship between (x_i, z) and $\bar{y}$. As an extreme example, consider the case of a nil variance $V_0 = 0$. Regressing $\bar{y}$ on the information set (x_i, z) would yield a R^2 of one and nil standard deviations: the relationship is deterministic. In this situation, the agent should fear no ambiguity, regardless of his concern for model misspecifications λ. Indeed, relative entropy assesses that all distributions are infinitely distant from a dirac distribution. The intuition carries over with a positive variance V_0: relative entropy assesses that as V_0 increases, the agent should be less and less confident in his model. This parallels how the standard deviations of the coefficients in a regression of $\bar{y}$ on (x_i, z) increase with the variance of noise V_0. To see the point more formally, note that the relative entropy between two gaussian distributions is:

$$D(f||f_0) = \frac{1}{2}\left(\frac{(E^f - E^{f_0})^2}{V^{f^0}}\right) + \frac{1}{2}\left(\frac{V^f}{V^{f^0}} - 1 + \ln\left(\frac{V^{f^0}}{V^f}\right)\right). \tag{22}$$

Normal distributions are characterized by their first two moments. The first term in the relative entropy (22) accounts for the distance between the two distributions due to the difference in their means; it is zero if and only if the two distributions have the same mean. The second term accounts for the distance between the two distributions due to differences in their variances; it is zero if and only if the two distributions have the same variance. What accounts for the issue at hands is that the mean term has the variance V^{f^0} at the denominator: two normal distributions with different means will be harder to distinguish if they have a higher variance.

Note that relative entropy focuses on the dependence of the measure of uncertainty in the variance of the unobservables but abstracts from its dependence in the sample size. In contrast, usual measures of the degree of confidence in one's estimates, such as standard deviations in classical statistics, also decrease as data accumulate. Whether agents who learn the model from data converge to rational expectations is an important question, addressed in the adaptive learning literature (Evans and Honkapohja (2001) and references therein). On the way to the asymptotic rational expectations limit however, standard deviations are non-zero. Whenever standard deviations are positive, agents face *some* ambiguity. How they react to this amount of ambiguity necessarily requires assumptions on their behaviors. Robust control reduces this assumption to the single parameter λ which can be interpreted as a preference parameter, much like risk aversion. Away from the asymptotic limit, the dependence of the degree of confidence in the sample size could matter for practical purposes if standard deviations were known to have decreased in macroeconometrics over time. It does not seem that they have however, partly because the possibility of structural breaks pushes macroeconometricians not to rely on data too far in the past, limiting the increase in the size of the sample from which they learn past regularities.

2.2 Best response

I now turn to how ambiguity—the worst-case scenario—affects the choice of the optimal action y_i. Straightforward algebra shows that:

$$E_i^f(y_i - \bar{y})^2 = C(V_0)^2(y_i - E_i^{f_0}(\bar{y}))^2 + C(V_0)V_0, \tag{23}$$

$$D(f||f_0) = \frac{1}{2}\left[V_0\left(\frac{\alpha}{\lambda}C(V_0)\right)^2(y_i - E_i^{f_0}(\bar{y}))^2 + \Big(C(V_0) - 1 - \ln(C(V_0))\Big)\right], \tag{24}$$

$$\text{where } C(V_0) \equiv \frac{1}{1 - \frac{\alpha V_0}{\lambda}} \geq 1. \tag{25}$$

Plugging these into the objective, minimizing the loss in y_i is equivalent to minimizing:

$$\min_{y_i}(1-\alpha)\left(\frac{(y_i - a_i)^2}{2}\right) + \alpha C(V_0)E_i^{f_0}\left(\frac{(y_i - \bar{y})^2}{2}\right) - \frac{\lambda}{2}\Big(C(V_0) - 1 - \ln(C(V_0))\Big). \tag{26}$$

As the last term is taken as given by agent i, the objective is of the same form as the one with no concern for robustness, but with an inflated weight on the uncertain target $\bar{y}$. It was α; it is now $\alpha C(V_0) \geq \alpha$. To understand why, remember that worst-case beliefs systematically act to exaggerate the distance of agent i's action y_i to its expectation of the uncertain target $\bar{y}$, all the more so that his action is far from the expected target. Without any concern for robustness, the agent shoots somewhere in between his two targets a_i and $\bar{y}$. But if he fears model misspecifications, he then considers alternative models that exaggerate the distance of his action to the uncertain target $\bar{y}$. He therefore adjusts his action closer to $\bar{y}$. As a result, he acts as if he put more weight on the uncertain target $\bar{y}$. Note that robustness matters only because of the existence of a trade-off between two targets. Were he to try to reach a unique uncertain target, agent i's concern for robustness would not affect his action: he would still act at the best forecast of his best model.

Because the objective is of the same form as in the rational-expectations benchmark, so does the best-response. It is of the same linear form as in (3):

$$y_i = (1 - \alpha'(V_0))a_i + \alpha'(V_0)E_i^{f_0}(\bar{y}|\omega_i). \tag{27}$$

with a weight:

$$\alpha'(V_0) \equiv \frac{\alpha C(V_0)}{1 - \alpha + \alpha C(V_0)}. \tag{28}$$

where condition (20) guarantees $\alpha'(V_0) \in (0, 1)$. The weight α' on the target $\bar{y}$ no longer reduces to strategic complementarities α alone: it now also depends on ambiguity. Both higher strategic complementarity and higher ambiguity are reasons to put more weight on $\bar{y}$. Ambiguity aversion therefore argues in favor of higher calibration for the parameter of strategic complementarity in models of price-setting and demand spillovers, a calibration which has been shown to be of great importance. The dependence of α' in ambiguity follows the same intuition as the dependence of the mean distortion (21): α' decreases in λ, but also increases in V_0.

