Reflexive Expectation Formation

Timo Ehrig [*] Jürgen Jost[†‡]

Paper to be presented at the meeting of the American Economic Association in Chicago, January 2012

Preliminary Draft Version

Abstract

How do economic agents form expectations regarding prices, when there is a fundamental uncertainty about the market process, like in or after financial crises? How does the formation of expectations fold back to the realized economic process, and in particular, to the selection of one of multiple possible equilibria in the economic process? We present a game theoretical model in which these questions can be studied in a systematic way.

We argue that financial expectation formation can be usefully studied as a game-like situation of mutual anticipation of economic agents, if the key uncertainties that drive the development are endogenous, that is, if there is uncertainty about the actions of other economic agents. For instance, after the financial crisis of 2008, there was uncertainty about how governments, banks, and firms would behave, and the anticipation of this behavior was key to the formation of financial expectations of individual agents.

In our model, individual agents entertain higher order beliefs regarding the expectations of other economic agents, which are the basis for their own expectation formation. These beliefs are updated in a learning process. The learning process is more complex as in REE models: Prices also contain endogenously generated information regarding the mutual expectations of different agents. Agents cannot fully distinguish if price movements are caused by a change in information about fundamentals, or a change in expectations of the other agents.

[*]Max Planck Institute for Mathematics in the Sciences, Leipzig, Germany

[†]Max Planck Institute for Mathematics in the Sciences, Leipzig, Germany

[‡]Santa Fe Institute, New Mexico, USA

1 Introduction

How do economic agents form expectations regarding asset prices and the development of macroeconomic quantities, when there is a fundamental uncertainty about the economic process, like after financial crises? How does this formation fold back to the realized economic process, and in particular, the selection of equilibria in the economic process? These questions are pressing, in particular after the financial crisis of 2008, and even today, as we do not yet know if we are in a stable economic trajectory.

In our view, a key shortcoming of the currently established modeling paradigms to study macroeconomic development, that is, Dynamic Stochastic General Equilibrium and Rational Expectations models, is that here price movements are assumed to be caused by events exogenous to the economy, but not by endogenously caused by changes in the higher order beliefs of the interacting agents. In many important contemporary market situations, the key uncertainties are generated endogenously, that is, agents are mutually uncertain about the actions of other agents that will determine future supply and demand conditions, and this results in mutual anticipation. Consequently, when forming expectations regarding asset prices, and macroeconomic quantities, agents need not only take into account their belief regarding fundamental values, but also regarding the expectations of other economic agents.

We hence suggest to study expectation formation as a game-like situation of mutual anticipation of economic agents. This thought has already been articulated by some macroeconomists (like in Morris and Shin, 2000). However, we argue that the analysis by Morris and Shin is not reaching far enough. In particular, Morris and Shin do not discuss how prices mediate the respective expectation formation and mutual anticipation of interacting agents.

'Global games' (ibid.) are a formal structure in which individual agents face coordination risks. However, a particular assumption regarding the structure of uncertainty is made in the global games literature. The assumption is that uncertainty is captured in a normally distributed random variable with noise that has the same structure for all individual agents. The mathematical reason why these assumptions are made is that then one can derive unique equilibriums using the Nash equilibrium as a prediction device.

One question of essential importance is if we can deviate from the Nash equilibrium approach to achieve predictability in macroeconomic models in which higher order beliefs play a central role.

In our own modeling we wish to understand how the history of the economy, and the corresponding alignment of higher order beliefs, brings about predictability. To be able to do this, we draw on epistemic game theory (developed by Adam Brandenburger and others, see e.g. [5]). We study price formation in a strategic market game setting. The role of higher order beliefs in price formation is not fully understood in theory. By drawing on strategic market games rather than a general equilibrium approach, we are able to make the agents aware of how they will reveal information by their trading behavior. Moreover, we are able to express almost any structure of uncertainty in our model, generalizing from the particular assumptions made in the global games literature. Another salient feature of the real world that we can study explicitly in our model is that some smart agents (like central banks) may try to influence the beliefs of other agents by their buying/ selling behavior, and plan ahead how the beliefs of others may evolve.

That prices mediate information is, of course, a commonplace. However, in REE models, the mediation by prices is oversimplified: By making the (rather strong) assumption of common knowledge of the price function, in these models, prices fully reveal the information of the participating agents. In reality, we argue, prices mediate expectations, but in far more complex ways. Prices reflect expectations of the interacting economic agents, but these expectations reflect if the market participants anticipate pessimistic or optimistic equilibria in the economic process, and how uncertain they are regarding the actions of others in the future. In some cases, some agents may even actively influence prices to shape (and even misguide) the expectation formation of other agents. In our model, prices hence become more informative than in REE models: Prices also contain endogenously generated information regarding the mutual expectations of different agents. The theoretical perspective that prices contain endogenous information, information regarding the mutual beliefs of interacting economic agents, is, to our knowledge, new to the literature.

We argue that this theoretical perspective promises to be helpful to analyze a range of pressing questions in a new way. Take the example of interest rates for government bonds of weak Euro countries like Ireland or Greece. The interest for these papers reflect the fear of default, but default is not an exogenous event. Default of government bonds is a complex outcome of actions of many participating agents. For instance, whether or not other Euro countries will support the weak countries heavily influences the possibility of default. But a decision whether to support or not is not an exogenous event, but an action from a deliberate government, who also anticipates the other participating agents. Hence, the interest rates for government bonds of weak countries are the outcome of mutual anticipation of different agents.

