

Heterogeneous Beliefs and Trading Inefficiencies*

William Branch
University of California, Irvine

Bruce McGough
University of Oregon

July 6, 2015

Abstract

This paper introduces heterogeneous beliefs and studies the implications for the existence of monetary equilibria, and its dynamic properties, in the monetary economy of Lagos and Wright (2005). An endogenous fraction of agents hold rational expectations and the remaining agents employ an adaptive learning rule similar to Evans and Honkapohja (2001) and Brock and Hommes (1997). Our model delivers three primary results that follow from the finding that heterogeneous beliefs can destabilize a stationary monetary equilibrium and lead to non-linear dynamics bounded around the monetary steady state. First, heterogeneous beliefs can lead to equilibria that are welfare reducing due, in part, to a lower acceptance rate in decentralized meetings. Second, when buyers, who are uncertain about their beliefs, behave like a Bayesian by placing a prior on sellers' beliefs, uncertainty impacts dynamic stability and welfare. Third, the model's unique predictions provide an explanation of new findings about the acceptance rate in monetary laboratory experiments.

1 Introduction

This paper introduces heterogeneous beliefs into the workhorse Lagos and Wright (2005) model of monetary theory. The question addressed in this paper is how heterogeneous beliefs affects the existence of a monetary equilibrium and its dynamic properties. The main finding is that heterogeneous beliefs can alter the nature of trade in bilateral markets by generating dynamic monetary equilibria with distinct implications for the intensive and extensive margins of trade, welfare, and for reconciling key experimental findings.

Models of monetary economies take into account the role assets play in facilitating exchange in decentralized markets.¹ For example, in the Lagos and Wright (2005) model

*This paper has benefited from comments by the Editor, Ricardo Lagos, two anonymous referees, Guillaume Rocheteau, Ryan Baranowski and Cars Hommes. We are especially grateful to John Duffy for sharing experimental data from Duffy and Puzzello (2013).

¹Primary references include Kiyotaki and Wright (1989), Diamond (1982), Shi (1997a), Lagos and Wright (2005), Rocheteau and Wright (2005).

agents trade in bilateral meetings, while limited commitment and imperfect monitoring prevents buyers from using credit in exchange for goods. A demand for fiat money can arise in equilibrium because it facilitates exchange in these bilateral meetings.

The set of monetary equilibria depends on the fundamentals of the economy such as the available technology, preferences, search frictions, and the trading protocol in bilateral matches. The beliefs of agents are also fundamental to the model. Sellers are only willing to exchange goods for fiat money if they expect that they will be able to exchange that money for goods in later periods. Similarly, buyers will only demand fiat money if they believe that it has value in the future. The monetary equilibria commonly examined in the literature feature a (potentially) delicate coordination of beliefs, and the principal theory of how beliefs are formed is the rational expectations hypothesis. Rational expectations, while a natural benchmark, is a strong assumption that requires agents to form their expectations optimally with respect to the true underlying conditional distributions that, in turn, depend upon agents' beliefs in a self-referential way.² We relax the rational expectations assumption by introducing heterogeneous beliefs – where buyers and sellers are distributed across a variety of forecasting models, including rational expectations – into the monetary search model.

We build on the framework in Lagos and Wright (2005).³ Each period is divided into two sub periods. In the first sub period, there is a decentralized market distinguished by bilateral exchange where buyers, who can consume but cannot produce, are matched with sellers, who can produce but cannot consume. When buyers and sellers are matched, the terms of trade are determined by take-it-or-leave-it offers made by buyers. In the second sub period, there is a centralized market where buyers and sellers can both consume, trade assets and produce.

We introduce heterogeneity in beliefs by assuming that there is a fraction of agents who hold rational expectations and the remaining agents are distributed across other “boundedly rational” forecasting models. As a baseline example, we assume that agents are either rational or form forecasts using an adaptive learning rule. The adaptive learning rule is in the spirit of Marcet and Sargent (1989) and Evans and Honkapohja (2001), who argue that an economic agent should behave like a good econometrician and form forecasts from a well-specified forecasting model that adapts to account for recent data.

The structure of heterogeneous beliefs in our model is specified so that in a steady-state equilibrium all beliefs coincide and the perfect-foresight steady-state rational expectations equilibrium obtains. The range of equilibrium outcomes and of economic dynamics depends on the distribution of individuals across forecasting models. An important aspect to our model is that the distribution of heterogeneity is an endogenous object and perfect foresight remains a choice available to all agents, though they must pay a computational cost for its use. Following Brock and Hommes (1997) and Branch and Evans (2006), the fraction of agents with perfect foresight is increasing in its forecast accuracy, net of the computational cost, relative to the accuracy of the adaptive learning rule. Empirical evidence in favor of this framework for expectation formation is provided by Branch (2004) and Hommes (2013).

²Ours is not the first search theoretic paper to point out that rational expectations is a strong assumption. For instance, see Gu, Mattesini, Monnet, and Wright (2012) (p.15) “Of course, we made some strong assumptions, including our assumption of perfect foresight, or rational expectations.”

³See Nosal and Rocheteau (2011) for an excellent exposition of search based monetary models.

The main results of this paper are as follows. First, heterogeneous beliefs can lead to inefficient trading outcomes with lower welfare than would arise in a rational expectations equilibrium. In a steady state, all individuals hold identical beliefs. Thus, the inefficiencies that stem from heterogeneous beliefs arise along dynamic equilibrium paths. The novelty in this paper is that we allow beliefs, and hence trading inefficiencies, to be determined as an endogenous object of the model.

Heterogeneous beliefs alter the set of monetary equilibria. Unlike rational expectations monetary equilibria which require full coordination of beliefs, we show the existence of monetary equilibria with heterogeneous, boundedly rational beliefs. Although the steady state of the monetary equilibrium we consider coincides with the steady state under rational expectations, heterogeneous beliefs can destabilize the steady state and lead to a variety of periodic and aperiodic cycles, including complicated dynamics, that remain bounded around the steady-state monetary equilibrium. Along these dynamic equilibria, the distribution of money holdings across agents is non-degenerate and time-varying.

An interesting implication of the framework in this paper is the role that within-match common knowledge assumptions play in bilateral trade. In random, anonymous decentralized meetings it is natural to assume that buyer and seller beliefs are not common knowledge within a match. To see how this impacts trading behavior, consider a take-it-or-leave-it offer made by a buyer to a seller. Under common knowledge, the buyer would select an offer that just meets the seller's participation constraint. Without common knowledge, the buyer must form expectations about the seller's participation constraint, and this constraint depends on the seller's beliefs about the future value of money. Thus when making a take-it-or-leave-it offer, a buyer must allow for the possibility that his beliefs do not coincide with the seller's. To account for this possibility, we model buyers as Bayesians who holds priors over sellers' beliefs and, therefore, whom explicitly acknowledge their uncertainty about the sellers' participation constraints. We call the bargaining offers that arise in this setting *Bayesian offers*.

Along a dynamic path, when a buyer and seller with different expectations about the value of money are matched, they may be unable to reach agreement on the terms of trade. Thus, heterogeneous beliefs can lead to an endogenous extensive margin of trade, which we measure as an *acceptance rate*, that is, the proportion of take-it-or-leave-it offers accepted by sellers. Naturally, the failure of some offers to be accepted in some matches, that is, an acceptance rate lower than one, magnifies search frictions.

We present results that demonstrate how heterogeneous beliefs reduce welfare and, in particular, we characterize how welfare is impacted by variation in buyers' uncertainty, parameterized in terms of the spread of their prior distribution. We identify three ways uncertainty impacts welfare. First, uncertainty affects the intensive margin of trade as buyers make "cautious" offers whereby they offer a higher payment in exchange for a smaller quantity. By altering the demand for money, this affects the equilibrium price. Second, these cautious offers, all else equal, increase the acceptance rate of sellers, i.e. the extensive margin of trade. Third, small changes in uncertainty can bifurcate the equilibrium leading to qualitatively distinct dynamics that can feature higher or lower welfare depending on the resulting dynamics for the acceptance rate.

Another main result of this paper offers the unique theoretical predictions from the model with heterogeneous beliefs as an explanation for the results from a Lagos-Wright laboratory environment in Duffy and Puzzello (2013) who provide evidence in favor of monetary equilibria. They also find, however, that

1. a large fraction of offers made by buyers are not accepted by sellers;
2. the accepted offers are different from the theoretical stationary equilibrium price and quantity;
3. there is a non-degenerate distribution of money holdings despite evidence of portfolio rebalancing in the centralized market;
4. higher-quantity (buyer) offers are more likely to be rejected.

These experimental results are inconsistent with the rational expectations equilibrium but are consistent with the theoretical predictions of the monetary search model with heterogeneous expectations. Furthermore, the model presented in this paper makes sharp predictions regarding the relationship between prices and the acceptance rate, as well as between prices and sellers' realized surpluses, which should be invariant to the price of money under rational expectations.

What is the intuition for why heterogeneous beliefs can destabilize the stationary monetary equilibrium? In a stationary equilibrium, rational expectations (perfect foresight) and the adaptive predictor deliver the same forecast. Depending on the parameterization, the perfect-foresight steady state corresponding to the monetary equilibrium may be determinate or indeterminate. When the steady state is indeterminate, the monetary equilibrium is a sink under perfect foresight dynamics and a source under adaptive expectations. In a neighborhood of the stationary equilibrium, rational expectations and the adaptive predictor will both forecast similarly and agents will be unwilling to pay the computational costs necessary for perfect foresight. As the fraction of adaptive agents increases, the dynamics push the economy away from the stationary equilibrium until it reaches a point at which the accuracy gains associated to rational expectations outweigh the costs required for its use; hence, the fraction of rational agents increases and the economy moves towards the stationary equilibrium. The resulting tension between the stabilizing (or, attracting) and repelling dynamics can yield periodic orbits and complex dynamics.

There is a substantial literature that studies dynamic monetary equilibria in versions of the Lagos-Wright model. Lagos and Wright (2003) show that periodic and aperiodic equilibria can arise when bilateral trade is negotiated via Nash bargaining. Nosal and Rocheteau (2011) show the existence of periodic cycles in a model very close to ours with buyer-takes-all bargaining. The novelty in this paper is that these fluctuations arise for a different reason, namely, the endogenous distribution of heterogeneous beliefs. Moreover, this paper has distinct implications for welfare, the acceptability of money in trade, and monetary policy. We also show that non-linear dynamics can arise under a broader range of model parameterizations.

It is also important to note that the framework employed here can be applied in more general settings such as any extension of the model that might include other assets, such as

stocks or bonds, that can have a liquidity role in over-the-counter markets. Fiat money of course is an asset without a payoff stream. Thus, the types of dynamics present in this paper might arise in other asset pricing settings with search frictions, and may provide realistic asset price dynamics such as bubbles and crashes and may help explain other asset pricing phenomena.

The paper proceeds as follows. Section 2 details the general model environment and introduces take-it-or-leave-it offers with belief uncertainty, *Bayesian offers*. Section 3 defines the monetary equilibrium with heterogeneous beliefs, studies the steady-state properties with belief uncertainty, and presents a set of benchmark results on the dynamic stability properties of heterogeneous beliefs. Section 4 presents the implications of heterogeneous beliefs for welfare. Section 5 explores the unique theoretical predictions of the model in the Duffy-Puzzello experimental data, while Section 6 concludes.

1.1 Related literature

The results in this paper relate to a literature in monetary theory which incorporates heterogeneous valuations in search models, see for instance Rocheteau (2011). In Jacquet and Tan (2011), an extension to Rocheteau (2011), buyers and sellers attach different values to money because of shocks to their disutility of effort in the centralized market. In their model, sellers value money more and so are willing to produce more in any given match. The results in this paper are closely related, but the difference in values arises because of endogenously heterogeneous beliefs. Moreover, we show that it is possible for output, within a match, to be inefficiently too low or too high. In Engineer and Shi (1998), Engineer and Shi (2001) and Berentsen and Rocheteau (2003) buyers and sellers have asymmetric demands for the other's goods and study whether monetary equilibria can exist even without the double coincidence problem. Berentsen and Rocheteau (2003) show that in settings with asymmetric demands that money can have value in equilibrium because agents value money symmetrically. In our model, buyers and sellers value money asymmetrically when they have heterogeneous beliefs. It is an interesting question, left for future research, about the choice of which assets to use as a medium of exchange in an environment where buyers and sellers disagree about the value of only a subset of assets.