2.3 Aggregate action and identification

The derivation in section 1.1 still applies, replacing α by $\alpha'(V_0)$:

$$\bar{y} = \Phi(V_0)\bar{a} + (1 - \Phi(V_0))z', \tag{29}$$

$$\text{where } \Phi(V_0) \equiv \frac{1 - \alpha'(V_0)}{1 - \alpha'(V_0)\mu} = \frac{1 - \alpha}{1 - \alpha + \alpha(1 - \mu)C(V_0)}. \tag{30}$$

As in the case of rational expectations, the coefficient Φ measures how much $\bar{y}$ depends on fundamental, as opposed to the public signal. The key difference is that Φ is now a function of V_0, because risk determines ambiguity, which in turn determines individual and aggregate behavior. But equilibrium imposes that the favored model of agent i be the one effective in equilibrium. From the expression (29) for $\bar{y}$, I verify that the equilibrium distribution of $\bar{y}$ conditional on ω_i is linear and gaussian:

$$\bar{y}|\omega_i \sim \mathcal{N}\left(\Phi\mu x_i + (1-\Phi\mu)z', \frac{\Phi^2}{\kappa}\right), \tag{31}$$

and coefficients can be identified:

$$E_i^{f_0}(\bar{y}|\omega_i) = \Phi\mu x_i + (1-\Phi\mu)z', \tag{32}$$

$$V_0 = \frac{\Phi^2}{\kappa}. \tag{33}$$

The risk V_0 in the endogenous variable $\bar{y}$ depends on the risk κ^{-1} in the exogenous fundamental, but also in how the equilibrium $\bar{y}$ loads on the fundamentals $\bar{a}$—the parameter Φ. Because Φ results from agents' actions, the model is one of endogenous risk: an economy that depends on imperfectly observed fundamentals is more risky than one that depends on observed public signals. But because risk V_0 affects how much ambiguity agents face, the model is also one of endogenous ambiguity. One agent's behavior affects his neighbor's utility—and ultimately his action—by affecting his neighbor's understanding of the world. If he responds to private signals that he is the only one to observe, his actions will be pure noise for others, creating a world that is difficult to understand. If instead he responds to public signals, his actions will become predictable, making the world easier to understand.

Equation (33) gives V_0 as a function of Φ—how the dependence on fundamentals determines risk—when equation (30) gives Φ as a function of V_0—how risk determines the dependence in fundamentals. Jointly, they determine Φ and V_0.

Lemma 3. *An equilibrium is characterized as* (Φ, V_0) *that satisfy equations* (30) *and* (33)*, with condition* (20)*.*

The best-response function $y_i = \phi_a a_i + \phi_x x_i + \phi_{z'} z'$ can then be recovered from Φ and V_0:

$$\phi_a \equiv 1 - \alpha', \tag{34}$$

$$\phi_x \equiv \Phi - (1-\alpha'), \tag{35}$$

$$\phi_{z'} \equiv 1 - \Phi. \tag{36}$$

3 Public information and ambiguity-aversion

Now that the equilibrium is characterized at the crossing of the two curves in lemma 3—plotted in figure 1—I consider how ambiguity-aversion affects agents' reliance on private and public information, and the dependence

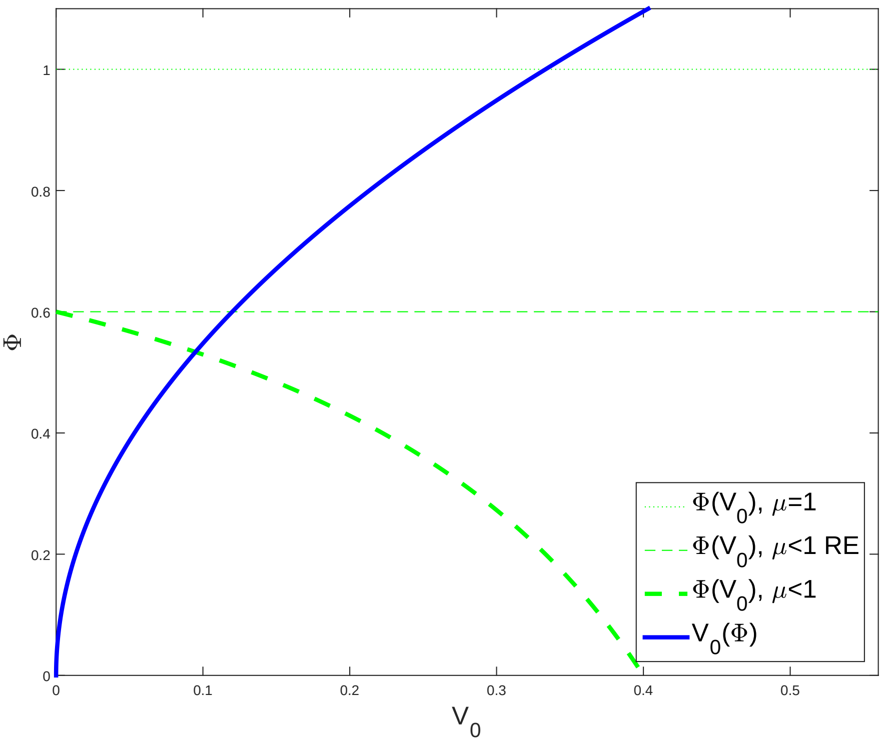

Figure 1: The equilibrium is at the crossing of equation (30) and equation (33). The figure is for $\alpha = 0.5$.

of the aggregate action $\bar{y}$ on the average fundamental $\bar{a}$.

3.1 Ambiguity-aversion as a driver toward public signals

There are two cases to distinguish: the equilibrium with private information only ($\mu = 1$), and the equilibrium with public information ($\mu < 1$). Because the unconditional mean m plays the role of a public signal, the case of private information corresponds to both $\kappa_z = 0$ and $\kappa_a = 0$: it requires a flat, uninformative prior. Although the case of a flat prior is of little practical relevance, considering the case of private information only is helpful to later interpret the effect of public information. With no public signal z', the aggregate action $\bar{y}$ cannot depend on z' and need be equal to the average fundamental $\bar{a}$: $\Phi = 1$—equation (30) is horizontal. The crossing of the two curves merely equates the risk in the endogenous action $\bar{y}$ to the risk in the exogenous fundamentals $\bar{a}$, $1/\kappa_x$. Risk and ambiguity do impact individual action however: the weight α' on the target $E_i^{f_0}(\bar{y})$.

Proposition 1. *With no public information $\kappa_z = \kappa_a = 0$, and for any $\alpha \in (0, 1)$, there exists a unique* *equilibrium whenever $\lambda > \alpha/\kappa_x$, given by $\Phi = 1$, $\phi_{z'} = 0$ and:*

$$\phi_x = \alpha' = \frac{\alpha}{\alpha + (1-\alpha)\left(1 - \frac{\alpha}{\lambda\kappa_x}\right)}, \tag{37}$$

$$\phi_a = 1 - \alpha'. \tag{38}$$

The condition $\lambda\kappa_x > \alpha$ is simply condition (20): for an equilibrium to exist, agents cannot be too ambiguity-averse.