It is often remarked that severe depressions are *not* primarily caused by exogenous events, like natural catastrophes, but rather by collective uncertainty regarding the expectations of other economic agents (see, for instance, Akerlof and Shiller, 2009).

Furthermore, consider the following example of verbalized reasoning. As a plausible example how practitioners may think, consider the following quote from the newspaper "the Economist " in 2010: "Interest rates will stay low only if growth remains slow. But if economies grow slowly, then profits will not rise fast enough to justify current share prices and incomes will not rise far enough to justify the prevailing level of house prices. If, on the other hand, the markets are right about the prospects for economic growth, and the current recovery is sustained, then governments will react by cutting off the supply of cheap money later this year." ("Bubble warning", The economist, Jan 7, 2010)

The journalist considers conditions of future supply and demand, about the expectations of other economic agents ("the market"), and about possible strategic actions ("governments cutting off the supply of cheap money"). The journalist also wondered if the majority of other agents ("the market") is right in terms of its expectations on growth.

The journalists reasoning is clearly a kind of boundedly rational reasoning. However, in some sense, his reasoning (and the resulting expectations) actually involve more complex reasoning as would be assumed to be in rational expectation models. In particular, he is aware that all economic agents have to cope with only partial awareness about future supply and demand conditions, and that, in consequence, agents strategically co-ordinate their expectations.

What are the dynamics in an economy in which all agents strategically anticipate their peers, if we picture all agents reasoning like the journalist in the quote above? Below, we present a model to analyze this question.

In our model, agents form expectations on the basis of mutual anticipation of their actions. They entertain beliefs on beliefs of the other agents. But moreover, agents plan ahead how their own actions will alter the beliefs of other agents over time.

With our model, we offer a structure to study the following questions in a systematic way: 1. What can agents read, theoretically, from prices, when knowledge is incomplete? 2. We allow the beliefs of the agents to be a priori inconsistent. Will the expectations of the agents become consistent over time?

In the literature, there exist several sophisticated models how agents can model the expectation of the market, that is the aggregate expectations of the collective of other agents. In this context, in particular the effect of random perturbations or shocks has been studied, with mathematical tools from the theory of stochastic processes. In existing models in which forecasting the forecasts of others is a key problem (Townsend, [36],Romer [31], Lorenzoni, [24]), forecasting others boils down to a problem of learning about information diffusion. In those models, agents have to solve an optimal inference problem. The key point is that price movements, e.g., as a result of external perturbations, reveal more information than absolute prices only, in particular about fundamental information that others may possess. At equilibrium, the expectations of the agents are consistent with the market that they generate. However, the possibility for the agents to solve their optimal inference problem hinges upon a general agreement on the uncertainties that unfold while the economy evolves. We argue that this general agreement often does not exist. Consequently, the mutual anticipation of the agents may, or may not be resolved in a self-consistent way.

We are (in long term perspective) motivated by a search for an alternative theory of economic expectation building. In the recent discussion, macroeconomic models based on dynamic stochastic general equilibrium models (that are based on the idea of rational expectations) were subject to criticism (Krugman, 2009, [23]; Colander et al. 2008 [9]). However, few theoretical alternatives to rational expectation models have been proposed. The suggestions in [9] are to use agent-based modeling, or to follow an "engineering" approach to macroeconomics. The suggestion to use agent based simulations is valuable, but we see the danger that such modeling is no longer analytically tractable, and it might be difficult to develop a coherent theory based on agent based simulation. In an "engineering" approach to policy making, one would dispense itself of a micro-foundation of the actions of economic agents, and a potential risk is that little theoretical process is made.

To maintain analytical rigor, as in DGSE and REE models, we argue, the analysis of economic expectation formation should be based on game theoretical analysis. However, there are some theoretical assumptions from these models that we do not maintain in our own modeling. We will now discuss these deviations.

In rational expectation equilibrium models (compare e.g. Dubey (1987) for a review) prices are fully revealing of the information of the agents. Ben-Porath and Heifetz ([6]) show that prices may not fully reveal states of nature when agents entertain heterogeneous beliefs on the price function, even though rationality and market clearing are common knowledge among the market agents.

We argue that heterogenous beliefs regarding the price function are the real-world condition we need to deal with. We regard Ben-Porath and Heifetz results as important and provocative, and argue that their paper raises further important questions: If prices are not necessarily revealing information, as a given price can be consistent with different combinations of fundamental values and higher order beliefs of the interacting agents, do the higher order beliefs of interacting agents become more consistent over time, due to a learning process of the involved agents? Does this then lead to the revelation of information via prices over time? Or can higher-order uncertainties lead to severe inefficiencies in markets, and may we need additional institutions in markets to ensure that prices reflect the information that individual agents posses?

We hence study an economy in which prices are, like in the Ben-Porath and Heifetz paper, consistent with different combinations of fundamentals and higher order beliefs. Our contribution is to study an economy in which the agents anticipate the (putative) evolution of beliefs of other agents, and base their own present decisions on this anticipation. This dynamic setting allows us to investigate under which conditions higher order beliefs get aligned (and if prices then become information revealing) or if in some cases higher order uncertainties lead to severe coordination failures.

We believe that this study is important, because it has consequences for policy-making. In particular, if prices are not necessarily fully revealing information this implies that economic coordination is not necessarily efficient.