The existence of cycles and non-linear dynamics in monetary models is well-known. For example, in addition to the aforementioned Lagos and Wright (2003), Gu, Mattesini, Monnet, and Wright (2012) demonstrate non-linear perfect foresight dynamics in a model with endogenous credit constraints. Most of these perfect foresight models require the dynamic pricing equation to be non-monotonic much like in the extensive literature that studies non-linear dynamics in overlapping generations models. In the present paper, interesting dynamics can arise under more general conditions and rely on the tension between attracting and repelling dynamics that are inherent to heterogeneous expectations. The non-linear dynamics that arise in our model are very close to those in Brock and Hommes (1997) and Hommes (2013).

There is an extensive literature that studies bounded rationality and learning in macroeconomic models. Adaptive learning models formulated by Marcat and Sargent (1989) and

Evans and Honkapohja (2001) are based on a “cognitive consistency principle” that states that economic agents should be modeled like economists and econometricians who specify models and revise their beliefs in light of data. Here, the choice of the forecasting model is an endogenous object. Moreover, most monetary models with adaptive learning do not feature bilateral bargaining. One contribution of this paper is to extend learning analysis to frameworks where higher-order expectations matter.

The results of this paper are also related to monetary search models with non-degenerate wealth distributions that arise from the idiosyncratic consumption and production possibilities in search markets. Shi (1997b) uses a large household model to construct a model of divisible money with a non-degenerate wealth distribution. Similarly, Berentsen, Camera, and Waller (2005) and Molico (2006) study the distributional consequences of monetary policy in monetary search models. Here, we work with a very tractable extension of the Lagos-Wright model, where the distribution of money holdings arises because of heterogeneous beliefs. Because of the non-degenerate monetary distribution, there will be an endogenous price dispersion reminiscent of Peterson and Shi (2004).⁴

There is also a literature on private information in payoffs to assets and the potential role of signaling and other strategic considerations in the decentralized market. See, for example, Nosal and Wallace (2007), Lester, Postlewaite, and Wright (2012), Li, Rocheteau, and Weill (2012), and Golosov, Lorenzoni, and Tysvinski (2014). Relatedly, Berentsen, McBride, and Rocheteau (2014) find that uncertainty and private information in the payoffs of an asset can account for the large number of rejected offers in the laboratory. This paper focuses on agents with heterogeneous beliefs – who make take-it-or-leave-it offers given their priors about the value of money – and abstracts from the strategic issues. These and other issues are discussed in greater detail below.

2 Model

This section incorporates heterogeneous expectations into the Lagos and Rocheteau (2005) formulation of the Lagos and Wright (2005) model with *ex-ante* buyers and sellers. The model’s preferences, production technologies, and market structure are standard. The novel elements we introduce are the forecasting models available to agents, Bayesian bargaining offers, and the endogenous distribution of agents across forecast models.

2.1 General Environment

The environment is non-stochastic. Time is discrete, and each time period is divided into two sub-periods. There are two types of non-storable goods: specialized goods are produced and consumed in a decentralized market (DM) that opens during the first sub period; and general goods (numeraire) that are produced and consumed in a competitive market (CM) during the

⁴Heterogeneous beliefs of the type under consideration here have been shown to be consistent with findings from learning-to-forecast experiments. See, for example, Anufriev and Hommes (2012) and Assenza, Heemeijer, Hommes, and Massaro (2011).

second sub-period. There are two types of agents: “buyers” and “sellers”. All agents produce CM goods using the same technology, but are distinguished by their ability to produce DM goods: sellers have access to a specialized-good production technology and buyers do not whereas only buyers receive utility from specialized-good consumption. Additionally, there is fiat money, a divisible, intrinsically worthless asset with outside supply M . We denote the price of fiat money (in terms of the numeraire) as ϕ . Fiat money is the only storable good in the economy.

Trading in the CM is centralized and coordinated by a Walrasian auctioneer. The DM is characterized by bilateral trading, which results in a double coincidence of wants problem. A seller and buyer meet with probability σ , capturing a standard search friction. When $\sigma = 1$, the search friction is shut down and each buyer is matched with certainty to a potential trading partner. Limited commitment and imperfect record-keeping preclude the use of unsecured credit in these meetings and, thereby, provide a role for fiat money to facilitate trade that otherwise would not occur. It is straightforward to model a setting where instead of money some other asset, such as a bond or a claim to a Lucas tree, can be used to secure trade in the DM.

We now describe the behavior of buyers and sellers in more detail. There is a unit mass of sellers who produce and consume the CM good, and produce but do not consume the DM good. The production of CM goods utilizes a technology that is linear in labor. Let q be the quantity of the DM good produced, and let x be the net quantity of the CM good consumed, where negative values of x arise when more is produced than consumed. The representative seller’s preferences are captured by $E \sum_t \beta^t U^s(q_t, x_t)$ where $U^s(q, x) = x - c(q)$, and the expectations operator here is taken to be rational – sellers always have perfect foresight in our model. Also, $c(q) \geq 0$ measures the disutility of labor associated to producing the quantity q ; the quasi-linearity in x , the net CM good consumed, greatly simplifies the analysis by eliminating wealth effects.

There is a unit mass of buyers, $i \in \mathcal{I}$, who produce and consume the CM good, and who consume but do not produce the DM good. In general, the buyer-specific reference $i \in \mathcal{I}$ will be suppressed in the notation except when predictor choice is discussed and tracking the decisions of specific agents is necessary for clarity. Again, the production of CM goods utilizes a technology that is linear in labor. Unlike sellers who are rational, buyers are boundedly rational, and must select a forecast model (also referred to as a “predictor” or “expectations type”) each period. Below, we are explicit about when and how a buyer chooses his predictor. As is common in the literature on bounded rationality, we model buyers as *anticipated utility maximizers*, which, within the context of our modeling environment, means they make decisions each sub period assuming their current predictor will be used indefinitely. See, for instance, Kreps (1998) and Cogley and Sargent (2008). Using the same notation as above, but now with q representing the quantity of the DM good consumed, the buyer’s instantaneous utility is given by $U^b(q, x, \tau) = x + u(q) + \Omega(\tau)$, where the τ captures expectations type, and the term $\Omega(\tau)$ represents the utility gained from the selection of a forecasting model of type τ . We will refer to a buyer who selects predictor τ as a “buyer of type τ ,” but it is important to remark that predictor selection, and hence buyer type, evolves over time. Finally, $u(q)$ measures the utility associated to consuming the quantity q . Assume as usual that $u' > 0$ and $u'' < 0$, and we also normalize $u(0) = 0$ for convenience.

We discuss in detail below how this specification of instantaneous utility embeds into the buyer’s dynamic decision problem.

Agents’ optimal money holdings and bilateral bargaining depend, in part, on beliefs about the (future) value of money. As mentioned above, we assume that sellers are rational in that they always have perfect foresight, whereas buyers are assumed to be boundedly rational, making decisions based on beliefs which are updated over time. The timing of the model has a buyer of type τ deciding on money holdings in the CM of period t given their beliefs about the future CM value of money and their expectations about the outcome of bilateral bargaining in the period $t + 1$ DM. Therefore, buyers in the period t CM must also formulate expectations about their own future expectations as well as expectations about the seller’s beliefs. It is this latter feature which leads to novel implications and is an aspect absent from the standard learning framework typically applied in monetary models without decentralized, bilateral trading between agents.

Though below we are much more specific about the particular timing assumptions, it is useful to provide a brief overview here. At the beginning of the first sub-period the DM opens. Buyers come into the first sub-period holding money and with pre-determined beliefs. Buyers and sellers are matched, bargaining ensues, and the first sub-period ends. The CM opens in the second sub period, buyers rebalance their portfolios and markets clear. Between the closing of the CM, after new data have been realized, and the opening of the DM buyers select their new expectations type and update their beliefs.⁵

2.2 The Centralized Market

Denote by W_t the buyer’s value function at the beginning of the second sub-period of period t (i.e. at the opening of the CM), and by V_t the value function at the opening of the DM. A superscript “s” signifies the associated value functions for the sellers.

Buyers enter the CM with a quantify of money m_t and predictor type τ_t . Buyers choose the net quantity x_t of the (CM) numeraire good to consume and how much money m_{t+1} to carry into the next period to use for DM trading. Thus, they solve the following problem:

$$W_t(m_t, \phi_t, \tau_t) = \Omega(\tau_t) + \max_{m_{t+1}, x_t} \{x_t + \beta E_t^{\tau_t} V_{t+1}(m_{t+1}, \phi_{t+1}, \tau_t)\}$$

$$x_t + \phi_t m_{t+1} = \phi_t m_t.$$

Note that while the buyer will select a new, and possibly different predictor before the opening of the DM in period $t + 1$, he makes decisions during the CM of period t assuming his predictor type tomorrow will be the same as his predictor type today, i.e. the predictor-type argument in V_{t+1} and Ω is τ_t , not τ_{t+1} . This reflects the anticipated-utility assumption mentioned earlier.

The problem may be written equivalently as

$$W_t(m_t, \phi_t, \tau_t) = \phi_t m_t + \Omega(\tau_t) + \max_{m_{t+1}} \{-\phi_t m_{t+1} + \beta E_t^{\tau_t} V_{t+1}(m_{t+1}, \phi_{t+1}, \tau_t)\}.$$

⁵Alternatively, we could identify a third sub-period in which predictor selection is made, but it seems more convenient notionally to specify a “between-period” decision.

Note that the optimization problem is separable in m_t , so that all buyers of the same predictor-type make the same rebalancing decision m_{t+1} , regardless of their entering money stock m_t .

Sellers solve an analogous problem, under the assumption of perfect foresight. If $\beta\phi_{t+1} \geq \phi_t$ then sellers may want to buy money in the CM. However, if $\beta\phi_{t+1} > \phi_t$ then money demand is infinite and no equilibrium exists, and if $\beta\phi_{t+1} = \phi_t$ then agents are indifferent across money holdings. Thus, we focus on price paths that have the property that $\beta\phi_{t+1} < \phi_t$, and check that this condition holds in the numerical analysis. In this case, the sellers supply money received in the DM inelastically in the CM.

2.3 The Predictor Choice

After the close of the CM in period- t , the buyer updates his predictor choice. There are N available predictors indexed by $\tau \in \{1, \dots, N\}$, with $\phi_{t+n}^e(\tau, t)$ being predictor-type τ 's forecast of the time- $t+n$ price of money for $n \geq 1$, formulated between the closing of the CM in period t and the opening of the DM in period $t+1$. For notational convenience, we denote $\phi_{t+n}^e(\tau) \equiv \phi_{t+n}^e(\tau, t)$ for point expectations, with the period when forecasts are formed being implied. A type τ predictor also comes with an uncertainty measure $F^\tau(\cdot, \Sigma)$, which may be important for bargaining in the DM: this will be discussed in detail in Section 2.4. In many of the examples, we place a particular structure on the set of forecasting models by assuming that buyers hold either perfect foresight beliefs (rational expectations) or beliefs summarized by an adaptive learning rule in the spirit of Evans and Honkapohja (2001).