With public information (including the case of public information only $\kappa_x = 0$), the existence and uniqueness of equilibrium is always guaranteed: equation (30) is decreasing from Φ^* to 0 as V_0 increases from 0 to $\frac{\lambda}{\alpha}$, while equation (33) is increasing.[8]

Proposition 2. *With public information $\kappa_z + \kappa_a > 0$, and for any $\alpha \in (0,1)$, there exists a unique equilibrium for any $\lambda > 0$.*

With public information, risk and the dependence on fundamentals become jointly determined in equilibrium in a meaningful way. Just as when there is only private information, higher risk induces higher ambiguity, which leads agents to track the unknown action of others more closely. But now, putting more weight on the action of others means putting more weight on the public signal, which does not get washed out when aggregating individual behaviors: higher risk V_0 induces a lower dependence on fundamentals. The decrease in the dependence on fundamentals in turn decreases risk, as $\bar{y}$ depends more on public signals that are known to everyone, but not enough to wholly counteract the effect: the equilibrium always feature less dependence on fundamentals than under rational expectations.

More generally, the more ambiguity-averse agents are—the lower λ—the less does the equilibrium depend on fundamentals. At the limit where the concern for model misspecifications becomes infinite, the equilibrium depends on public signals only, even when there is little strategic complementarities.

Proposition 3. *With public information $\kappa_z + \kappa_a > 0$,*

- *Φ is increasing in λ, from zero as λ tends to zero to Φ^* as λ tends to infinity.*
- *ϕ_a is increasing in λ, from zero as λ tends to zero to $\phi_a^* = 1 - \alpha$ as λ tends to infinity.*
- *$\phi_{z'}$ is decreasing in λ, from 1 as λ tends zero to $\phi_{z'}^*$ as λ tends to infinity.*

[8] There is no convenient closed-form solution for the equilibrium. Plugging in equation (33) into (30), the equilibrium Φ appears as the unique solution to a cubic equation that satisfies condition (20). Cubic equations do have closed-form solutions, but these are messy and not very insightful.

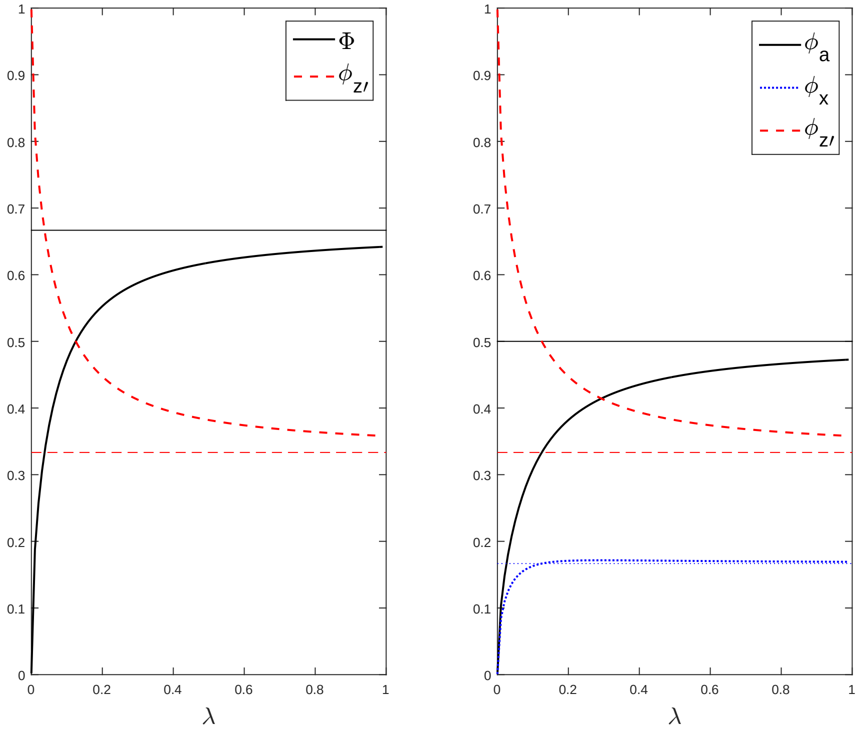

Figure 2: Equilibrium best response coefficients as a function of confidence λ. Dotted horizontal lines are the coefficients under no concern for robustness. The figure is for $\alpha = 0.5$ and $\kappa_a = \kappa_{z'}$.

- *ϕ_x can be non-monotonic, since as λ increases, agents focus less on the uncertain target, but also more on private relative to public information. Nevertheless, it always tends to zero as λ tends to zero, and to its value under no concern for robustness ϕ_x^* as λ tends to infinity.*

Figures 2 illustrates the comparative statics results for a value of strategic complementarity $\alpha = 0.5$ and equal precision of public and private information $\kappa_x = \kappa_{z'}$. It plots the three coefficients ϕ_a, ϕ_x, $\phi_{z'}$ as the concern for model misspecification increases (as λ decreases to zero).

3.2 Releasing public information

Under rational expectations, the equilibrium dependence of $\bar{y}$ on fundamentals $\bar{a}$ depends only on the relative quantity of private vs. public information μ, not on the absolute levels of information κ_x and $\kappa_{z'}$—equation (8). When agents fear model misspecifications, this is no longer the case, as the quantity of information κ affects the risk, hence the ambiguity, in $\bar{y}$—equation (33). When private information κ_x increases, both curves shift up in figure 1: as under rational expectations, equation (30) shifts up because the weight on private information increases; but in addition, equation (33) shifts up because risk, hence ambiguity decreases. As a result, the dependence on fundamental unambiguously increases.

Of greater interest are the comparative statics in the precision of the public signal κ_z, since it corresponds to the policy experiment of releasing public information on the state of the economy, an issue central bankers, for instance, repeatedly face. Under rational expectations, the sole effect of releasing public information is to decrease μ—(30) shifts down—therefore decreasing the dependence of $\bar{y}$ on fundamentals. It is this move away from fundamentals that some authors, the first of whom Morris and Shin (2002), have worried could be detrimental to welfare. However, when agents fear model misspecifications, there is a countervailing effect: when information is released—either private or public—risk diminishes. It is then easier to learn the model of the economy, and agents face less ambiguity. As a result, the are less worried to miss the target of the aggregate action of others $\bar{y}$, and track their individual fundamental a_i more. Formally, equation (33) shifts up. The overall impact of a release of public information is thus ambiguous. The following proposition states that the second effect dominates when ambiguity-aversion is high enough.

Proposition 4. *The dependence on fundamentals Φ increases as more public information is released if and only if $\lambda\kappa_x < \alpha(1-\alpha)^2$.*

Figure 3 illustrates these comparative statics in both cases.