In our own model, we maintain the assumption of full rationality. That is, in our model, rationality as defined in Savage (1954) of all agents is common knowledge among the agents. But as has been pointed out by other scholars already (compare e.g. Angeletos (2011), Ben-Porath and Heifetz (2010) even given common knowledge of rationality the expectation formation of economic agents can be complex, and lead to endogenous fluctuations in the economic process. Endogenous uncertainties and fluctuations can arise even given fully rational agents, and our paper is one contribution in the emerging literature that attempts to understand how these uncertainties and fluctuations can be tamed.

Before we proceed, we briefly discuss the key technical contributions of this paper.

While market interaction and price building processes have already been studied as strategic market games, incomplete knowledge of the interacting agents has so far only been studied in the spirit of Harsanyi (1967). While Harsanyi's seminal contribution allowed huge steps toward analyzing market interaction given informational uncertainty, we argue that we need to move beyond Harsanyi to fulfill the promises we made above.

In Harsanyi's 'Bayesian games', *the structure of the uncertainty that is to be resolved* is assumed to be common knowledge. Even though agents are allowed to entertain different degrees of uncertainty, in Harsanyi (1967) each agent's uncertainty needs to be expressible in form of a partition of a state space, *and the agents need to entertain identical prior belief about the 'objective' probability distribution over the singleton elements in the state space.* By making this assumption, agents' beliefs are *by construction* consistent with the process that they generate. We think that a characterizing element of real economies is that the agents would not be able to agree on a prior distribution over fundamental values.

The assumptions of Harsanyi have already been critically addressed in the literature, for instance, in Kadane ([21]). In our approach, agents do share some knowledge about the economy (the equations of the model), but the disagreement on the probability distributions over fundamentals makes the expectation formation problem of the agents more complex. In particular, they need to plan ahead in their assumptions regarding the evolution of the beliefs of the other agents.

The paper will continue as follows. In the next section, we will present our model. In the subsequent section, we present an older version of our model, which we include, as here some ideas, like the anticipation of the agents regarding the evolution of the beliefs of other agents are sketched which are not yet included in the newer version of the model. We will then discuss our model, and finally conclude the paper.

2 A Model of Reflexive Expectation Formation

2.1 A strategic market game

The basic model works with integer time t and contains a financial institution, called, for simplicity, 'bank', that lives as long as it does not default, and at each t, N short term creditors, called, again for simplicity, 'investors', that lend money to the bank which the bank then returns with interest to them at $t+1$, unless it defaults. When the bank defaults at time $t+1$, the investors from time t loose their money, and the bank foregoes its opportunity for future profits. In order not to default at $t+1$, the bank must find sufficiently many new investors so that the bank can pay out the previous ones.

We model the interaction between the bank and the investors as a strategic market game in the sense of [12].

At time t, the bank (agent 0) chooses how much to pay back at time $t+1$ for the short term credit issued to her at time t. The amount is called β^{t+1}. A choice of the quantity of this amount is the bank's strategy at time one. $S_0^t = \{\beta^{t+1} \in \mathbb{R}^+\}$

At time t, an investor i choses how much of her endowment to spend to buy short term debt, given the anticipation of a price p^t (to be defined below), which it will only learn about ex post. There are N investors. The supply of money from selling short term debt to the bank at t is $\sigma^t = \sum_i \sigma_i^t$. Here, the constraint is that the investment σ_i^t of agent i cannot exceed her endowment (which will be described below). This defines i's strategy space S_i^t. The investors and the bank choose the corresponding variables σ_i^t and β^{t+1}, resp., independently of each other. The price of short term debt then is

$$p^t = \frac{\beta^{t+1}}{\sigma^t}. \tag{1}$$

The precise market mechanism is quite important for any epistemic analysis: who sees what and when? what are the precise information flows? In the paper of Shin [34] it is made clear that the choice of the market mechanism will drive higher order dynamics. Without entering into these details here, we only point out that the model assumption employed here captures the idea that the bank has to offer its investors the prospect of a sufficient future gain, and the business becomes more costly for the bank the fewer investment money it can generate. At the moment in Europe, this is not the way actual credit markets function, as the unsecured inter-bank market almost dried up and is replaced by the collateralized interbank market and increased lending of the central bank to the different financial institutions. However, the model may regain relevance once this market is again in place. In any case, the model is a theoretical study of the mediation of aggregate coordination risks by prices, a theoretical study thus far missing in the literature.

Thus, the investors, having learned p^{t-1}, decide about their investments σ_i^t before knowing p^t. Likewise, at time t, the bank knows p^{t-1}, and since (in contrast to the investors) it also knows β^t, it can deduce the total investment volume σ^{t-1}. When deciding β^{t+1}, it does not yet know p^t either.

The balance of the bank at t is

$$b^t := \rho(t)b^{t-1} + \sigma^t - \beta^t, \tag{2}$$

where $\rho(t)$ is a random factor whose distribution will be discussed below. The precise value of $\rho(t)$ is not revealed to the investors, but they know its distribution. The bank learns the value of $\rho(t)$ at time t, but does not know it before t; again, of course, the bank knows its distribution.

Thus, the bank receives σ^t from its investors, pays out β^t to the investors from the previous time period, and earns money in proportion to its cash reserves from previous periods.

The bank defaults at time t when $b^t < 0$. If the bank defaults at t, then by definition, $R^\tau = 1$ for all $\tau \geq t$, otherwise $R^t = 0$.