Each buyer $i \in \mathcal{I}$ selects his predictor type τ_{t+1} by solving the following problem:

$$\begin{aligned} \tau_{t+1} &= \operatorname{argmax}_{\tau \in \{1, \dots, N\}} \Omega_i(\tau, \phi_t^e(\tau), \phi_t) \\ \text{where} \quad \Omega_i(\tau, \phi_t^e(\tau), \phi_t) &= -(\phi_t^e(\tau) - \phi_t)^2 - C_{it}(\tau). \end{aligned} \quad (1)$$

The objective $\Omega_i(\tau, \phi_t^e(\tau), \phi_t)$ may be viewed as a utility function with two components: $-(\phi_t^e(\tau) - \phi_t)^2$ captures the past performance of predictor τ ; and $C_{it}(\tau)$ is an idiosyncratic preference shock measuring buyer i 's ease, or lack thereof, with using predictor τ , i.e. the ‘‘cost’’ to adopting predictor τ by buyer i . The cross-sectional distribution of this shock determines the proportion of buyers using a particular predictor in a particular period: see Brock and Hommes (1997) for details. We return to this point in Section 3.1.

2.4 The Decentralized Market

At the opening of the DM, buyers and sellers are randomly matched. Given a match, the buyer trades money for the specialized good produced by the seller, and the terms of trade are negotiated through bargaining. We emphasize here take-it-or-leave-it offers made by buyers given their beliefs about the value of money and their prior over sellers' beliefs, i.e. the sellers' participation constraints. We also consider a Nash bargaining protocol in the Appendix. Nash bargaining facilitates a simple and interesting result, Proposition 6, that illustrates a trading inefficiency that arises when there is a belief disagreement between buyer

and sellers.⁶ Buyer-takes-all bargaining has two advantages. First, it facilitates analytic results in the model with heterogeneous beliefs. Second, it is a convenient device to alter the implications of common knowledge assumptions in bilateral trade.

2.4.1 Buyers' Beliefs

The beliefs of a buyer selecting predictor τ at the closing of the CM in period t have two components: point-forecasts $\phi_{t+n}^e(\tau)$ and bargaining uncertainty $F_{t+1}^\tau(\cdot, \Sigma)$.⁷ While a buyer of type τ believes with certainty, at the opening of the CM in period t , that the $t+1$ -price of money will be $\phi_{t+1}^e(\tau)$, he must also formulate expectations about the seller's beliefs. This is where bargaining uncertainty comes in to play, and it is best to begin with a discussion of common knowledge.

The terms of trade depend on the buyer's and the seller's forecasts of the price of money, since this value will be determined in the subsequent CM. When buyers and sellers with different beliefs about the future value of money are matched, the bargaining process will evidently depend on the extent to which beliefs are common knowledge within the match. Under the take-it-or-leave-it protocol there is no extended interaction between agents. The buyer, who is taken to be the first mover, simply makes an offer to the matched seller and the seller decides whether to accept. To make his offer, the buyer must forecast the price expectations of the seller and the lack of interaction precludes the extraction of useful information from the seller about his/her beliefs. Under this protocol, we do not assume common knowledge of beliefs; instead we assume that buyers use their own beliefs as proxies for the forecasts of the sellers. However, while a buyer of type τ is confident that $\phi_{t+1} = \phi_{t+1}^e(\tau)$, he is not certain that the seller holds a similar forecast. Instead, a buyer of type τ behaves like a Bayesian, acknowledging the uncertainty, and placing a prior on a seller's price forecast via a distribution $F_{t+1}^\tau(\cdot, \Sigma)$, where the mean of this distribution is $\phi_{t+1}^e(\tau)$, and the variance is measured by Σ . We interpret Σ as parameterizing buyers' *uncertainty* about sellers' beliefs. As adaptive agents learn over time about ϕ_t , they update their prior accordingly.⁸

2.4.2 Take-it-or-leave-it protocol

To simplify exposition in these next few subsections, we drop time subscripts: all the action is happening within period. Also, we suppress the predictor type, as it is taken as given by the buyer when offers are made. Buyers and sellers are randomly matched and buyers make (q, d) offers under a take-it-or-leave-it protocol, where q is the quantity of the good contracted as deliverable and d is the quantity of money offered in exchange. Denote by ϕ^b and ϕ^s the point expectations of the price ϕ held by buyers and sellers, respectively, at the opening of the DM. To be accepted by a matched seller, the corresponding buyer's offer

⁶There is a close connection between Nash bargaining and Bayesian offers, that is explored below.

⁷Only the point forecasts $\phi_{t+1}^e(\tau)$ are needed for bargaining in the period $t+1$ DM.

⁸The Appendix also considers a formulation of the model where adaptive agents also specify an updating rule for Σ .

must satisfy a participation constraint given by

$$W^s(d, \phi^s) - c(q) \geq W^s(0, \phi^s).$$

Since W^s is quasi-linear in ϕ^s , this constraint may be equivalently written as $\phi^s d \geq c(q)$.

Under homogeneous rational expectations, the buyer knows ϕ^s and therefore will never offer more than the “reservation wage” $c(q)/d$. Our modeling environment is distinguished by heterogeneity and bounded rationality. In particular a matched buyer does not know with certainty the corresponding seller’s perceived price ϕ^s , and therefore cannot determine the reservation wage. As discussed in Section 2.4.1, a buyer takes ϕ^s as a random variable with mean ϕ^b and distribution given by F_Σ . In this initial discussion, we assume buyers hold dogmatic beliefs: $\Sigma = 0$ and $\phi^s = \phi^b$. The case of buyers with dogmatic beliefs is useful for providing intuition and analytic results. In the next section we consider the more general case in which $\Sigma > 0$.

Buyers are first movers under this bargaining arrangement. A matched buyer is assumed to know the preferences of the corresponding seller but does not know the seller’s price forecasts. As discussed above, the buyer uses his own price forecast to form an expectation of the seller’s price forecast. More specifically, under dogmatic priors, a matched buyer makes an offer by solving the following problem

$$\begin{aligned} \max_{q, d \leq m} \quad & u(q) + W_t(m - d, \phi^b) \\ & \phi^b d \geq c(q), \end{aligned} \tag{2}$$

where the inequality is the seller’s participation constraint, and to aid notation we suppress here, and in the following section, the dependence of the CM value function on τ . Since the buyer will always offer their forecast of the seller’s reservation price, this participation constraint binds, which leads to the following solution:

$$\begin{aligned} q(m, \phi^b) &= \begin{cases} q^* & \text{if } \phi^b m \geq c(q^*) \\ c^{-1}(\phi^b m) & \text{else} \end{cases} \\ d(m, \phi^b) &= \frac{c(q(m, \phi^b))}{\phi^b}. \end{aligned}$$

Here, q^* is the efficient level of output, as determined by $u'(q^*) = c'(q^*)$: see the Appendix for more details.

2.4.3 Bayesian offers

We now relax the dogmatic assumption that buyers believe $\phi^s = \phi^b$, and instead allow buyers to acknowledge their uncertainty about sellers’ beliefs and behave like a Bayesian with a subjective prior distribution F_Σ . A Bayesian take-it-or-leave-it offer, then, balances maximizing a buyer’s surplus against the subjective probability an offer will be accepted. We call this form of take-it-or-leave-it bargaining *Bayesian offers*.

Let $\chi(q, d, \phi^s)$ be the characteristic function identifying an acceptable offer, that is,

$$\chi(q, d, \phi^s) = \begin{cases} 1 & \text{if } \phi^s d \geq c(q) \\ 0 & \text{else} \end{cases}$$

The buyer's offer solves

$$\begin{aligned} & \max_{q, d \leq m} E_{F_\Sigma} \left((u(q) + W^b(m - d, \phi^b)) \chi(q, d, \phi^s) + W^b(m, \phi^b) (1 - \chi(q, d, \phi^s)) \right), \\ \text{or } & \max_{q, d \leq m} (u(q) - \phi^b d) \left(1 - F_\Sigma \left(\frac{c(q)}{d} \right) \right), \end{aligned} \quad (3)$$

where the second expression exploits quasi-linearity, $u(0) = 0$, and computes expectations. The notation E_{F_Σ} emphasizes that we are computing expectations over ϕ^s against the distribution F_Σ , which is centered at ϕ^b .⁹

In case the liquidity constraint $d \leq m$ is not binding, the buyer's offer satisfies the following conditions:

$$\begin{aligned} \frac{u'(q)}{c'(q)} &= \frac{\phi^b d}{c(q)} \\ u'(q) &= (u(q) - \phi^b d) \left(\frac{c'(q)}{d} \right) h_\Sigma \left(\frac{c(q)}{d} \right), \end{aligned} \quad (4)$$

where $h_\Sigma(x) = dF_\Sigma(x)(1 - F_\Sigma(x))^{-1}$ is the hazard function associated to the distribution F_Σ . In this case $h_\Sigma \left(\frac{c(q)}{d} \right)$ measures the instantaneous rate of rejection given that the offer (q, d) is accepted. The first condition in (4) equates the marginal rate of substitution with the expected ratio of relative valuations of money.¹⁰ The latter condition in (4), equates the marginal benefit to a buyer from an offer with the expected surplus lost by rejection, i.e. it's the surplus $u(q) - \phi^b d$ times the conditional probability of rejection, the latter being given by $\left(\frac{c'(q)}{d} \right) h_\Sigma \left(\frac{c(q)}{d} \right)$.

The condition (4) is familiar. In case of dogmatic priors, that is, $\Sigma = 0$, the buyer's offer will correspond to the perceived reservation wage of the seller: $\phi^b d = c(q)$. In this case, the quantity offered, q , will satisfy the usual condition for efficiency: $u'(q) = c'(q)$. The response of the buyer's offer to changes in uncertainty is challenging to analyze at the current level of generality. Below, we present two specific examples.

Now consider the constrained case. Imposing $d = m$ reduces the buyer's decision problem

⁹We restrict buyers to making a single (q, d) offer in any given meeting. Also, our focus is on dynamic equilibria where prices (hence, ϕ^s) vary over time. Since sellers' beliefs are never fully observed, all interactions are anonymous and infrequent, we therefore rule out buyers who engage in experimentation.

¹⁰Notice that $c(q)/d$ is the seller's marginal value of money. Thus this F.O.C. is analogous to the condition in Nash bargaining with heterogeneous beliefs which equates $u'(q)/c'(q) = \phi^b/\phi^s$, as shown in the Appendix. In the dogmatic case, when the constraint is not binding the offer satisfies $u'(q)/c'(q) = 1$. The key difference between the Bayesian offers and the Nash bargaining is the second condition in (4), which accounts for the perceived acceptance rate.

to one dimension.¹¹ The objective function $(u(q) - \phi^b m) \left(1 - F_\Sigma \left(\frac{c(q)}{m}\right)\right)$ is not globally concave. However, existence and uniqueness can be established under a variety of different assumptions. The following proposition provides the assumptions we will use in the sequel:

Proposition 1 *If F_Σ has finite second moments and if $h'_\Sigma > 0$ then the constrained problem*

$$\max_{q \geq 0} (u(q) - \phi^b m) \left(1 - F_\Sigma \left(\frac{c(q)}{m}\right)\right) \quad (5)$$

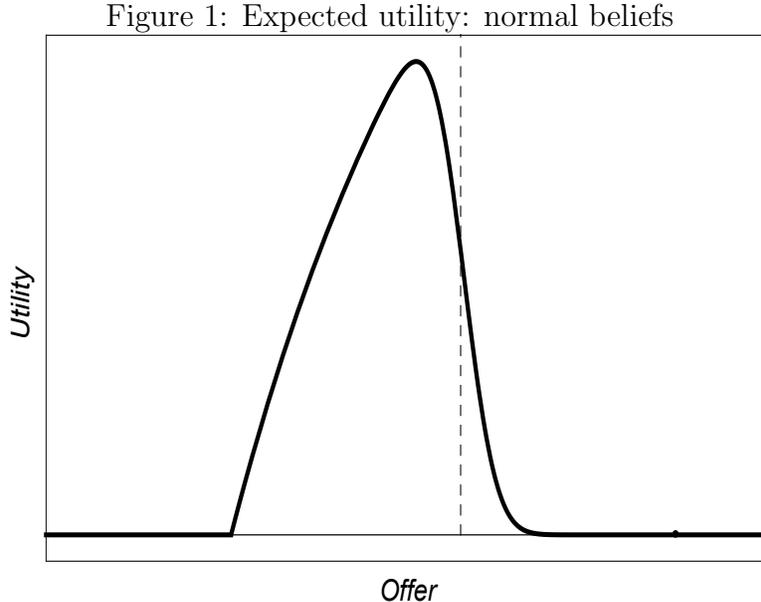
has a unique solution.