3.3 Risk interpretation

I have focused so far on the standard interpretation of the multiplier preferences of robust control as ambiguity-aversion. As mentioned in the previous section however, these preferences are consistent with expected utility, and thus can alternatively be interpreted as modelling risk-aversion only. Under this interpretation of preferences, the higher weight on the uncertain target $\bar{y}$ in the objective (26) is to be read differently. In arbitrating between reaching his two targets, the agent trades off the certain loss of missing a_i against the uncertain loss of missing $\bar{y}$. Therefore, the more risk-averse the agent is—the lower λ—the more he targets the aggregate action $\bar{y}$. This effect of risk-aversion is however missed in the literature on dispersed information that relies on quadratic preferences, because the certainty equivalence applies: best-response functions are independent of risk. The positive results in this section can be read as the consequences of risk-aversion on the equilibrium reliance on private and public information when going beyond the convenient but ultimately restrictive case of the certainty equivalence: risk-aversion increases the dependence of the economy on public signals, away from fundamentals.

Why exactly do risk-aversion and ambiguity-aversion play the same role in this model? The exact equivalence between the robust control multiplier preferences and expected utility is a bit of a technical coincidence—the

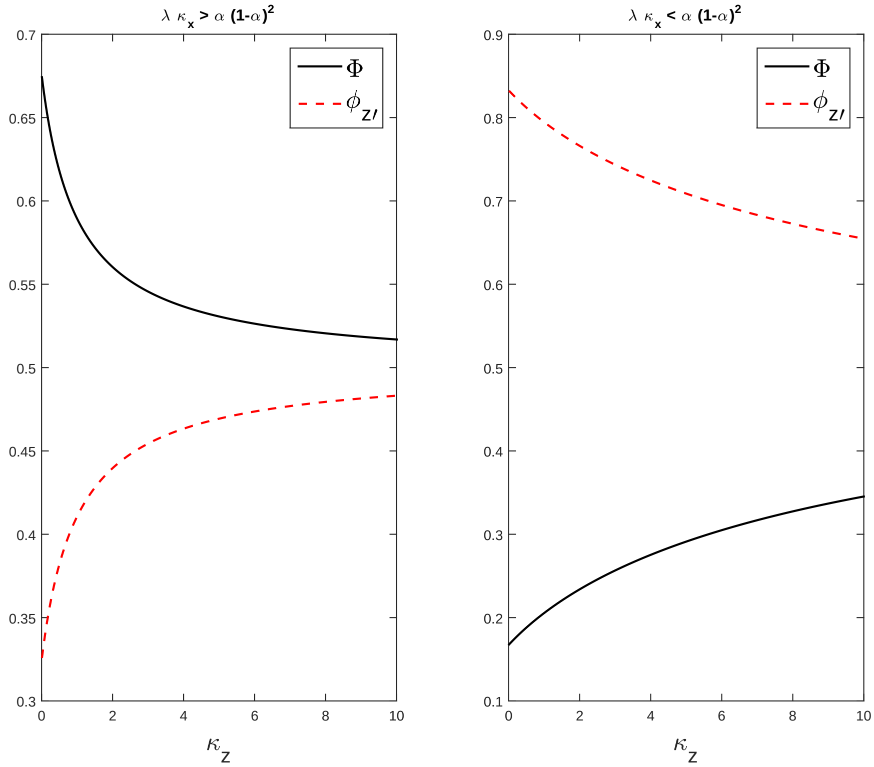

Figure 3: Equilibrium best response coefficients as a function of the precision of public information κ_z.

equivalence does not apply if one considers a constraint problem that bounds the relative entropy between models instead of a multiplier problem. There is however a deep reason why agents who dislike ambiguity ends up disliking risk, a reason the model sheds light on. The relative entropy metric takes into account that agents are more unsure of the model when models are hard to distinguish empirically; learning from data is in turn harder in an economy with more unpredictable fluctuations—a riskier economy. Whether agents dislike risk per se, or because risk makes it hard to learn the true model, does not matter for the predictions of the model.

4 The Overproduction of Ambiguity

Ambiguity drives the equilibrium away from fundamentals relative to the rational expectations equilibrium, as agents react to the ambiguity in $\bar{y}$ by putting more weight on $\bar{y}$, and therefore on public information. In doing so, agents jointly reduce the amount of ambiguity in equilibrium. It is however a non-intended consequence of everyone's actions: no individual agent internalizes his impact on ambiguity. Is the equilibrium amount of ambiguity efficient? I adjust the welfare criterion used under rational expectations to incorporate ambiguity aversion. To make it relevant, the planner cannot directly affect the worst-case scenario that each agent fears, since it is part of their preferences: the planner cannot just tell agents that they have nothing to worry about and hope to be believed on his word—the criterion rules out the confidence fairy. However, the planner can engineer a less risky allocation which will have agents worry less about model misspecifications. Relative to the welfare objective under rational expectations, this is precisely this new concern that I want to focus on. Formally, an efficient allocation is solution to:

$$\min_{\substack{(y_i(x_i,z))_i \\ \bar{y}(\bar{a},z)}} E^0\left(\max_{(f_i(\bar{y}|\omega_i))_i}\{(1-\alpha)\frac{(y_i-a_i)^2}{2}+\alpha E_i^f\left(\frac{(y_i-\bar{y})^2}{2}\right)-\lambda D(f_i||f_i^0)\}\right), \quad (39)$$

$$\text{st.}\bar{y}=\int y_i d_i.$$

Again, I restrict to solutions where agent i's favored model is Gaussian with variance V_0 independent of his information set ω_i, as in equation (15). Worst-case beliefs of agent i given $(y_i, \bar{y})$ are still given by equation (19), provided condition (20). Injecting these beliefs in the objective (39), the problem of the planner reduces to:

$$\min_{\substack{(y_i(x_i,z))_i \\ \bar{y}(\bar{a},z)}} (1-\alpha)E^0\left(\frac{(y_i-a_i)^2}{2}\right)+\alpha C(V_0)E^0\left(\frac{(y_i-\bar{y})^2}{2}\right)-\frac{1}{2}\lambda\Big(C(V_0)-1-\ln(C(V_0))\Big), \quad (40)$$

$$\text{st.}\bar{y}=\int y_i d_i.$$

Lemma 4. *The efficient best-response of agent i satisfies:*

$$(1-\alpha+\alpha C(V_0))y_i = (1-\alpha)a_i + \alpha C(V_0)E_i(\bar{y}) - \delta E_i(\bar{y}-\bar{E}(\bar{y})), \tag{41}$$

$$\delta = C'(V_0)\alpha[E(y_i-\bar{y})^2 - V_0] \geq 0, \ \textit{and} > 0 \textit{ provided } \lambda < \infty. \tag{42}$$