The investors are short-lived, they enter the market at t and consume and leave at $t+1$. Each investor i at time t has an endowment $\rho_I(t)e_i$ in the time period that she is born. Here, $\rho_I(t)$ is simply a suitable growth factor to make the model consistent. We assume that this factor is the same for all investors. It may be interpreted as some kind of macroeconomic growth rate. We shall also need this factor's expectation value $\bar{\rho}_I(t)$. E_i is the space of possible discounted endowments of agent i, $E = E_1 \times \cdots \times E_N$ is the space of possible discounted endowments of all agents.

The expected payoff of investor i at $t+1$ is

$$\pi_i^{t+1} = (1 - R^t)p^t\sigma_i^t - \bar{\rho}_I(t+1)\sigma_i^t. \tag{3}$$

The payoff measures the reward or loss of the investor who has made a decision what to invest on the basis of the observed history of the play at t.

The investor, of course, has the alternative not to invest at all in which case her pay-off would be 0. Thus, she will invest only if she believes that the bank will not default ($R^{t+1} = 0$) and that

$p^t > 1$. Thus, the bank needs to repay more at $t+1$ than it has collected at t, i.e. $\beta^{t+1} > \sigma^t$. The bank can do that, however, only when it collects more than it pays at time t, i.e., $\sigma^t > \beta^t$. Therefore, the balance of the bank needs to grow over time, and the bank thus needs a positive factor $\rho(t)$ in (2).

At each time t, the aim of the bank is to maximize its profits b^{t+1} at $t+1$. We assume here that the bank does not possess a longer time horizon than 1, for instance, because that is the period of the term of office of its CEO.

We assume that all equations of our model, and the utility functions, *the distribution of growth factors $\rho(t)_I$ and $\rho(t)$, the complete price history* and the agents' endowments are common knowledge.

*What is **not** common knowledge are the $\rho(t)$'s and the b^t's. We will make use of epistemic game theory to analyze the corresponding mutual uncertainty. We will discuss in what ways this analysis is more general than the analysis in global games.*

2.2 Strategy for Analysis

In the following, we will analyze the game with the help of epistemic game theory (Brandenburger 2007, [5]). A basic assumption of such an approach is that every agent can put himself into anybody else's perspective and reason from that perspective on the basis of putative assumptions. Assumptions of common knowledge of rationality or of common belief in it can serve to make the reasonings of the different participants consistent and constrain and align them, but this need not necessarily be the case.

An important feature of our epistemic analysis will be the coupling of different times in the optimization schemes and belief dynamics of the agents. Our epistemic analysis will not be conclusive, that is, we shall not be able to deduce a unique collection of rational and consistent actions. We consider this, however, not as a deficit of our model, but rather as an aspect to be expected also in more realistic settings. *More precisely, we assume that for an analysis of macroeconomic coordination problems to be conclusive, we need to make additional assumptions on the current cognitive hypotheses of the interacting agents. For a model to be predictive, we argue, we need to include assumptions of the cognitive hypothesis into the description of the game that is a model of the macroeconomic coordination. The cognitive hypotheses cannot be derived from common knowledge* *about certain aspects of the interaction situation, as we will discuss.*

The essential questions for the agents are the following. For i, a time t agent:

1. Will there be enough supply of money from investors at both times t and $t+1$ to ensure that the bank will not default at either period?

2. How many other investors will be willing to invest at time t and how much of their endowment are they willing to invest, that is, what is the price p^t that will determine my gain at $t+1$ in the non-default case?

For the bank:

3. How should I set β^{t+1} in order to ensure that on one hand, enough investors will think investment to be profitable and that on the other hand I will earn a profit? Here the challenge is the following. The CEO at t will only influence the price that is seen ex post for investors at t; he will influence the ex ante expectations of the investors at $t+1$. The investors' expectations at t are influenced by their predecessors. However, the CEO at t only cares about the performance of the bank at $t+1$, in the same manner as the CEO at $t-1$ only cared about the performance of the bank at t, etc. What do we assume that the CEO at t knows? He should know b^{t-1} and $\rho(t)$. About the risk of default at t, this CEO can do nothing about. He can only hope that his predecessors cared about it, but he can neither influence $\rho(t)$ nor σ^t, as he cannot give any information to the investors at t. His choice of β^{t+1} will influence the thinking of investors at $t+1$, but not at t. Set up this way, the only reason for a CEO at t not to set β^{t+1} to zero is the following: He knows that the investors know the stochasticity of $\rho(t)$. Thus, he knows that if he chooses a β^{t+1} which is very low (which would only be 'fair' given the unlikely case of a very positive $\rho(t)$, investors will observe a low price and not invest at $t+1$, which increases his default risk (and this CEO cares about the balance at $t+1$, as this is what is payoff relevant for him. However, if this CEO knows that b^t is actually quite high, and makes a bet that even given potential losses from $\rho(t+1)$, he might well set β^{t+1} to zero, as so, he has no 'cost' of paying back at $t+1$. Of course, doing that, he also takes into account that investors at $t+1$ will not invest, leaving no additional short term income to him.

Clearly, these questions are interlinked, and an epistemic analysis needs to take this into account.

We have set up the model in such a manner that it is invariant under time shift (after discounting the growth factor inherent in the model), and therefore at each time t, there is a baseline epistemic scenario: Each investor i at t reasons that she is in the same structural situation as investors at previous times, and that the investors at $t+1$ (on whose willingness to invest it depends that i will get her pay-off) are likewise in the same structural situation. Also, this then also holds for all her coinvestors at time t. Thus, since the scheme has worked in the past, i.e., the bank has not defaulted, and investors had invested because from their prior history they had expected a pay-off, it should continue to work now and in the future, and therefore, every investor now and in the future should invest in the same manner that past investors had. And since i believes that everybody will reason in the same manner as she herself, she should go ahead and invest. If, indeed, every investor reasons in that way, this is self-confirming and self-fulfilling.