The proof is in the Appendix.

To aid intuition, consider Figure 1. Here we plot the objective

$$(u(q) - \phi^b m) \left(1 - F_\Sigma \left(\frac{c(q)}{m}\right)\right)$$

under the assumption that beliefs are normally distributed about ϕ^b and that u is an affine-transformed CRRA utility function. The vertical dashed-line indicates the offer that results from dogmatic priors. A robust feature of the monetary equilibria that we consider is that, provided the uncertainty is not too large, the buyer's objective function is maximized at an offer that is to the left of the dashed line, i.e. a more cautious offer.

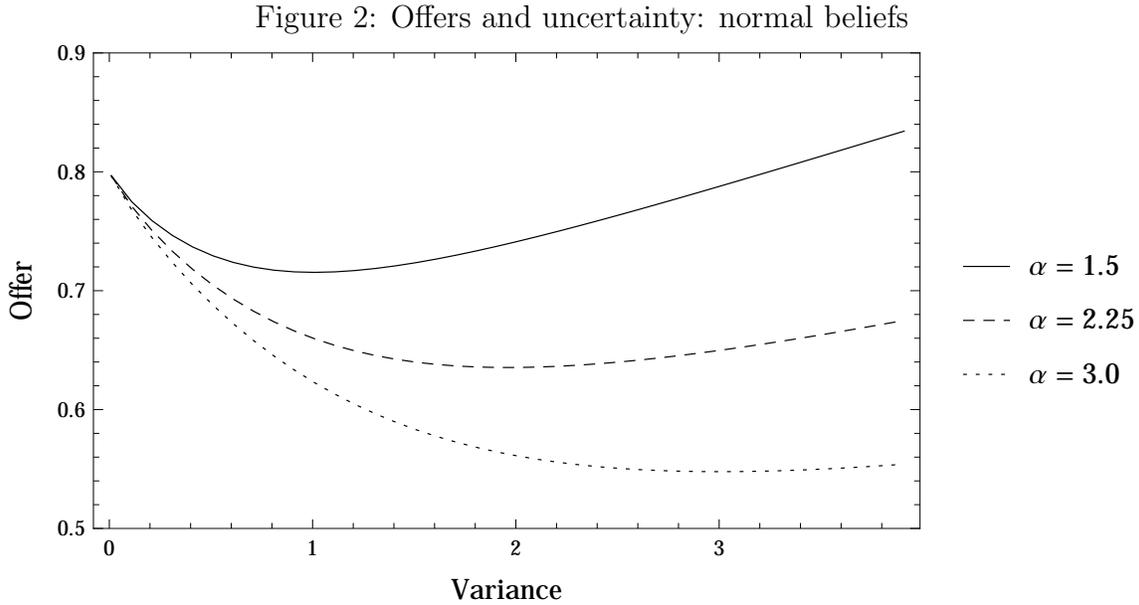


The intuition underlying that Bayesian offers leads to, all else equal, a lower offer is that for fixed uncertainty Σ , a higher offer is more likely to be rejected and a risk-averse buyer

¹¹We assume, for this discussion, that there is positive surplus at mean beliefs: if $c(q) = \phi^b m$ then $u(q) > c(q)$. This assumption is stronger than is needed to guarantee a matched buyer makes a non-trivial offer, and also is innocuous in the following sense: along the equilibrium time-paths that we consider, buyers will never hold more money than they intend to offer, and therefore will never hold so much money that if it is all offered the surplus is negative.

will make offers that take the acceptance probability into account. However, the logic is subtle as higher uncertainty could also lead buyers to gamble and make a higher offer if the resulting surplus is sufficiently high to compensate the buyer for the possibility the offer is rejected by the seller.

Continuing to assume a normally distributed prior, the response of the buyer's offer to a change in uncertainty is ambiguous. Figure 2 below plots the offer as a function of the uncertainty parameter Σ ; different offer curves correspond to different levels of risk aversion α . As the variance increases from zero the offer initially falls; however, as the variance continues to rise, the slope of the offer curve becomes positive, and if the buyer is not too risk averse, a high variance may lead to offers above the dogmatic-priors level.



What is the intuition for the finding that the scale of uncertainty can affect whether Bayesian offers are cautious? Consider the following simple illustrative example. Let ε_l and ε_h be (smallish) positive numbers and suppose that the buyer believes the seller's expected price ϕ^s is distributed in the finite set $\{\phi^b - \varepsilon_l, \phi^b, \phi^b + \varepsilon_h\}$ according to the density (π_l, π, π_h) .¹² Further, assume that $\pi_l \varepsilon_l = \pi_h \varepsilon_h$ so that this distribution has mean ϕ^b . It follows that

$$\pi_l = (1 - \pi) \frac{\varepsilon_h}{\varepsilon_l + \varepsilon_h} \text{ and } \pi_h = (1 - \pi) \frac{\varepsilon_l}{\varepsilon_l + \varepsilon_h}.$$

Finally, for simplicity, assume money-holdings $m = 1$. With this simple setup, the buyer's behavior is easily characterized: he chooses $q \in \{\phi^b - \varepsilon_l, \phi^b, \phi^b + \varepsilon_h\}$ according to which of these choices solves

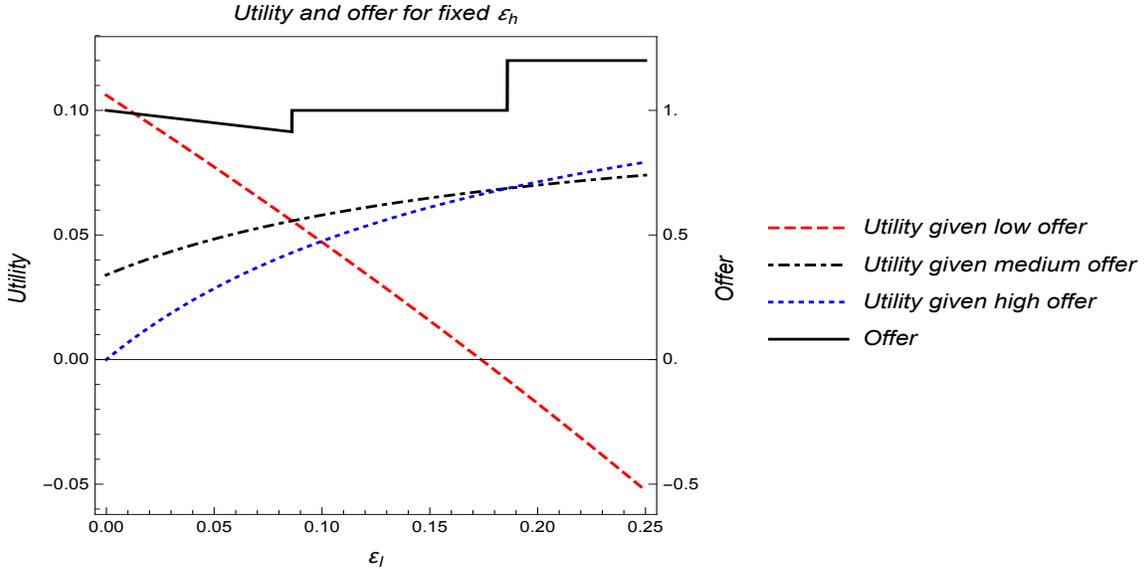
$$\max \left\{ \begin{array}{l} u(\phi^b - \varepsilon_l) - \phi^b \\ (u(\phi^b) - \phi^b)(1 - \pi_l) \\ (u(\phi^b + \varepsilon_h) - \phi^b)\pi_h \end{array} \right\}.$$

If $\pi < 1$ and $\varepsilon_l = \varepsilon_h \equiv \varepsilon$ is small, then $q = \phi^b - \varepsilon$, i.e. the buyer makes an offer that will be accepted with probability 1. Now fix ε_h and allow ε_l to increase. Initially, the

¹²This example is illustrative and admittedly does not match our assumptions above.

value $u(\phi^b - \varepsilon_l) - \phi^b$ of the low offer continues to dominate, so that the offer q is falling as ε_l increases. However, as ε_l continues to increase, the value $u(\phi^b - \varepsilon_l) - \phi^b$ of a low offer continues to fall, and the probability $1 - \pi_l$ that a medium or high offer is accepted increases. Eventually, one of the latter dominates and the buyer switches offers, so that q is (discontinuously) increasing in ε_l . This scenario is exhibited in Figure 3 below. Note that the quantity offered is indicated on the RHS scale.

Figure 3: Offers and uncertainty: discrete beliefs



The simple illustration just considered demonstrates the complex trade-offs at play when beliefs change. The situation is even more complex when continuously supported beliefs are imposed, and in general, analytic results are not available. For the computational work in the sequel, we assume that buyers' beliefs are normally distributed about ϕ^b , and we use the variance of the distribution as a measure of uncertainty. The corresponding hazard function does not have a closed-form expression, but can be shown to be increasing, as demanded by our assumptions.

These examples illustrate the comparative static effects of uncertainty on the Bayesian offers. Of course, in equilibrium ϕ and m are endogenous objects that depend, in turn, on the optimal Bayesian offer. Below, the monetary equilibria we consider have the property that higher uncertainty leads to lower q in Bayesian offers.

2.5 Money demand

We now derive money demand given the bargaining solutions in the previous sections. Recall that $\phi_t^e(\tau_t)$ is the price point expectation of a buyer of type τ_t at the opening of the time t DM. Denote by $(q(m_t, \phi_t^e(\tau_t)), d(m_t, \phi_t^e(\tau_t)))$ the period t offer of an agent of type τ_t . The

type- τ_t buyer's DM value function may be written

$$\begin{aligned}
V_t(m_t, \phi_t^e(\tau_t), \tau_t) &= \sigma E_t^{\tau_t} \left(\chi(q(m_t, \phi_t^e(\tau_t)), d(m_t, \phi_t^e(\tau_t)), \phi^s) \left\{ u(q(m_t, \phi_t^e(\tau_t))) \right. \right. \\
&\quad \left. \left. + W_t(m_t - d(m_t, \phi_t^e(\tau_t)), \phi_t^e(\tau_t), \tau_t) \right\} \right) \\
&\quad + \sigma E_t^{\tau_t} \left((1 - \chi(q(m_t, \phi_t^e(\tau_t)), d(m_t, \phi_t^e(\tau_t)), \phi^s)) W_t(m_t, \phi_t^e(\tau_t), \tau_t) \right) \\
&\quad + (1 - \sigma) W_t(m_t, \phi_t^e(\tau_t), \tau_t),
\end{aligned}$$

where $E_t^{\tau_t}$ represents the expectation taken over ϕ^s against the distribution $F_{\Sigma}^{\tau_t}$, which is centered at $\phi_t^e(\tau_t)$. Using $W_t(m_t, \phi_t, \tau_t) = \phi_t m_t + \Omega(\tau_t) + W(0, \phi_t, \tau_t)$, it is straightforward to show that a household of expectations type τ_t chooses its money m_{t+1} , in the centralized market, to solve

$$\begin{aligned}
W_t(m_t, \phi_t, \tau_t) &= \phi_t m_t + \Omega(\tau_t) + \max_{m_{t+1}} \left\{ -\phi_t m_{t+1} + \beta E_t^{\tau_t} V_{t+1}(m_{t+1}, \phi_{t+1}^e(\tau_t), \tau_t) \right\} \\
&= \phi_t m_t + \Omega(\tau_t) + \beta W_{t+1}(0, \phi_{t+1}^e(\tau_t), \tau_t) + \max_{m_{t+1}} \left\{ -(\phi_t - \beta \phi_{t+1}^e(\tau_t)) m_{t+1} \right. \\
&\quad \left. + \beta \sigma \left(1 - F_{\Sigma}^{\tau_t} \left(\frac{c(q(m_{t+1}, \phi_{t+1}^e(\tau_t)))}{d(m_{t+1}, \phi_{t+1}^e(\tau_t))} \right) \right) \right. \\
&\quad \left. \times \left(u(q(m_{t+1}, \phi_{t+1}^e(\tau_t))) - \phi_{t+1}^e d(m_{t+1}, \phi_{t+1}^e(\tau_t)) \right) \right\}.
\end{aligned}$$

The solution to this optimization problem determines the buyer's money demand as a function of current price and expectations.