Agent i's efficient best-response corrects the equilibrium best-response (27) by the term $-\delta E_i(\bar{y}-\bar{E}(\bar{y}))$, which is the marginal social cost of increasing y_i due to the resulting increase in ambiguity. It is the product of δ—the marginal social cost of increasing V_0—and $\bar{y}-\bar{E}(\bar{y})$—the marginal increase in V_0 with $\bar{y}$, which increases when y_i increases. Agent i's action impacts everybody else's understanding of the model through his impact on V_0, which in turn affects everybody else's welfare because agents are ambiguity-averse. Efficiency commands that he should take into account this understanding externality and attempt to close the gap between $\bar{y}$ and $\bar{E}(\bar{y})$, that is to reduce the average unexpected component of $\bar{y}$. Because agent i does not observe the gap $\bar{y}-\bar{E}(\bar{y})$, efficiency commands that he should take his best expectation of the term. For instance, if agent i expects others to hold on average too pessimistic expectations on $\bar{y}$ ($\bar{y} > \bar{E}(\bar{y})$), he should temper his private incentives to act so as not to widen the gap $\bar{y}-\bar{E}(\bar{y})$ further.

The presence of the corrective term in the best-response function hints at the inefficiency of the equilibrium. To prove it, I solve for the efficient allocation implied by the efficient best-response. Averaging (41) over i, the aggregate action $\bar{y}$ satisfies:

$$(1-\alpha+\alpha C)\bar{y} = (1-\alpha)\bar{a} + \alpha C\bar{E}(\bar{y}) - \delta\bar{E}(\bar{y}-\bar{E}(\bar{y})). \tag{43}$$

The efficient aggregate action follows a second-order linear difference equation, as opposed to the first-order difference equation the equilibrium allocation satisfies. Using the higher-order polynomial H, this can be written $\bar{y} = P(H)\bar{a}$, with:

$$P(H) = (1-\alpha)\Big((1-\alpha+\alpha C)I - (\alpha C-\delta)H - \delta H^2\Big)^{-1}. \tag{44}$$

Since $H^k\bar{a} = \mu^k\bar{a} + (1-\mu^k)z'$, it follows that $\bar{y} = \Phi\bar{a} + (1-\Phi)z'$ with:

$$\Phi = P(\mu) = \frac{1-\alpha}{1-\alpha+\alpha C(1-\mu)+\delta\mu(1-\mu)}. \tag{45}$$

Just as for equilibrium, the (common) conditional variance of $\bar{y}$ needs satisfy (20) and (33). An efficient allocation is characterized by equations (45) and (33), provided condition (20). Simply put, the efficient allocation can be seen at the crossing of two curves, just the as equilibrium allocation. A technical difficulty however is that the new term δ depends on $E(y_i - \bar{y})^2$, which is itself a function of Φ:

$$E(y_i - \bar{y})^2 = \left(\Phi^2 - \phi_a^2\right)\frac{1}{\kappa_x} + \phi_a^2\frac{1}{\kappa_b}, \tag{46}$$

$$\phi_a(V_0) = \frac{1-\alpha}{1-\alpha+\alpha C(V_0)}. \tag{47}$$

Using equation (33) to replace Φ^2 by its equilibrium value $V_0\kappa$ in equation (46), equation (45) expresses Φ as an explicit function of V_0:

$$\Phi = \frac{1-\alpha}{1-\alpha+\alpha C(V_0)(1-\mu)+\alpha C'(V_0)\mu(1-\mu)\left(\phi_a^2(V_0)\left(\frac{1}{\kappa_b}-\frac{1}{\kappa_x}\right)+V_0\left(\frac{\kappa}{\kappa_x}-1\right)\right)}. \tag{48}$$

It follows that:

Lemma 5. *The efficient allocation is characterized as* (Φ, V_0) *that satisfy equations* (48) *and* (33)*, with condition* (20).

Figure 4 illustrates the determination of the efficient allocation at the crossing of the two curves. Comparing equation (30) and (48), the equilibrium appears efficient if and only if $\delta\mu(1-\mu) = 0$. More generally, the appendix shows that equation (48) always lies below equation (30), and strictly so except when $\delta\mu(1-\mu) = 0$.

Proposition 5. *The efficient allocation is unique.*

- *Under rational expectations ($\lambda = \infty$), the equilibrium is efficient (lemma 2).*
- *If there is only private information ($\mu = 1$), the equilibrium is efficient.*
- *If there is only public information ($\mu = 0$), the equilibrium is efficient.*
- *In all other cases, the equilibrium is inefficient. The equilibrium features too much dependence on fundamentals to the detriment of public information ($\Phi^{eq} > \Phi^{eff}$), and too much conditional volatility ($V_0^{eq} > V_0^{eff}$).*

With only private information, there is no way to move the economy toward a better understood world, so the efficiency of the equilibrium is not surprising. But since the prior always acts as a public signal, the absence of public information requires the unlikely assumption of a flat prior $\kappa_a = 0$, which makes the case irrelevant for practical concerns. With only public information, there is no uncertainty whatsoever in this model—all agents face the same fundamental shock $\bar{a}$ that they all observe—so that the efficiency of the equilibrium is not surprising: again, the case is irrelevant for practical concerns.