Likewise, the bank CEO at t may simply assume that his predecessors in the past had already optimized the scheme and that the investors will behave in the same manner as their ancestors in the past did, and so, he should choose β^{t+1} such that $p^t = p^{t-1}$.

This baseline epistemic analysis needs the assumption that the distribution of the stochastic growth factor $\rho(t)$ in (2) is invariant and uncorrelated over time, that is, the $\rho(t)$ are i.i.d.

We assume that the bank knows the value of $\rho(t)$ at time t, that is, when specifying β^{t+1}, but not earlier.

The other agents are not told that value, but may attempt to infer it from p^t.

So, let us put into the position of the bank CEO at time t. According to (2), he wishes to maximize

$$\begin{aligned} b^{t+1} &= \rho(t+1)b^t + \sigma^{t+1} - \beta^{t+1} \\ &= \rho(t+1)(\rho(t)b^{t-1} + \sigma^t - \beta^t) + \sigma^{t+1} - \beta^{t+1} \\ &= -\beta^{t+1} + \sigma^{t+1} + \rho(t+1)\sigma^t - \rho(t+1)\beta^t + \rho(t+1)\rho(t)b^{t-1}. \end{aligned} \tag{4}$$

In this equation, the CEO at time t chooses β^{t+1}, and he thereby influences the investors' investment σ^{t+1}, that is, the second term. The investment σ^t in the third term is determined independently by the investors at time t, but based on information overlapping with the one that the CEO possesses.

The remaining terms have already been determined in the past and therefore are no longer under our CEO's control. Nevertheless, they will not only contribute to his target variable b^{t+1}, but may also determine the investment decisions $\sigma_i(t)$.

Thus, the CEO has to model the investment decisions of the investors i at t and $t+1$.

b^t will grow in time at a higher rate than ρ as long as $\sigma^t > \beta^t$. And when $\rho > \rho_I$, any such positive difference $\sigma^t - \beta^t$ will make a growth contribution at a higher rate thant the investors can achieve by alternative schemes.

2.3 Analysis of possible growth dynamics

What follows is an analysis of the optimal strategies for the investors given that there is no risk of default and that they knew β^{t+1}. While this assumption does not hold in our model, the following analysis nevertheless should be helpful to understand the dynamics of the model.

To understand the default risk dynamics we need to understand the mutual anticipation of the agents which we will analyze which the help of epistemic game theory below. However, the analysis that follows here (for the risk free case), is an important preparation for this analysis. As our agents are all assumed to know the equations of the model, and know that the equations are common knowledge, the agents are assumed to be able to do the analysis of the risk free case that we do here. Thus, we can make use of some of our results from analyzing the risk free case in the analysis of the mutual anticipation of the agents below.

In the situation where i at t expects no default, by (1) and (3), we have the target variable $\pi_i^{t+1} = \beta^{t+1}\frac{\sigma_i^t}{\sum_j \sigma_j^t} - \bar{\rho}_I(t+1)\sigma_i^t$. This should be positive for an investment to make sense, and the investor hopes to achieve the optimal value of π_i^{t+1}. Since the investment decision of i is σ_i^t, we need to evaluate the derivative

$$\frac{d\pi_i^{t+1}}{d\sigma_i^t} = \frac{\beta^{t+1}}{\sigma^t}\frac{\sum_{k\neq i}\sigma_k^t}{\sum_j \sigma_j^t} - \bar{\rho}_I(t+1) \tag{5}$$

which in the case when all N investors invest the same amount becomes

$$= \frac{\beta^{t+1}}{\sigma^t}\frac{N-1}{N} - \bar{\rho}_I(t+1). \tag{6}$$

At investment optimum, this has to vanish, which is the case iff

$$\beta^{t+1} = \frac{N}{N-1}\bar{\rho}_I(t+1)\sigma^t. \tag{7}$$

Thus, if all investors picked their optimal investment volume in response to a given β^{t+1}, the following price would result:

$$p_{\text{eq}}^t = \frac{N}{N-1}\bar{\rho}_I(t+1). \tag{8}$$

Thus, by (7), when β^{t+1} is larger than $\frac{N}{N-1}$ times the total endowment of all (assumed to be equally endowed) investors, they may all wish to invest all the money they have in the bank *given that there is no risk of default of the bank.* If $p^t_{\min} := \bar{\rho}_I(t+1)\sigma^t < \beta^{t+1}$, investment is already desirable for the investors, as they receive a positive pay-off, but if $\beta^{t+1} < \frac{N}{N-1}\bar{\rho}_I(t+1)\sigma^t$, then they can increase their pay-off by reducing their investment. In this case, a coordination problem between the investors arises. For instance, they could all invest the same share of their money until equality is reached in (7). Or some could invest all their money while others completely abstain. Thus, this coordination obviously does not possess a unique solution. For the CEO, however, what is of interest is simply the existence and stability of a solution, but not the distribution of investments between the various investors. Of course, it may be possible that the investors are not able to solve this coordination problem, resulting in a lower price.

From (8), we see that the equilibrium price has to grow at a higher rate than the economy growth rate ρ_I. But this is a finite size effect that disappears in the limit $N \to \infty$.