There are three cases to consider:

- If $\phi_{t+1}^e(\tau_t) > \beta^{-1} \phi_t$, then there is no solution since expected rate of return exceeds the rate of time preference.
- If $\phi_{t+1}^e(\tau_t) = \beta^{-1} \phi_t$, then they are indifferent and any $m \geq c(q^*)/\phi_{t+1}^e(\tau_t)$ is a solution, and $d = c(q^*)/\phi_{t+1}^e(\tau_t)$.
- If $\phi_{t+1}^e(\tau_t) < \beta^{-1} \phi_t$, $d = m_{t+1}$ and m_{t+1} solves the interior F.O.C.

$$\begin{aligned}
0 &= -\phi_t + \beta \phi_{t+1}^e(\tau_t) + \beta \sigma \left(1 - F_{\Sigma}^{\tau_t} \left(\frac{c(q(m_{t+1}, \phi_{t+1}^e(\tau_t)))}{m_{t+1}} \right) \right) \\
&\quad \times \left(u'(q(m_{t+1}, \phi_{t+1}^e(\tau_t))) \times \frac{\partial}{\partial m_{t+1}} (q(m_{t+1}, \phi_{t+1}^e(\tau_t))) - \phi_{t+1}^e(\tau_t) \right) \\
&\quad - \beta \sigma d F_{\Sigma}^{\tau_t} \left(\frac{c(q(m_{t+1}, \phi_{t+1}^e(\tau_t)))}{m_{t+1}} \right) \times \frac{\partial}{\partial m_{t+1}} \left(\frac{c(q(m_{t+1}, \phi_{t+1}^e(\tau_t)))}{m_{t+1}} \right).
\end{aligned} \tag{6}$$

In the paper we focus only on price paths for which this last case holds, so that we may use the interior FOC to determine money demand. This last condition arises whenever

buyers are liquidity constrained in the sense that they do not hold enough money balances to purchase the efficient level of DM consumption. As equation (6) shows, a buyer’s optimal money holdings depends on the expected holding return (cost) $\beta\phi_{t+1}^e(\tau_t)$ and the second term that reflects the liquidity premium associated to the transaction role of money in bilateral trade. This latter term is positive only when money is scarce so that the liquidity constraint $d \leq m$ binds. This first-order condition implicitly defines the demand for money in terms of current price and expectations: we denote this demand by $m_{t+1} = m(\phi_t, \phi_{t+1}^e(\tau))$. While closed form solutions do not, in general, exist in the case of heterogeneous expectations, it is possible to numerically solve for the money demand function.

3 Heterogeneous Expectations and Monetary Equilibria

The Rocheteau and Wright (2005) and Lagos and Rocheteau (2005) models already have heterogeneous agents in the form of *ex-ante* buyers and sellers. Our imposed belief structure introduces heterogeneity within buyers and potentially between buyers and sellers. Notice that, under the maintained assumptions, the money market equilibrium (and so equilibrium ϕ_t) only depends on the expectations of the buyers since they are the ones who will choose to hold money in the centralized market. On the other hand, the equilibrium quantity of DM good produced depends, in part, on the expectations of sellers.

3.1 Heterogeneous Expectations

As discussed above, we assume that buyers have an array of predictors, or forecasting models, available to them. The particular set of forecasting models that we consider is motivated by the econometric learning literature which models economic agents as econometricians who take seriously the applied econometric problem of choosing a satisfactory specification for the forecasting model. Obviously, the best performing forecasting model is a correctly specified model that captures the precise dynamic properties for ϕ_t , in this case it would be the perfect foresight predictor. However, formulating perfect foresight is likely to incur a computational and cognitive cost. Instead, the advice of many econometricians is to choose “parsimonious” forecasting models that project recently observed data into a forecast for future values of ϕ_t that minimizes mean-squared forecast errors. Whether a particular buyer is willing to pay the cost of acquiring perfect foresight depends on its benefits relative to simpler forecasting models. A key aspect of the approach in this paper is that the distribution of agents across forecasting models is an endogenous object, thereby, preserving many of the cross-equation restrictions that are salient features of rational expectations models.

We assume that buyers select their forecasting models by weighing past performance, measured in terms of mean-squared forecast errors (the Ω term), against the cost of usage in terms of utility. Recall that there are N predictor types, $\tau \in \{1, \dots, N\}$, with point-forecasts denoted by $\phi_{t+n}^e(\tau)$ for $n \geq 1$. Note that at the opening of the DM in period t , a buyer will hold a predictor of type τ_t and will use point forecasts $\phi_t^e(\tau_t)$ when engaged in bargaining.

Recall from Section 2.3 that between the close of the time $(t - 1)$ CM and the opening of the time t DM, buyer $i \in \mathcal{I}$ selects his predictor by maximizing the objective

$$\Omega_i(\tau, \phi_{t-1}^e(\tau), \phi_{t-1}) = -(\phi_{t-1}^e(\tau) - \phi_{t-1})^2 - C_{it-1}(\tau).$$

We now assume $C_{it}(\tau) = C(\tau) + \nu_{it}(\tau)$. Provided that the $\nu_{it}(\tau)$ are serially and cross-sectionally i.i.d and hold an extreme value distribution, the proportion of buyers, $n_t(\tau)$, using predictor τ in period t to forecast ϕ_t and ϕ_{t+1} is given by the MNL map (7):¹³

$$n_t(\tau) \equiv n_t\left(\tau, \phi_{t-1}, \{\phi_{t-1}^e(\omega)\}_{\omega=1}^N\right) = \frac{\exp(\gamma \cdot \Omega(\tau, \phi_{t-1}^e(\tau), \phi_{t-1}))}{\sum_{\omega=1}^N \exp(\gamma \cdot \Omega(\omega, \phi_{t-1}^e(\omega), \phi_{t-1}))}, \quad (7)$$

where

$$\Omega(\tau, \phi_{t-1}^e(\tau), \phi_{t-1}) = -(\phi_{t-1}^e(\tau) - \phi_{t-1})^2 - C(\tau).$$

The MNL approach has a long and venerable history in discrete decision making and is a natural way of introducing randomness in forecasting into the present environment. Young (2004) argues that randomness in forecasting has a similar interpretation to mixed strategies in actions in that it provides robustness against forecasting model uncertainty. The parameter γ is typically called the “intensity of choice” parameter and is inversely related to the variance of the cross-sectional distribution of $\nu_{it}(\tau)$. Since the MNL map is derived from a random utility setting, finite values of γ parameterize deviations from full utility maximization. In our analysis below, we compare the equilibrium dynamics across various values of γ . The “neoclassical” case is $\gamma \rightarrow \infty$ where agents choose only the best performing forecasting model.

The timing assumptions here require special discussion. We follow the adaptive learning literature in assuming that current values of the endogenous state variables are not directly observable when forming boundedly rational forecasts. This is usually assumed to avoid a simultaneity in least-squares parameter estimates and the endogenous variables. In this setting, the assumption preserves logical consistency for adaptive agents. Rational agents (who have one-step-ahead perfect foresight) know the current value for the price, but the adaptive agents do not. Under this natural assumption, predictor selection takes place at the end of the period, after updating information from the CM, and then forecasts are made. As discussed previously, boundedly rational forecasts are not updated after trading in the DM. The approach taken here assumes that agents have a menu of predictor choices. Each agent looks to the most recent forecasting performance which, combined with their random-utility preference shock underlying the MNL map, chooses a predictor that they will use in the present period.¹⁴

At this point, a brief remark is warranted about the interpretation of the money-demand function $m(\phi_t, \phi_{t+1}^e(\tau))$ for buyers who do not hold rational expectations. It was previously

¹³See Brock and Hommes (1997) for details.

¹⁴At the end of period t , a buyer of type τ chooses their money-holdings optimally given their beliefs. We do not impose that the individual correctly foresees how their beliefs might change and evolve in subsequent periods, in particular, at the beginning of the following period when predictors are selected by buyers. These behavioral assumptions follow the anticipated utility approach emphasized by Sargent (1999) who adapted the concept from Kreps (1998).

assumed that a buyer with boundedly rational beliefs does not observe contemporaneous ϕ_t when forming expectations. We interpret the optimization problem as determining a money demand schedule that the agents turn into the Walrasian auctioneer in the centralized market and the auctioneer sets the price to clear the market.¹⁵

An equilibrium price path is a sequence of prices ϕ_t satisfying market clearing:

$$\sum_{\tau=1}^N n_t \left(\tau, \phi_{t-1}, \{\phi_{t-1}^e(\omega)\}_{\omega=1}^N \right) m(\phi_t, \phi_{t+1}^e(\tau)) = M. \quad (8)$$

Recall that the pair

$$(q(m_{t+1}, \phi_{t+1}^e(\tau)), d(m_{t+1}, \phi_{t+1}^e(\tau)))$$

represents the *offer* made by the buyer, but this offer is only accepted if it meets the seller's participation constraint under perfect foresight. Thus we define

$$\hat{q}_{t+1} = \hat{q}(m_{t+1}, \phi_{t+1}^e(\tau), \phi_{t+1})$$

to be the realized trade, where

$$\hat{q}_{t+1} = \begin{cases} q(m_{t+1}, \phi_{t+1}^e(\tau)) & \text{if } c(q(m_{t+1}, \phi_{t+1}^e(\tau))) \leq d(m_t, \phi_{t+1}^e(\tau))\phi_{t+1} \\ 0 & \text{else} \end{cases}.$$

Then the corresponding equilibrium supply of specialized goods is given by

$$q_{t+1} = \sum_{\tau=1}^N n_t \left(\tau, \phi_{t-1}, \{\phi_{t-1}^e(\omega)\}_{\omega=1}^N \right) \hat{q}(m(\phi_t, \phi_{t+1}^e(\tau)), \phi_{t+1}^e(\tau), \phi_{t+1}). \quad (9)$$

This expression for trade in the DM takes into account the possibility that the offers of some types of buyers may not be accepted by sellers. We are now ready to define a heterogeneous expectations equilibrium.

Definition 2 *A heterogeneous expectations equilibrium are the sequences $\{q_t, \phi_t, n_t(\tau)\}$ satisfying (7)-(9).*

The remainder of the paper focuses on the implications of heterogeneous beliefs for the equilibrium dynamics. It is convenient now to make a specific assumption about the set of predictors available to buyers. As a simple example, we assume that buyers form their forecasts from an adaptive learning rule that specifies any deviations from the monetary steady state as mean reverting. In particular, we assume that agents perceive the price process to be given by the perceived law of motion:

$$\phi_t = \bar{\phi} + \theta(\phi_{t-1} - \bar{\phi}). \quad (10)$$

We restrict $\theta \in [0, 1]$. This adaptive learning rule is well-specified as it nests the monetary steady state and is able to capture deviations from steady state along a dynamic equilibrium

¹⁵This is reminiscent of the temporary general equilibrium theory of Grandmont (1977).

path.¹⁶ Buyers form forecasts by projecting this forecasting model given the observed values for ϕ . Buyers forecasts, made after the close of the time $t - 1$ CM, are given by

$$\phi_t^e \equiv \phi_t^e(\phi_{t-1}) \equiv \bar{\phi} + \theta(\phi_{t-1} - \bar{\phi}) \quad (11)$$

$$\phi_{t+1}^e \equiv \phi_{t+1}^e(\phi_{t-1}) \equiv \bar{\phi} + \theta^2(\phi_{t-1} - \bar{\phi}). \quad (12)$$

Notice that $\phi_{t+1}^e(\bar{\phi}) = \bar{\phi}$ which implies that the adaptive predictor fixes the steady-state equilibrium.