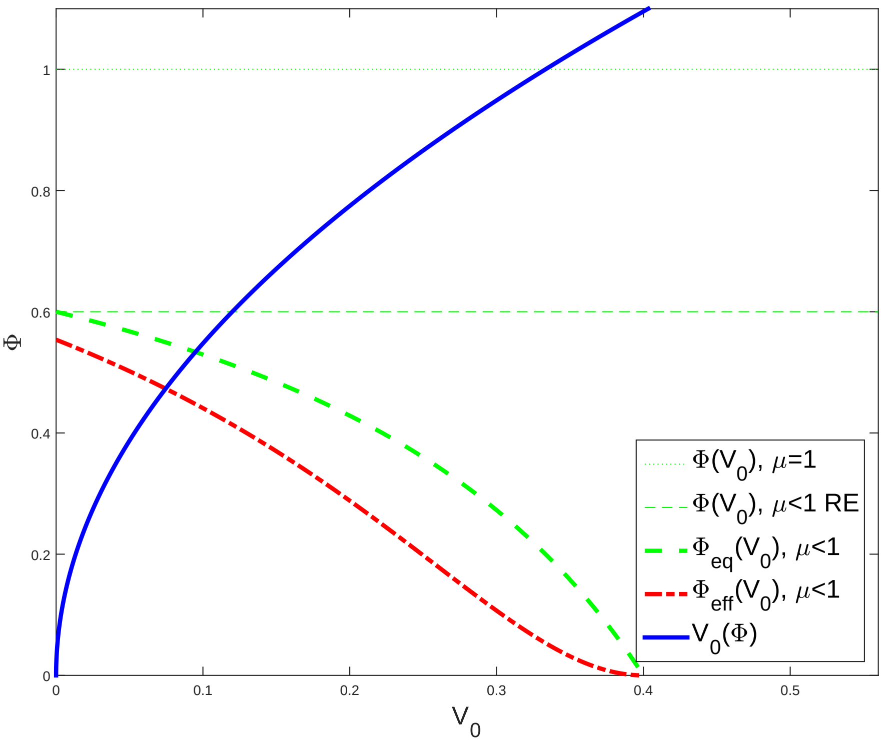

Figure 4: The efficient allocation is at the crossing of equation (48) and equation (33). The figure is for $\alpha = 0.5$.

In all relevant cases, the equilibrium is inefficient when agents fear model misspecifications. Agents do not internalize the effect of their actions on how easy it is to learn the equilibrium. Agents react to their private information too much, creating too much noise for others, creating too much ambiguity on the model governing $\bar{y}$. Simply put, there is an overproduction of ambiguity in equilibrium: agents would rather live in a world they understand better, even if it means living in a world that does not track fundamentals as well.

Again, it is possible to interpret the model as one of risk-aversion instead of ambiguity-aversion. The inefficiency is then to be interpreted differently. Agents do not internalize that in reaction to their private signals, they make the economy more risky for everybody else. The efficiency of the equilibrium in the benchmark case of quadratic preferences ($\lambda = \infty$) appears as the tree that hides the forest. Away from the knife-edge case of the certainty equivalence, the equilibrium is generically inefficient. This matters since much of the literature on the efficient use of information, such as Angeletos and Pavan (2007), has focused on the case of quadratic preferences (of a more general form than the ones covered here).

5 Conclusion

I have looked at a beauty contest with ambiguity-averse agents to reassess the role of public information in coordinating agents' actions. When agents do not fully understand the world they live in, public information differs from private information along a new dimension. A world in which others respond to signals that I also observe is one that I can easily understand; a world in which others respond to signals only they can see is one whose functioning I have trouble uncovering. In equilibrium agents fail to internalize the negative externality of their reliance on private information and private fundamentals on everybody else's understanding of the model: the world is inefficiently ambiguous.

The key feature of the model is that ambiguity is endogenous: it increases with risk, which depends on whether others behave in a predictable way. Alternatively however, the model can be interpreted as one where agents directly care about risk, simply because they are risk-averse. A world where others react to signals I observe is one that is less risky, something I can value per se. The reliance on private information and fundamentals makes for a riskier world, which has positive and normative implications when we move away from the certainty equivalence of the quadratic framework popular in the literature.

The present model is a static one. In a dynamic setting, an increase in risk can be expected to have long-lasting effects on ambiguity: even a short-lived episode of macroeconomic turmoil—volatile GDP and asset prices—could lead agents to become less confident in their understanding of the economy for an extended period of time. A dynamic extension of the present framework may be able to model the often vaguely-defined notion of a "crisis of confidence". It could do so without considering it as an exogenous disturbance, but as an endogenous and long-lasting increase in ambiguity following a short-lived increase in risk. The endogenous reaction of ambiguity could generate long-lasting effects of risk shocks, something the uncertainty literature—which most often equates uncertainty to risk only—typically has difficulty finding (Bloom (2009)).

A Proof of lemma 2

The program of the planner is:

$$\min_{(y(x_i,z')),(\bar{y}(\bar{a},z'))} \frac{1-\alpha}{2}\int (y_i(x_i,z')-a_i)^2 P(x_i,z')dx_i dz' + \frac{\alpha}{2}\iint (y_i(x_i,z')-\bar{y}(\bar{a},z'))^2 P(x_i,z',\bar{a})dx_i dz' d\bar{a},$$
$$st.\forall(\bar{a},z'), \bar{y}(\bar{a},z') = \int y_i(x_i,z')P(x_i|\bar{a},z')dx_i. \tag{49}$$

It is concave. Noting $\beta(\bar{a},z')P(\bar{a},z')$ the Lagrange multiplier on the $(\bar{a},z')$ constraint, the Lagrangian is:

$$\mathcal{L} = \frac{1-\alpha}{2}\int (y_i(x_i,z')-a_i)^2 P(x_i,z')dx_i dz' + \frac{\alpha}{2}\iint (y_i(x_i,z')-\bar{y}(\bar{a},z'))^2 P(x_i,z',\bar{a})dx_i dz' d\bar{a}$$
$$-\int \beta(\bar{a},z')\bar{y}(\bar{a},z')P(\bar{a},z')d\bar{a}dz' + \iint \beta(\bar{a},z')y_i(x_i,z')P(x_i,z',\bar{a})dx_i dz' d\bar{a}. \tag{50}$$

The first-order conditions are:

$$/y_i : (1-\alpha)(y_i-a_i) + \alpha(y_i - E_i(\bar{y})) + E_i(\beta) = 0, \tag{51}$$
$$/\bar{y} : \beta = 0. \tag{52}$$

Therefore, $y_i = (1-\alpha)a_i + \alpha E_i(\bar{y})$, which is the same best-response function as the equilibrium one, therefore yielding the same allocation.