In the general case, the vanishing of (5) leads to the equation

$$\sigma^t_i = \sqrt{\sum_{k\neq i} \sigma^t_k} \left(\sqrt{\frac{\beta^{t+1}}{\bar{\rho}_I(t+1)}} - \sqrt{\sum_{k\neq i} \sigma^t_k} \right). \tag{9}$$

When N is large and hence the influence of each investor is small, this is approximated by

$$\sigma^t_i = \sqrt{\sigma^t} \left(\sqrt{\frac{\beta^{t+1}}{\bar{\rho}_I(t+1)}} - \sqrt{\sigma^t} \right). \tag{10}$$

Thus, when the adjusted growth rate of β exceeds that of σ, the investors will like to invest. While the investors at time t do not know the actual β^{t+1}, they may try to infer a trend from observed past price movements and base their investment decisions on an extrapolation of that trend.

If $\sigma^\tau = \beta^\tau$ for all times τ, then the pay-off of the bank is always zero, but it does not default. If in this case, we have (7), then

$$\sigma^{t+1} = \frac{N}{N-1}\bar{\rho}_I(t+1)\sigma^t, \tag{11}$$

and the investments grow geometrically in time. When the bank wants to make a profit at the equilibrium (8), then, by looking at the first four times in (4), we must have

$$\sigma^{t+1} + \rho(t+1)\sigma^t > \frac{N}{N-1}\bar{\rho}_I(t+1)(\sigma^t + \rho(t)\sigma^{t-1}), \tag{12}$$

that is, the investments have to grow at an even higher rate.

Of course, this is sustainable only in the asymptotic limit $N \to \infty$, as for finite N it is not matched by the growth of the endowments of the investors.

To summarize: If the β's grow slower than the σ's, it cannot be a sustainable strategy over time to invest. As long as the endowments of the investors are not fully invested and as long as the mean of the $\rho(t)$ is bigger than the mean of the $\rho_I(t)$, the value pie can be enlarged and there are more profits to be distributed among bank and investors. But if endowments are fully invested, it is not possible that both bank and investors sustainable profit from trade exceeding their own natural growth rates.

In the case where the investments also grow geometrically at a rate of $\frac{N}{N-1}\bar{\rho}_I(t)$ and where everything is invested at time t with the condition (7), the bank might be able to absorb a shortfall $\sigma^{t+1} < \beta^{t+1}$ if there are some earnings from previous periods left, i.e., $b^t > 0$, and if the random factor $\rho(t+1)$ is sufficiently large. If not, the bank may get into trouble. Now, if in this situation, one of the investors starts to worry about $\rho(t+1)$ being too small, she may get cold feet and decide to not invest. In that case, σ^{t+1} becomes still smaller, and the likelihood of a default increases. Thus, if one investor worries that another one might worry and will not invest, he should not invest either. In this manner, a downward cascade may set in that leads to the sure default of the bank. The bank CEO, however, may try to adjust β^{t+1} so as to influence the beliefs of the investors and induce them to invest after all.

In short, the situation for the investors is the following. They already can make a profit at any price above $p^t_{\min}$, leading to a growth rate larger than $\bar{\rho}_I$, but in order the maximize their pay-off, they may reduce their investments to get a better price. At the condition (7) that determines their optimal investment value, a growth rate of investments (11) is needed that is not sustained by the growth of the endowments. Therefore, the investors may worry that β is too high and stop investing for fear of a crash. And in fact, they may also stop because they fear that somebody else worries about the possibility of a crash. In contrast, when β is too low, first of all, their investment may not pay enough, and secondly, they may worry that the bank will get into trouble and default. And, of course, they may fear that others worry about this possibility.

If an investor knew β^{t+1} and the investments of her colleagues, she could determine her optimal investment σ^t_i from the equilibrium growth condition (7). For the bank CEO, the situation is different, as his pay-off will depend on the future investments σ^{t+1}, and he can only try to influence the expectations of future investors by choosing his current β^{t+1}.

2.4 Cognitive Dynamics and Epistemic Analysis

So far, we analyzed the case when the b^t's are high, and there is no risk of default. Then we are in a standard market regime. However, when the b^t's are small, and there is the risk of default, the dynamics get far more complex. To analyze these dynamics, we make use of epistemic game theory.

We consider continuous mappings of the form:

$$\lambda_i^t : H_i^t \to M(p^t \times b^{t-1} \times \rho(t) \times p^{t+1} \times b^t \times \rho(t+1) \times b^{t+1} \times S^t \times S^{t+1} \times H^{t,-i} \times H^{t+1,-i}) \quad (13)$$

Technically spoken, hypotheses $h_i^t \in H_i^t$ are types as in epistemic game theory. However, it may be a better wording to call them the hypotheses of the agents. M is a probability measure. Given a compact metrizable space ω, write $M(\omega)$ for the space of all Borel probability measures on ω, where $M(\omega)$ is endowed with the topology of weak convergence (and so is again compact metrizable).

Each H_i^t is a compact metrizable space. Members of H_i^t are called hypotheses for player i at time t. Members of $p^t \times b^{t-1} \times \rho(t) \times p^{t+1} \times b^t \times \rho(t+1) \times b^{t+1} \times S^t \times S^{t+1} \times H^{t,-i} \times H^{t+1,-i}$ are called states (of the world) *at t*.

The hypotheses express that each agent at t needs to reason about all other agents at both t and t+1 in deriving their own strategies. The hypotheses express all possible configurations of beliefs of the agents, both about relevant quantities of the world, and about the respective beliefs of other agents, that is, their cognitive states from which they will then derive their strategies. Epistemic game theory helps us to analyze which cognitive states, together with additional assumptions like common knowledge of rationality, are consistent with which actions of the players. Using the type mapping, we can express all possible configurations of higher order beliefs.