A brief remark about this adaptive predictor. The forecasting rule (10) is consistent with a Bayesian learning rule in a stochastic environment for appropriate priors on the underlying stochastic process (see Ljungqvist and Sargent (2004)).¹⁷ The present environment, though, is non-stochastic and so it is not natural to have agents econometrically estimate the parameters of their model.¹⁸ Thus, the forecast rule (10) is in the spirit of adaptive learning rules where agents know the steady-state value of ϕ but are uncertain about the transition path.

The remainder of the paper makes the following simplifying functional assumptions. We assume that $c(q) = q$ and that $u = \frac{(q+b)^{1-\alpha} - b^{1-\alpha}}{1-\alpha}$, $\alpha > 0$.¹⁹ In order to highlight the role of beliefs, the subsequent analysis sets $\sigma = 1$, shutting down the search friction. Thus, when positive rejection rates arise, i.e. partial acceptability of money, it is clear that the endogenous role heterogeneous beliefs play in magnifying search frictions.

3.2 Steady-state Properties

In Section 4, we study the dynamic properties of the model with Bayesian offers. In that section, we consider the case in which buyers subjective distribution over sellers' beliefs ϕ^s is distributed $N(\phi^b, \Sigma)$, where Σ parameterizes buyer uncertainty over sellers' participation constraint. The baseline analysis allows adaptive agents to learn about the mean of the distribution, i.e. they update ϕ^b , but treat Σ as a fixed belief parameter. We take this to be the baseline case for two reasons: first, we are interested in how uncertainty affects monetary equilibria; second, for technical convenience. We do show, however, that many of the qualitative results are robust to an extension where adaptive agents also update Σ in real-time as the conditional variance of ϕ_t . In the latter case, when buyers learn about the first and second moments of the CM price process, the dynamic system with heterogeneous

¹⁶There is considerable empirical evidence that subjects in “learning to forecast” experiments make use of simple rules along the lines of (10): see Hommes (2013).

¹⁷Alternatively, one could assume that $\bar{\phi}$ is not known and reinterpret (11) as an AR(1) econometric model for price with $\bar{\phi}, \theta$ estimated recursively in real-time. Baranowski (2012) considers the case where buyers and sellers generate forecasts from a recursively estimated model of the form (11) and finds that the monetary steady state is locally stable under this alternative learning rule.

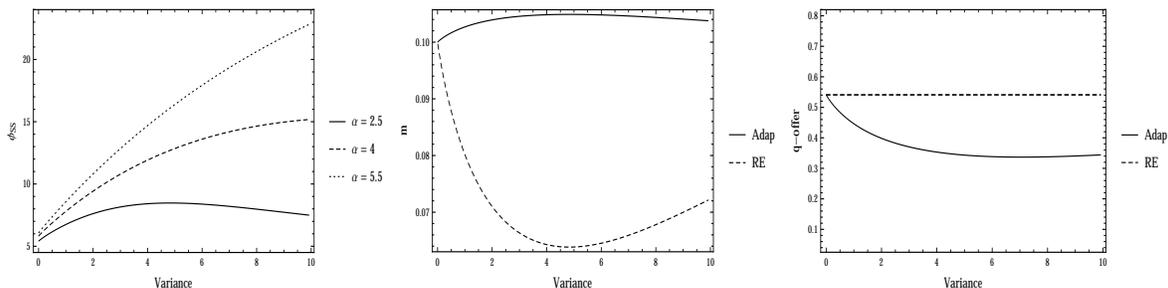
¹⁸Additionally, if agents were try to econometrically estimate the slope coefficient in an AR(1) model in a non-stochastic environment they would suffer from an extreme multicollinearity problem.

¹⁹Taking u to be CRRA can be problematic because utility is not defined at zero. This issue arises in case there is a non-zero search friction or if heterogeneity of beliefs leads to a match where no trade occurs. The nature of equilibrium dynamics in search models are sensitive to the value of b . Throughout, we choose values of b so that the only stable perfect foresight equilibrium in the limit is the steady-state. That is, to consider model parameterizations that do not lead to stable periodic and complicated dynamics under rational expectations.

beliefs has the same steady-state as the model with rational expectations since in a steady-state agents learn that $\Sigma = 0$. In the baseline version of the model, though, adaptive buyers no longer make the same offer as buyers with perfect foresight, which leads to different steady-state implications for the model. As a means of providing intuition for the results that follow in Section 4, this section examines the comparative static effects of uncertainty on the steady state.

Figure 4 plots the comparative statics of a change in uncertainty, measured by the variance parameter Σ . To generate this figure, we set $M = .1$, $\beta = .5$, and $\gamma = 2$. We normalize the cost to the adaptive predictor to zero, and set the cost of the rational predictor to $C = 1$, implying a steady-state fraction of rational agents $\bar{n} \approx .12$. The left-most plot computes the steady-state CM price $\bar{\phi}$ as a function of Σ under three different specifications for the curvature of the utility function $\alpha = 2.5, 4, 5.5$. The middle and right plots compute the money holdings and DM offers separated by agent-type (i.e. adaptive or rational) when $\alpha = 4$.

Figure 4: Effects of uncertainty in steady state



Section 2.4.3 anticipates the non-monotonic effects of uncertainty demonstrated in Figure 4. A higher degree of uncertainty, provided that the level of uncertainty is not too great, leads to a higher steady-state price. Moreover, when there is more curvature in the DM utility function (higher α) the effect of uncertainty on price is more pronounced. Thus, uncertainty leads to something analogous to an “uncertainty premium” in the price of money. What accounts for the effect on steady-state price? The middle and right-most panels provide the intuition. For the model parameterization generating Figure 4, the steady-state fraction of adaptive agents is 0.88 and, according to the figure, these agents hold more money in exchange for a lower quantity than the rational agents. Moreover, for sufficiently low levels of uncertainty, the comparative static effect of uncertainty is, as expected, to increase money holdings in exchange for a lower quantity. These two effects are equivalent to a shift in the money-demand function for adaptive agents, which leads to the higher price in the figure – the uncertainty premium. Unlike the typical search friction, these agents do not hold more money balances as a way to self-insure against an idiosyncratic liquidity shock. They hold demand for money as insurance against the acceptance shock that arises when the buyer makes an offer that the seller rejects. Interestingly, in the right-most plot of Figure 4, the rational agents continue to make the same steady-state offer as if they were in the rational expectations equilibrium: an increase in steady-state ϕ leads to a proportional decrease in money holdings by rational agents. This implies that aggregate DM consumption is lower in the steady-state with heterogeneous buyers who make Bayesian offers than in the rational

expectations equilibrium.

3.3 Dynamic monetary equilibria

We now turn to the dynamic implications of our model. We begin by adopting the take-it-or-leave-it protocol with dogmatic prior beliefs about seller's beliefs in order to maintain tractability, provide a set of analytic stability results, and provide intuition for the main results that follow. We note that even though this section assumes buyer-takes-all bargaining, all qualitative stability results are robust to Bayesian offers and Nash bargaining.²⁰

3.3.1 Baseline Case: Homogeneous Expectations

As a baseline, this section assumes homogeneous expectations across buyers: we shut down the predictor dynamics. With homogeneous buyers' beliefs, the buyer's money demand is given by

$$m(\phi_t, \phi_{t+1}^e) = \beta^{\frac{1}{\alpha}} \phi_t^{-\frac{1}{\alpha}} (\phi_{t+1}^e)^{\frac{1-\alpha}{\alpha}}. \quad (13)$$

Since perfect foresight and adaptive learning forecast the same price in steady state, the monetary steady state in both cases is computed as the solution to $m(\bar{\phi}, \bar{\phi}) = M$, which yields $\bar{\phi} = \beta^{\frac{1}{\alpha}} M^{-1}$.

Under perfect foresight, buyers perfectly predict future prices, i.e. $\phi_{t+1}^e = \phi_{t+1}$. Setting demand equal to M and solving, we get a univariate dynamic system that determines equilibrium price paths:

$$\phi_t = \beta^{\frac{1}{\alpha-1}} M^{\frac{\alpha}{1-\alpha}} \phi_{t-1}^{\frac{1}{1-\alpha}} \equiv f(\phi_{t-1}). \quad (14)$$

Evaluated at the steady state, the derivative of the right-hand-side is $\frac{1}{1-\alpha}$.

Under adaptive learning, equilibrium dynamics are derived by inserting (11) into money demand (13) and imposing market clearing. We obtain

$$\phi_t = \frac{\beta}{M^\alpha} (\bar{\phi} + \theta^2(\phi_{t-1} - \bar{\phi}))^{1-\alpha} \equiv f(\phi_{t-1}, \alpha). \quad (15)$$

Notice that $\bar{\phi}$ is still a steady state of the system.²¹ Differentiating, we get

$$f_\phi(\bar{\phi}, \alpha) = (1 - \alpha)\theta^2,$$

which implies that for $\theta \in (0, 1]$, the steady state is dynamically stable provided that

$$0 < \alpha < \frac{1 + \theta^2}{\theta^2} \equiv \alpha_c.$$

²⁰In other search environments, the bargaining protocol can make a difference for the existence of periodic and aperiodic dynamics. For example, in a search model with imperfectly enforced credit contracts, Gu, Mattesini, Monnet, and Wright (2012) show that periodic and complicated dynamics can arise under Nash bargaining or competitive pricing but not under buyer-takes-all or proportional bargaining. In our model, it is the attracting-repelling dynamics of beliefs that lead to complex dynamics and these forces are present under all of the commonly employed bargaining protocols.

²¹Depending on the calibration, there may be other steady states; we neglect these as their attainment would require not only systematic forecast errors, but precisely the same error, every period.

At $\alpha = \alpha_c$, the steady state destabilizes as the system's eigenvalue crosses negative one.

The following result is an immediate consequence.

Proposition 3 *Under perfect foresight or adaptive learning, there exists two steady-state equilibria, an autarky steady state where $\bar{\phi} = 0$ and a monetary steady state with $\bar{\phi} = \beta^{1/\alpha} M^{-1}$.*

1. **Perfect foresight.** *When $0 < \alpha < 2$, the monetary steady state is locally determinate and the autarkic steady state is locally indeterminate. When $\alpha > 2$, the monetary steady state is locally indeterminate and the autarkic steady state is locally determinate.*
2. **Adaptive learning.** *There exists $\alpha_c \equiv 1 + \theta^2/\theta^2$ such that, for $0 < \alpha < \alpha_c$, the monetary steady state is locally stable, and for $\alpha > \alpha_c$ the monetary steady state is unstable.*

For $\alpha < 2$, the monetary steady state is determinate and there is a (locally) unique perfect foresight equilibrium given by $\phi_t = \bar{\phi}$. Because it is determinate, the steady state is also dynamically unstable with price paths starting near, but not equal to $\bar{\phi}$ that will, at least initially, move away from the steady state. On the other hand, for $\alpha > 2$, the steady state is indeterminate implying that there are, locally, many perfect foresight equilibria: there is an open set U of the steady state so that for any open set $\bar{\phi} \in V \subset U$ and any initial condition $\phi_0 \in V$, the associated equilibrium price path stays in U (and even converges to $\bar{\phi}$). Indeed, because it is indeterminate, the steady state is dynamically stable.

Intuitively, indeterminacy and, hence, a dynamically stable perfect foresight steady-state monetary equilibrium, arises when there is strong expectational feedback. All values of α in our model correspond to negative expectational feedback: agents want to smooth DM-good consumption, and an expected price increase reduces the required money holdings needed to finance trade in the DM, i.e. the liquidity property of money increases so that buyers do not need to carry high money balances. Larger values of α induce stronger feedback and if $\alpha > 2$ then this negative feedback is sufficiently strong to generate indeterminacy.