B Proof of the EU representation of robust multiplier preferences

For any density f^*,

$$D(f||f_0) = D(f||f^*) + E^f \log\left(\frac{f^*}{f_0}\right). \tag{53}$$

Pick f^* defined such that $\log\left(\frac{f^*(\bar{y})}{f_0(\bar{y})}\right) = \frac{L(\bar{y})}{\lambda} - \log\left(E^{f_0}(e^{\frac{L(\bar{y})}{\lambda}})\right)$. Then:

$$E^f(L) - \lambda D(f||f_0) = -\lambda D(f||f^*) + \lambda\log(E^{f_0}(e^{\frac{L}{\lambda}})). \tag{54}$$

The robust multiplier preferences maximize this function in f. Note that only the first term of the right-hand side depends on f, and has a maximum of zero reached in $f = f^*$. Hence:

$$\max_f E^f(L) - \lambda D(f||f_0) = \lambda \log(E^{f_0}(e^{\frac{L}{\lambda}})). \tag{55}$$

C Proof of proposition 3

Note that λ is a shifter of (30) only, through its negative impact on $C(V_0) = \frac{1}{1-\frac{\alpha V_0}{\lambda}}$. As λ increases from 0 to ∞, Φ increases from 0 to Φ^*, and V_0 increases. The coefficient $\phi_z = 1 - \Phi$ decreases, from 1 to ϕ_z^*. Given equation (30), Φ^* increases means that C decreases. Therefore, $\phi_a = \frac{1-\alpha}{1-\alpha+\alpha C}$ increases, from 0 to ϕ_a^*. The last coefficient $\phi_x = \Phi - \phi_a$ is a bit more tricky as it can be non monotonic. Differentiating $\phi_a = \frac{1-\alpha}{1-\alpha+\alpha C}$ and $\Phi = \frac{1-\alpha}{1-\alpha+\alpha C(1-\mu)}$, it appears that $\frac{\partial \phi_x}{\partial C}$ has the sign of $(1-\alpha)^2 - (\alpha C)^2(1-\mu)$. Since C is decreasing in λ, $\frac{\partial \phi_x}{\partial \lambda}$ has the sign of $(\alpha C)^2(1-\mu) - (1-\alpha)^2$. This last expression is positive as $\lambda \to 0$ and decreasing in λ. It follows that ϕ_x is non-monotonically increasing in λ if and only if $\alpha^2(1-\mu) < (1-\alpha^2)$.

D Proof of proposition 4

Differentiating equation (30) and (33) in Φ, V_0 and κ_z, and eliminating dV_0 between the two equations, the implicit function theorem gives:

$$\frac{\partial \Phi}{\partial \kappa_z} = \frac{\frac{\Phi \alpha C}{\kappa^2}\left(\frac{\alpha}{\lambda}(1-\mu)CV_0\kappa - \kappa_x\right)}{1-\alpha+\alpha(1-\mu)C + \frac{\alpha^2}{\lambda}(1-\mu)C^2\frac{2\Phi^2}{\kappa}}. \tag{56}$$

It follows that $\frac{\partial \Phi}{\partial \kappa_z}$ has the sign of $\alpha V_0 - \mu\lambda$, that is $\alpha\Phi^2 - \kappa_x\lambda$. Equation (30) and (33) can be combined to characterize the equilibrium Φ as the unique root that lies between 0 and 1 of the polynomial of degree three:

$$P(X) = \alpha X^3 - \alpha X^2 - \lambda\kappa\left(\frac{1-\alpha\mu}{1-\alpha}\right)X + \lambda\kappa. \tag{57}$$

Multiplying $\alpha\Phi^2 - \kappa_x\lambda$ by $(1-\Phi)$ and using the fact that Φ satisfies equation (57), we get that $\frac{\partial \Phi}{\partial \kappa_z} > 0$ if and only if $\Phi < 1-\alpha$. Because $1-\alpha \in (0,1)$, $\frac{\partial \Phi}{\partial \kappa_z} > 0$ if and only if $P(1-\alpha) < 0$, that is if and only if $\alpha(1-\alpha)^2 > \lambda\kappa_x$.

E Proof of lemma 4

Using the fact that posterior variance is constant across agents, it can be written $V_0 = E^0 E_i^{f^0}(\bar{y} - E_i^{f_0}(\bar{y}))^2$. The program of the planner is:

$$\begin{aligned}
\min_{(y(x_i,z')),(\bar{y}(\bar{a},z'))} \; & \frac{1-\alpha}{2}\int (y_i(x_i,z') - a_i)^2 P(x_i,z')dx_i dz' + \frac{\alpha}{2}C(V_0)\iint (y_i(x_i,z') - \bar{y}(\bar{a},z'))^2 P(x_i,z',\bar{a})dx_i dz' d\bar{a} \\
& - \frac{1}{2}\lambda(C(V_0) - 1 - \ln(C(V_0))), \\
& st.\forall(\bar{a},z'), \bar{y}(\bar{a},z') = \int y_i(x_i,z')P(x_i|\bar{a},z')d\bar{a}dz', \\
& st.V_0 = \iint (\bar{y}(\bar{a},z') - E_i(\bar{y}))^2 P(x_i,z',\bar{a})dx_i dz' d\bar{a}.
\end{aligned} \tag{58}$$

It is concave. Noting $\beta(\bar{a},z')P(\bar{a},z')$ the Lagrange multiplier on the first constraints, and $\delta/2$ the constraint on the second constraint, the Lagrangian is:

$$\begin{aligned}
\mathcal{L} =& \frac{1-\alpha}{2}\int (y_i(x_i,z') - a_i)^2 P(x_i,z')dx_i dz' + \frac{\alpha}{2}C(V_0)\iint (y_i(x_i,z') - \bar{y}(\bar{a},z'))^2 P(x_i,z',\bar{a})dx_i dz' d\bar{a} \\
& - \frac{1}{2}\lambda(C(V_0) - 1 - \ln(C(V_0))) \\
& - \int \beta(\bar{a},z')\bar{y}(\bar{a},z')P(\bar{a},z')d\bar{a}dz' + \iint \beta(\bar{a},z')y_i(x_i,z')P(x_i,z',\bar{a})dx_i dz' d\bar{a} \\
& - \frac{\delta}{2}\left(V_0 - \iint (\bar{y}(\bar{a},z') - E_i(\bar{y}))^2 P(x_i,z',\bar{a})dx_i dz' d\bar{a}\right).
\end{aligned} \tag{59}$$

The first-order conditions are:

$$/y_i : (1-\alpha)(y_i - a_i) + \alpha C(V_0)(y_i - E_i(\bar{y})) + E_i(\beta) = 0, \tag{60}$$

$$/\bar{y} : \beta = \delta(\bar{y} - \bar{E}(\bar{y})), \tag{61}$$

$$/V_0 : \delta = \alpha C'(V_0)E(y_i - \bar{y})^2 - \lambda\left(C'(V_0) - \frac{C'(V_0)}{C(V_0)}\right). \tag{62}$$

Injecting equations (61) and (62) in (60) yields equation (41).