These will be helpful in our analysis. For instance, the bank has a belief over the strategies of the investors at time t (take the marginalize the type of the bank to S^t to get this belief). Thus, it has a belief about the price that will come about at t, conditional on a choice of $beta^{t+1}$ that the bank executes. Thus, the bank can plan ahead which price p^t it believes will be likely to be visible. Now the bank also has a prior belief on how the investors at $t+1$ will react to the observation of this price. Marginalize the type of the bank at t to $p^t \times H^{t+1,-i}$ to get this belief. All these information can be conveniently coded in the type construction that we outline above. For general information about this construction, consult Brandenburger (2007). Note that we can also encode in this hypotheses space how the bank believes to be able to strategically influence the hypotheses of the investors at time t+1 by manipulating the price. More general, the hypotheses space construction enables us to express and analyzes reflexive dynamics among hypotheses of different market participants, which is, to our knowledge, a novelty to the literature.

But the investors, in our model, are similarly smart agents as the bank. The investors can also entertain beliefs how to strategically influence the beliefs of the CEO at t+1 by manipulating the price at t. However, there is an asymmetry in the model, which is that the bank at t knows about b^{t-1} and ρ^t, while the investors do not know about this.

However, epistemic game theory does not tell us where the concrete hypotheses of the agents come from. We need to make additional assumptions at this point. From our analysis of the case without default above, we can already take some useful hints for this analysis.

2.5 Analysis of a specific case

This section is draft and incomplete.

We analyze the following specific case. The $\rho(t)$ are uniformly distributed on the interval $[0.6 - 1.6]$. The rho_I is fixed to 1.05.

We assume that for the time periods up to $t = 0$, $\beta^t = \sigma^t$ and $b^t = 0$. We assume that it is common knowledge among investors that $b^t = 0$ for all time periods up to time zero.

We definite $\beta^1 = \rho_I \sigma^0$, and σ^0 is a fixed value.

Now we assume that innovative types appear at time one. The innovative CEO at time one sets $\beta^2 = 1.02\rho_I^2\sigma^0$. The investors stay with the default strategy at time one: $\sigma_i^1 = \rho_I \sigma_i^0$. The results in a price $p^1 = 1.02\rho_I$.

At $t = 2$ the investors know that the bank will default, if they do not adjust their strategies. They know that they need to invest at least $1.02\rho_I^2\sigma_i^0$ to keep the bank alive. They also know from the assumption of common knowledge of the equations of the model and the common knowledge of the distribution of the $\rho(t)$'s and ρ_I, that 'a growth period could start now': There is room to increase overall profits from both bank and investors, as the investors are not yet fully invested with all their endowments, and the bank has a higher growth rate as the investors. However, from common knowledge of the equations of the model, all agents also know (as they can do the same analysis as in this paper under 'possible growth dynamics' that this growth period is not sustainable forever.

However, the investors do not know exactly when the endowments will be fully invested, as they can only observe the price and cannot fully deduce the σ^t's from the observation, as they have a lack of knowledge of the b^t from time one on (as the process of the $\rho(t)$ is stochastic, they cannot deduce the value of the b^t's from their knowledge that $b^t = 0$ for the periods up to $t = 0$).

Thus they do know know exactly when the growth dynamic will break down. They also don't know the precise value of the balance cushion of the bank that it might accumulate over time.

In the type space, there is an implicit recursion on all hypotheses until the infinite future. The agents thus reason through all time periods.

2.6 What can the agents read from prices

To be completed. Please watch out for a new version of the paper.

3 Discussion and Conclusion

How do economic agents form expectations regarding prices, when there is a fundamental uncertainty about the market process, like after financial crises? How does the formation of expectations fold back to the realized economic process, and in particular, to the selection of one of multiple possible equilibria in the economic process?

The main contribution of this paper is to present a model in which these questions can be addressed in a systematic way. We provide a structure for conceptualizing and understanding how the fundamental economic world of achieved trades and production and the "cognitive" world of the (mutual) expectations of the economic agents interact in a dynamical manner.

In our model, individual agents entertain higher order beliefs regarding the expectations of other economic agents, which are the basis for their own expectation formation.

The agents can observe prices as the aggregate result of the collective activity of all agents, but they do not know the full details about other agents. The agents hence learn indirectly about the expectations of the other agents by observing prices and their movements over time.

We argue that in context of global crisis such as the financial crisis of 2008, information about the likelihood that the economic process favors one equilibrium (for instance, an optimistic one) over another (for instance, a pessimistic one) is often more important than information about fundamentals. This type of information is endogenous, because it is basically information about the higher order beliefs of the other agents. With our model, we offer a structure to study how this type of information can be extracted from prices.

With our model, from a technical point of view, we also contribute to the literature on learning in games. In the literature, most models of learning in games study situations in which the actions of other agents can be observed. In our model, the agents cannot observe the actions of other agents, but they learn about the actions of others indirectly via the observation of aggregate quantities like prices.

In real life, we argue, beliefs of economic agents are often inconsistent. The model helps to study such situations, and if and how beliefs become consistent over time via the common observation of prices.

Uncertainty in higher order beliefs can cause inefficiencies in the coordination of the agents. Are such uncertainties one cause of inefficiencies after crises, when the market is 'confused'?