When $\alpha < 2$, the perfect-foresight steady state is locally unique; however there may be equilibrium paths which remain relatively near the steady state. The assessment of these paths can be achieved using bifurcation analysis. As a dynamic system, the steady state destabilizes when α crosses the critical value 2 from above. Since $f'(\bar{\phi}) = -1$ when $\alpha = 2$, a flip (or, period doubling) bifurcation is suggested, but the system does not meet the regularity conditions needed to identify this primary bifurcation.²² Numerical analysis suggests that for $\alpha < 2$, all non-steady state perfect foresight paths are globally as well as locally explosive.

Other parameterizations of the model can produce period-doubling bifurcations. For example, Nosal and Rocheteau (2011, chp. 4) is identical to the perfect foresight version of the model under consideration here, and they show the existence of 2-cycles with $\alpha > 2$ and the coefficient b in the utility function is sufficiently positive. Lagos and Wright

²²If $\phi \rightarrow f(\phi, \alpha)$ captures the dynamics (14) then a generic flip bifurcation requires $1/2f_{\phi\phi}^2 + 1/3f_{\phi\phi\phi} \neq 0$, which fails in our model.

(2003) demonstrate the possibility of periodic and aperiodic cycles in the version with Nash bargaining.

Comparing the regions of dynamic stability associated to perfect foresight and adaptive learning yields a nice example of an important phenomenon in the adaptive learning literature: stability reversal. Intuitively, perfect foresight requires forward-looking behavior and adaptive forecasting relies on backward-looking behavior. Therefore, forward stability indicates backwards instability, and vice-versa. As a precise and extreme case, consider $\theta = 1$, so that agents have naive expectations: $\phi_{t+1}^e = \phi_{t-1}$. In this case the adaptive model is dynamically stable when $\alpha < 2$, which is the complement of the stability region for the perfect foresight model. The possibility of stability reversal plays a prominent role in explaining the dynamic behavior associated to the model with heterogeneous expectations: the tension between stability and instability creates regimes that alternate between attraction and repulsion, and in the process, create complex dynamics.

Corollary 4 *Under adaptive learning, with $\theta < 1$, as α crosses α_c from below, there is a flip bifurcation and there exists a locally unique cycle of period two.*

The proof is contained in the Appendix.²³ This corollary provides a nice first example of new, complex dynamics induced by expectations heterogeneity, even with expectations homogeneity among buyers. When buyers hold perfect-foresight beliefs and $\alpha < 2$ so that the monetary steady state is dynamically unstable, there are no additional non-explosive equilibrium paths. In contrast, when buyers have adaptive expectations and when $\alpha_c < \alpha < \alpha_c + \varepsilon$ for some $\varepsilon > 0$, so that the monetary steady state is dynamically unstable, there is an additional non-explosive equilibrium path which is characterized as a two-cycle in prices.

3.3.2 Heterogeneous expectations

In the previous section, we observed that for $\alpha > \alpha_c = \frac{1+\theta^2}{\theta^2}$, the adaptive price path does not converge to the steady state; on the other hand, since for all $\theta \in (0, 1]$, $\alpha_c \geq 2$, it follows that when $\alpha > \alpha_c$ the perfect foresight price path does converge to steady state. This dichotomy suggests a role for predictor selection in generating dynamics. The intuition is as follows: assume $\alpha > \alpha_c$ and that most agents have perfect foresight. Then the steady state is dynamically stable, and the price path converges toward it. Assume also that it is costly to formulate perfect foresight. As ϕ_t approaches $\bar{\phi}$, the less costly adaptive predictor forecasts nearly as well as the costly rational predictor, and so an increasing proportion of agents begin using it; when this proportion is high enough, the steady state becomes unstable, and the price path is driven away from $\bar{\phi}$. But, far from $\bar{\phi}$, the rational predictor has strong advantages over the adaptive predictor – strong enough to overcome its higher cost. Therefore more and more agents switch to the rational predictor, which ultimately stabilizes the steady state, and the process starts over. This section combines analytic and numerical analysis to demonstrate that these attractor-repellor dynamics exist in the monetary search model with heterogeneous

²³Numerical analysis indicates that, as α continues to increase, the amplitude of the 2-cycle becomes too large to guarantee $\phi_t > \beta\phi_{t+1}^e$.

beliefs. The subsequent section shows that these non-stationary equilibrium dynamics have important welfare implications and provide insights into experimental findings.

Normalizing the cost of adaptive beliefs to zero, denoting the cost of the rational predictor by C , and letting n_t be the proportion of buyers choosing the rational predictor at time t , we have

$$n_t = n(\phi_{t-1}, \phi_{t-2}) = \frac{e^{-\gamma C}}{e^{-\gamma C} + e^{-\gamma[\phi_{t-1}^e(\phi_{t-2}) - \phi_{t-1}]^2}}, \quad (16)$$

where $\phi_{t-1}^e(\phi_{t-2}) = \bar{\phi} + \theta(\phi_{t-2} - \bar{\phi})$ is the most recent one-period ahead adaptive forecast. Market clearing is given by

$$n(\phi_{t-1}, \phi_{t-2})m(\phi_t, \phi_{t+1}) + (1 - n(\phi_{t-1}, \phi_{t-2}))m(\phi_t, \phi_{t+1}^e(\phi_{t-1})) = M,$$

where m is money demand. Using the functional form for demand given by (13), and solving for ϕ_{t+1} , we obtain the following dynamic system:

$$\phi_{t+1} = g(\phi_t, \phi_{t-1}, \phi_{t-2}), \quad (17)$$

where

$$\begin{aligned} g(\phi_t, \phi_{t-1}, \phi_{t-2}) &= \beta^{\frac{1}{\alpha-1}} \phi_t^{\frac{1}{1-\alpha}} n(\phi_{t-1}, \phi_{t-2})^{\frac{\alpha}{\alpha-1}} \\ &\times \left\{ M - [1 - n(\phi_{t-1}, \phi_{t-2})] m(\phi_t, \phi_{t+1}^e(\phi_{t-1})) \right\}^{\frac{\alpha}{1-\alpha}}. \end{aligned}$$

The monetary search model with heterogeneous expectations is a 3rd-order non-linear difference equation, and the types of stationary and non-stationary equilibria can be studied in the usual way. As expected, $\bar{\phi}$ is a fixed point of this system. Moreover, in a stationary equilibrium the value of $\bar{n} = n(\bar{\phi}, \bar{\phi})$ is given by

$$\frac{e^{-\gamma C}}{e^{-\gamma C} + 1}.$$

The relative cost of rational expectations, C , and the intensity of choice, γ , control the steady state fraction of rational agents. Larger values of γ lead to a lower steady-state fraction of rational agents. Thus, we expect that the dynamic stability of monetary equilibria will depend on γ , making it a natural bifurcation parameter.

In the special cases of steady-state or naive expectations, i.e. $\theta = 0$ and $\theta = 1$, respectively, we are able to provide a precise analytic steady-state stability result.

Proposition 5 *Consider the model (16)-(17) with perfect foresight and adaptive beliefs.*

1. *When agents dynamically choose between perfect foresight and steady-state expectations, i.e. $\theta = 0$, $\exists \gamma_c > 0$ such that the monetary steady state is locally stable if and only if $\gamma < \gamma_c$. Moreover, as γ crosses γ_c from below there is a flip bifurcation and there exists a locally unique stable two-cycle.*
2. *When agents dynamically choose between perfect foresight and naive expectations (i.e. $\theta = 1$), the monetary steady state is locally unstable.*

The proof is in the Appendix.

Consider the first part of Proposition 5, when $\theta = 0$, so that adaptive agents always forecast the steady state. Also, here we assume that $\alpha > 2$ so that the perfect-foresight steady state is dynamically stable. Writing the system as $\phi_t \rightarrow h(\phi_t, \gamma)$ to emphasize that we now treat the intensity of choice, γ , as the bifurcation parameter, we may compute

$$h_\phi(\bar{\phi}, \gamma) = \frac{1 + e^{\gamma C}}{\alpha - 1},$$

which yields stability provided that

$$\gamma < \gamma_c \equiv \frac{\log(\alpha - 2)}{C}.$$

Intuitively, low values of γ impart two-fold stabilizing effects: first, even in steady state, there is still a large proportion of rational agents – $n = 1/2$ when $\gamma = 0$ – and thus the dynamic stability of perfect foresight is a dominant force; second, small γ imparts only a moderate inclination to switch predictors even when, net of the cost, the adaptive predictor is superior. Below γ_c , these stabilizing tendencies prevail and guide the economy to the monetary steady state; however, as γ crosses γ_c from below, the intensity of choice is higher and buyers will switch more quickly to the better performing forecasting model and this can destabilize the steady state. The strong switching intensity also implies that as the economy moves further from the steady state they are quick to switch predictors again. These repelling and attracting forces, that arise with a high intensity of choice γ , can lead to complex dynamics. We demonstrate this possibility with numerical analysis.

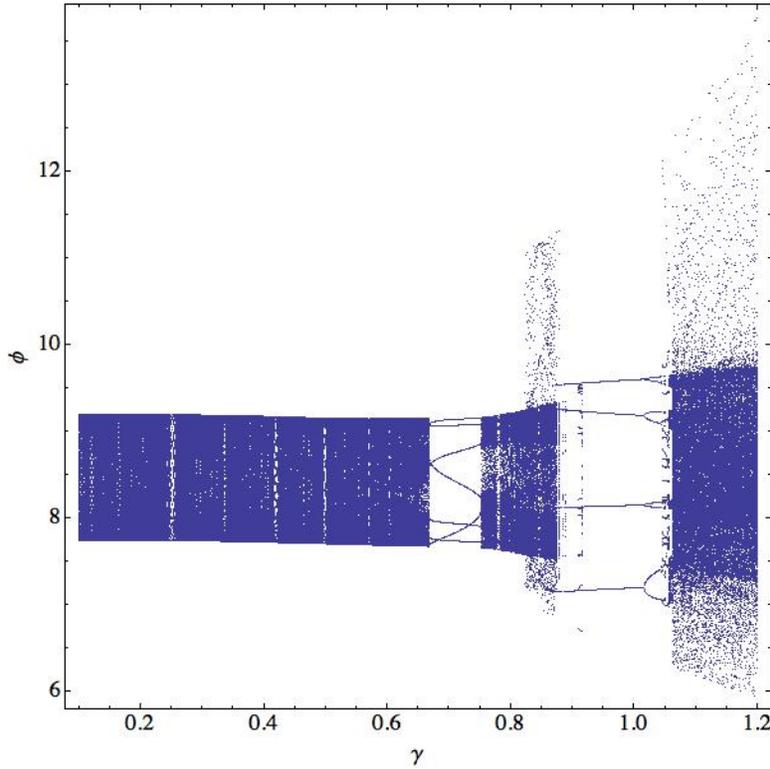
Now consider the second part of Proposition 5, the case of a “naive” predictor: $\phi_{t+1}^e = \phi_{t-1}$, i.e. $\theta = 1$. In this case, and in the sequel, the system’s dimension is three, so that identification of the primary bifurcation requires a center manifold reduction that is intractable in the present environment.²⁴ However, the proposition establishes that the steady-state is unstable.

Lower values of α – specifically, $\alpha < 2$ – tend to result in globally explosive behavior. This makes sense since when $\alpha < 2$ the rational model is determinate, and thus dynamically unstable. As the economy moves further away from the steady state buyers will switch to the rational predictor which will drive the economy even farther from the steady state. On the other hand, higher values of α yield bounded, complex price paths.²⁵ This is illustrated in Figure 5 which provides a bifurcation diagram where the intensity of choice γ is the bifurcation parameter. Here $M = .1$, $\beta = .5$, $C = 1$ and $\alpha = 4$. Notice that for all values of γ , the equilibrium price path remains bounded around the steady-state monetary equilibrium, and as γ varies, the dynamics shift between periodic and more complicated patterns.

²⁴The center manifold is an invariant manifold of the (re-centered) dynamic system, whose tangent space at the steady state is spanned by the eigenspaces of the system’s derivative corresponding to the eigenvalues of unit modulus. The center manifold theorem shows that the dynamics of any solution near the steady state is determined by the dynamics projected onto the center manifold. For a system of coupled dynamic equations one can characterize these solutions through a reduction of the dimension of the system by restricting paths to lie on the center manifold. In practice, such a center manifold reduction can be intractable to implement.