The expression of δ can be simplified to $\delta = \alpha C'(V_0)[E(y_i - \bar{y})^2 - V_0]$. To show that it is positive, first notice that $C'(V_0) = \frac{\alpha}{\lambda}C^2(V_0) \geq 0$, with equality if and only if $\lambda = \infty$, ie. under non-robust rational expectations. Besides, $E(y_i - \bar{y})^2 - V_0 = E(y_i - E_i(\bar{y}))^2 \geq 0$.

F Proof of proposition 5

An efficient allocation is any crossing of the two curves (48) and (33) that satisfies condition (20). Equation (48) is a decreasing function of V_0 because δ is an increasing function of V_0. Indeed, write δ as:

$$\delta = C'(V_0)\alpha \left(\frac{1-\alpha}{1-\alpha+\alpha C(V_0)} \right)^2 \left(\frac{1}{\kappa_b} - \frac{1}{\kappa_x} \right) + C'(V_0)\alpha V_0 \left(\frac{\kappa}{\kappa_x} - 1 \right). \tag{63}$$

It is easy to check that both terms are increasing in V_0. Besides, the limit of equation (48) as $V_0 \to 0$ is positive, and the limit as $V_0 \to \frac{\lambda}{\alpha}$ is zero. It follows that the efficient allocation exists and is unique. That the equilibrium is efficient when $\mu = 0$, $(1-\mu) = 0$, or $\delta = 0$ (rational expectations) is immediate. In all other cases, the curve of equation (48) is strictly below the curve of equation (30), so that $\Phi^{eq} > \Phi^{eff}$ and $V_0^{eq} > V_0^{eff}$.

References

Angeletos, G.-M. and La'O, J. (2009). Noisy business cycles. *NBER Macroeconomics Annual 2009*, 24:319–378.

Angeletos, G.-M. and La'O, J. (2013). Sentiments. *Econometrica*, 81(2):739–779.

Angeletos, G.-M. and Werning, I. (2006). Crises and prices: Information aggregation, multiplicity, and volatility. *American Economic Review*, 96(5):1720–1736.

Angeletos, S. and Pavan (2007). Efficient use of information and social value of information. *Econometrica*, 75(4):1103–1142.

Backus, D., Ferriere, A., and Zin, S. (2015). Risk and ambiguity in models of business cycles. *Journal of Monetary Economics*, 69:42–63.

Bellman, R. (1956). Minimization problem. *Bulletin of the American Mathematical Society*, 62(3):270.

Benhabib, J., Wang, P., and Wen, Y. (2013). Uncertainty and sentiment-driven equilibria. mimeo.

Benhabib, J., Wang, P., and Wen, Y. (2015). Sentiments and aggregate demands fluctuations. *Econometrica*, 83(2):549–585.

Bergemann, D., Heumann, T., and Morris, S. (2015). Information and volatility. *Journal of Economic Theory*, 158:427–465.

Bidder, R. and Smith, M. (2012). Robust animal spirits. *Journal of Monetary Economics*, 59:738–750.

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

Bryant, J. (1983). A simple rational expectations keynes-type model. *Quarterly Journal of Economics*, 98(3):525–528.

Cooper, R. (1994). Equilibrium selection in imperfectly competitive economies with multiple equilibria. *Economic Journal*, 104(426):1106–1122.

Cooper, R. (1999). *Coordination Games*. Cambridge University Press.

Cooper, R. and John, A. (1980). Coordinating coordination failures in keynesian models. *Quarterly Journal of Economics*, 103(3):441–463.

Cover, T. M. and Thomas, J. A. (2006). *Elements of Information Theory*. Wiley.

Evans, G. W. and Honkapohja, S. (2001). *Learning and Expectations in Macroeconomics*. Princeton University Press.

Hansen, L. and Sargent, T. (1995). Discounted linear exponential quadratic gaussian control. *IEEE Transactions on Automatic Control*, 40(5):968–971.

Hansen, L. and Sargent, T. (2008). *Robustness*. Princeton University Press.

Hart, O. (1982). A model of imperfect competition with keynesian features. *Quarterly Journal of Economics*, 97(1):109–138.

Heller, W. P. (1986). Coordination failure under complete markets with applications to effective demand. In *Equilibrium Analysis, Essays in honor of Kenneth J. Arrow*, volume 2, pages 155–176. Cambridge University Press.

Ilut, C. and Schneider, M. (2014). Ambiguous business cycles. *American Economic Review*, 104(8):2368–99.

Kiyotaki, N. (1988). Multiple expectational equilibria under monopolistic competition. *Quarterly Journal of Economics*, 103(4):695–713.

Klibanoff, P., Marinacci, M., and Mukerji, S. (2009). Recursive smooth ambiguity preferences. *Journal of Economic Theory*, 144:930–976.

Knight, F. (1921). *Risk, Uncertainty and Profits*. Houghton Mifflin.

Lorenzoni, G. (2009). A theory of demand shocks. *American Economic Review*, 99(5):2050–2084.

Mankiw, G. and Reis, R. (2002). Sticky information versus sticky prices: A proposal to replace the new keynesian phillips curve. *Quarterly Journal of Economics*, 117(4):1295–1328.

Morris, S. and Shin, H. S. (1998). Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review*, 88(3):587–597.

Morris, S. and Shin, H. S. (2002). Social value of public information. *American Economic Review*, 92(5):1521–1534.

Murphy, K. M., Shleifer, A., and Vishny, R. W. (1989). Industrialization and the big push. *Journal of Political Economy*, 97(5):1003–1026.

Roberts, J. (1987). An equilibrium model with involuntary unemployment at flexible competitive prices and wages. *American Economic Review*, 77(5):856–874.

Savage, L. J. (1954). *The Foundations of Statistics*. Dover Publications.

Schaal, E. and Taschereau-Dumouchel, M. (2015). Coordinating business cycles. mimeo.

Shleifer, A. (1986). Implementing cycles. *Journal of Political Economy*, 94(6):1163–1190.

Tallarini, T. (2000). Risk-sensitive real business cycles. *Journal of Monetary Economics*, 45:507–532.

Wald, A. (1945). Statistical decision functions which minimize the maximum risk. *Annals of Mathematics*, 46(2):265–280.

Woodford, M. (2003). Imperfect common knowledge and the effects of monetary policy. In P. Aghion, R. Frydman, J. S. and Woodford, M., editors, *Knowledge, Information and Expectations in Modern Macroeconomics*. Princeton University Press.

Woodford, M. (2010). Robustly optimal monetary policy with near-rational expectations. *American Economic Review*, 100(1):274–303.