We see these questions as crucial for developing a more realistic understanding of current real world economic phenomena. We hope that our model can provide a basis for the study of these questions.

References

[1] Akerlof, G. 1970. The Market for "Lemons": Quality Uncertainty and the Market Mechanism. The Quarterly Journal of Economics, Vol. 84, No. 3, pp. 488-500

[2] Akerlof, G. Shiller, R. 2009. Animal Spirits. Princeton: Princeton University Press.

[3] Angeletos, G.M., La'O, J. 2011. Decentralization, Communication, and the Origins of Fluctuations. Working Paper. MIT.

[4] Arthur, B. 1994. Inductive reasoning and bounded rationality. American Economic Review, Vol.84, pp.406-411

[5] Brandenburger, B. 2007, ”The Power of Paradox: Some Recent Developments in Interactive Epistemology” International Journal of Game Theory, Vol. 35, 2007, 465-492

[6] Ben-Porath, E., Heifetz, A. 2010. Common Knowledge of Rationality and Market Clearing in Economies with Asymmetric Information. Northwestern University: Discussion Paper.

[7] Bernheim, D. 1984. Rationalizable Strategic Behavior. Econometrica, Vol. 52, No. 4, pp. 1007-1028.

[8] Carlsson, H., van Damme, E. 1993. Global Games and Equilibrium Selection. Econometrica, Vol.61, pp.989-1018

[9] Colander, D., Howitt, P., Kirman, A., Leijonhufvud, A., and Mehrling, P. 2008. Beyond DSGE Models: Toward an Empirically Based Macroeconomics. American Economic Review: Papers & Proceedings 2008, 98:2, 236240

[10] Debreu, G. 1959. Theory of Value. Yale: Yale University Press.

[11] Diamond, P., Fudenberg, D. 1989. Rational Expectations Business Cycles in Search Equilibrium. The Journal of Political Economy, Vol. 97, No. 3, pp. 606-619.

[12] Dubey, P. Geanakoplos, J. and Shubik, M. 1987 The relevation of information in strategic market games. Journal of Mathematical Economics 16, pp 105-137.

[13] Ehrig, T. Jost, J. 2010a. Strategic Rational Expectations. Conference Paper. Midwest Theory Meeting, Northwestern University: Chicago.

[14] Ehrig, T. Jost, J. 2010b. A Game Theoretical Approach to Expectation Formation in Markets. Conference Paper. Symposium of the Spanish Economic Association, Madrid.

[15] Farmer, D. and Geanakoplos, J., The Virtues and Vices of Equilibrium and the Future of Financial Economics, Complexity, Vol. 14, no. 3, 2008

[16] Feinberg, Y. 2005. Games with Unawareness. Working Paper. Stanford: Stanford Business School.

[17] Foster, S. Young, P. 2003. Learning, Hypothesis Testing and Nash Equilbrium. Games and Economic Behavior Volume 45, Issue 1, Pages 73-96.

[18] Fudenberg, D. and Levine, D., 1998. The Theory of Learning in Games. MIT Press: Cambridge, MA.

[19] Geanakoplos, J. 2004. The Arrow-Debreu Model of General Equilibrium. Cowles Foundation Paper. Yale: Yale University.

[20] Harsanyi, J. 1967. Games with Incomplete Information Played by 'Bayesian' Players, Part I. The Basic Model. Management Science, 14(3), pp. 159-82.

[21] Kadane, J and Larkey, P. 1982. Subjective Probability and the Theory of Games. Management Science 28 (2): 113-120.

[22] Kawamura, E. 2005. Competitive equilibrium with unawareness in economies with production. Journal of Economic Theory 121, pages 167191

[23] Krugman, P. 2009. How did Economists get it so wrong?, New York Times, September 2, 2009

[24] Lorenzoni, G. A Theory of Demand Shocks. American Economic Review 2009, 99:5, 20502084

[25] Makowski, L., Ostroy, J. 1995. A revision of the first theorem of welfare economics, The American Economic Review, 85, 808-827

[26] Malevergne, Y., Sornette, D. 2006. Extreme financial risks, Springer

[27] Milgrom, P. Stokey, N. 1980. Information, Trade and Common Knowledge. Journal of Economic Theory, 26, 17-27, 1982

[28] Morris, S., Shin, H. 2000. Rethinking Multiple Equilibria in Macroeconomic Modeling. NBER Macroeconomics Annual, Vol. 15, pp. 139-161.

[29] Morris, S., Shin, H. 2006. Global Games: Theory and Applications. Chapter 3 (pp.57-114) in: Cambridge Collections Online, Cambr.Univ.Press

[30] Muth, J. 1961. Rational Expectations and the Theory of Price Movements. Econometrica, Vol. 29, No. 3 (Jul., 1961), pp. 315-335

[31] Romer, D. 2006, Advanced Macroeconomics, McGraw Hill

[32] Rosu, Ioanid. 2006. Asymetric Information: Walrasian Equilibria, and Rational Expectations Equilibria. Notes. Chicago: GSB.

[33] Savage, L. 1954. The foundations of statistics. John Wiley and Sons.

[34] 1996. Shin, H. Comparing the Robustness of Trading Systems to Higher Order Uncertainty. Review of Financial Studies, 63, 39-59.

[35] Sornette, D. 2003. Why stock markets crash, Princeton Univ.Press

[36] Townsend, R. 1983. Forecasting the forecasts of others. Journal of Political Economy, Vol. 91, pp.546-588

[37] Yildiz, M. forthcoming. Walrasian Bargaining. Games and Economic Behavior