²⁵This case is anticipated by the non-linear dynamics in monetary search models with perfect foresight. In these models, large values of α yield a non-monotonic dynamic map: see Nosal and Rocheteau (2011).

Figure 5: Bifurcation Diagram: $\theta = 1$



More generally, by setting $\theta \in (0, 1)$, we are now assuming that non-rational agents adopt an adaptive forecasting model. As in the naive case, the system is three dimensional, however here the stationary monetary equilibrium may be stable. Furthermore, the nature of the primary bifurcation will depend on the adaptive parameter θ .

As an illustration, consider small values of θ . Recall that for $\theta = 0$, the primary bifurcation is a flip (period-doubling). We would expect to see a flip bifurcation for small positive θ . This appears to indeed be the case: setting $\alpha = 4$ and $\theta = .1$ (and holding $M = .1$, $\beta = .5$ and $C = 1$) we may compute that

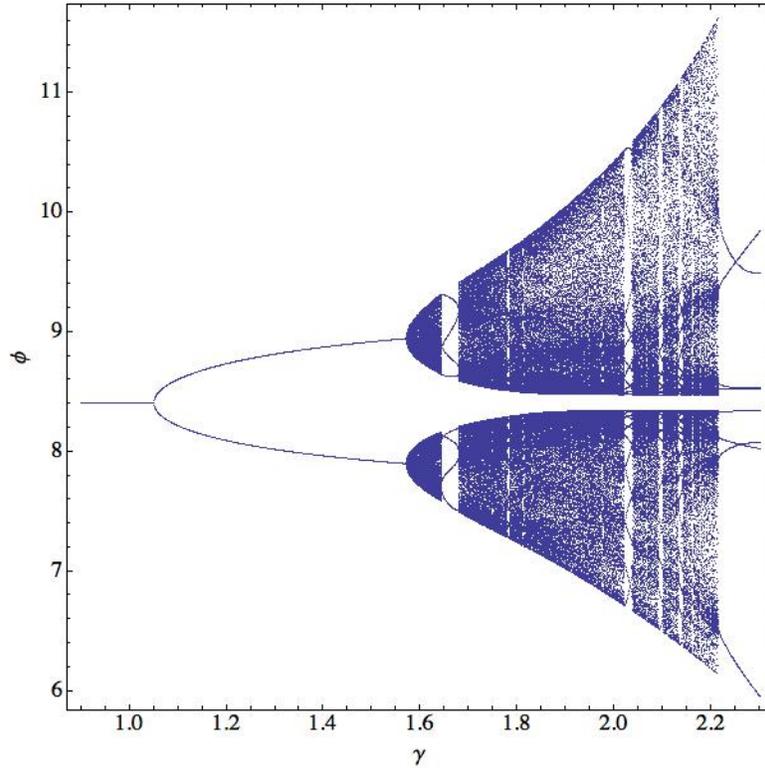
$$\gamma_c = 2 \log(2) + \log(5) - \log(7) \approx 1.05.$$

As in the steady-state beliefs case, for $\gamma < \gamma_c$ the steady-state equilibrium is stable; at γ_c one of the non-zero eigenvalues crosses the unit circle at -1 , which destabilizes the steady state consistent with a flip bifurcation. A precise characterization again requires an intractable center manifold reduction, and, as before we use numerical analysis to identify the primary bifurcation. Numerical evidence indicates that a stable two-cycle does emerge: see Figure 6. Also, as in the case $\theta = 0$, as γ increases, the bifurcation diagram in Figure 6 indicates that more complex dynamics arise.

For larger values of θ , destabilization obtains as the eigenvalues cross the unit circle with non-trivial imaginary components, which suggests a Neimark-Sacker bifurcation.²⁶ A

²⁶A Neimark-Sacker bifurcation is the discrete time equivalent of a Hopf bifurcation, where a pair of complex eigenvalues cross the unit circle.

Figure 6: Bifurcation Diagram: $\theta = .1$



Neimark-Sacker bifurcation implies the existence of a stable closed invariant curve. And, in fact, this is precisely what we see. Setting the parameter values as above, but letting $\theta = 1/2$ yields

$$\gamma_c = \log(2) \approx 0.69.$$

Figure 7 provides the associated bifurcation diagram, and, as anticipated, as γ crosses γ_c from below, the steady state destabilizes, and complex behavior emerges.²⁷ As γ increases further, the orbit structure shifts between periodic and more complex patterns, as in previous examples. Figures 8 and 9 provide sample attractors.

3.4 Discussion

The wide range of equilibrium results presented above warrants a brief review and discussion. We began our analysis by looking at the steady-state properties of the equilibrium with Bayesian offers. There we saw that buyer uncertainty about sellers' beliefs leads buyers to make more cautious offers with higher payments in exchange for lower quantities. In equilibria with a large enough fraction of adaptive agents, this uncertainty leads to a higher steady-state price of money relative to a steady-state without uncertainty, i.e. an uncertainty premium. While such a steady-state would be unlikely to be observed in practice since one

²⁷In fact, the orbits may be periodic depending on the “rotation angle,” see Kuznetsov (1998) for details.

Figure 7: Bifurcation Diagram: $\theta = 1/2$

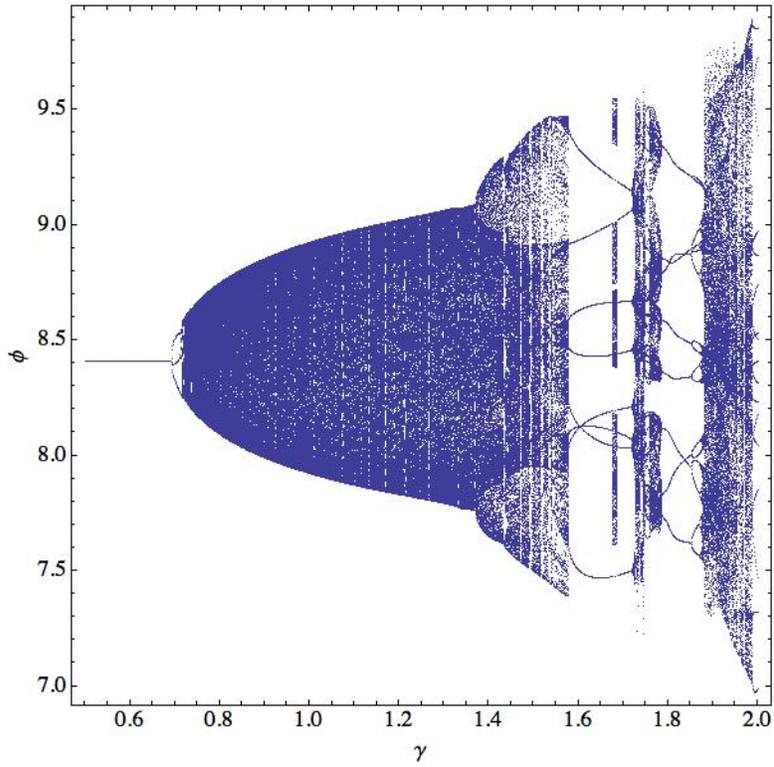
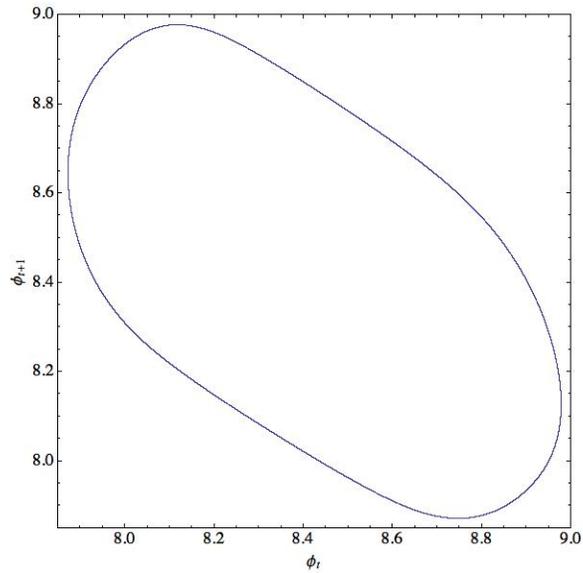


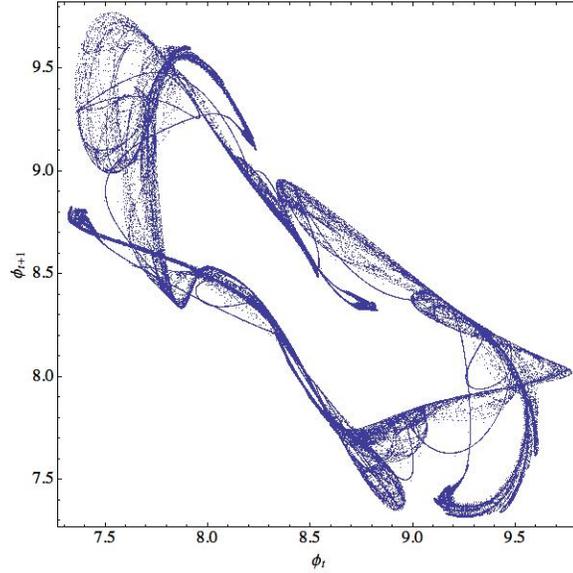
Figure 8: Attractor: $\theta = 1/2, \gamma = 1.10$



would expect Buyer uncertainty to vanish in a stationary equilibrium, these results are useful for interpreting the results in Sections 4 and 5.

Turning next to the analysis of equilibrium dynamics, we adopted the take-it-or-leave-it

Figure 9: Attractor Diagram: $\theta = 1/2, \gamma = 1.92$



trading protocol when buyers have dogmatic beliefs. Under homogeneous rational expectations, the implied perfect-foresight dynamics are quite simple: if $\alpha < 2$ there is a unique non-explosive equilibrium price path given by the model's steady state, and if $\alpha > 2$ there are many price paths, and they all converge to the steady state. In particular, the model's long run dynamics are trivial and independent of α . This tame behavior is in sharp contrast to that implied by heterogeneous expectations. To organize our review, we form a list:

1. **Buyer homogeneity.** Assuming all buyers adopt the forecast model given by (10), the equilibrium dynamics of the price path are determined by $\alpha_c \geq 2$. Small values of α , corresponding to buyers who are less concerned about consumption smoothing, lead to convergent price paths: the long run behavior is coincident to that implied by homogeneous expectations among buyers and sellers with large α . When $\alpha > \alpha_c$, that is, when buyers are quite concerned about consumption smoothing, more complicated dynamics may prevail: in particular, near α_c , a stable two-cycle emerges, which implies non-trivial long run behavior.
2. **Buyer heterogeneity: steady-state beliefs.** Now suppose that buyers select between the costly perfect-foresight model and the cheaper steady-state forecast: $\theta = 0$. If $\alpha > 2$, complex long run dynamics may emerge: as the intensity of choice parameter γ increases, the steady state destabilizes and a stable two-cycle emerges. Numerical analysis suggests that for larger values of γ , complex, and even chaotic behavior ensues.
3. **Buyer heterogeneity: naive beliefs.** Now suppose that buyers select between the costly perfect-foresight model and the naive forecast: $\theta = 1$. Here, when $\alpha > 2$, our numerical analysis suggests that complex dynamics emerge for all values of the intensity of choice parameter.
4. **Buyer heterogeneity: adaptive beliefs.** Now suppose that buyers select between

the costly perfect-foresight model and an adaptive forecast: $\theta \in (0, 1)$. In this case, and again when $\alpha > 2$, complex, and even chaotic dynamics result for γ above some threshold level.

For large α , items (1) - (4) may be collectively summarized as follows: homogeneous rational expectations imparts convergent short run dynamics and trivial long run dynamics; heterogeneous expectations robustly leads to complex short run dynamics and nontrivial long run dynamics.

4 Implications for welfare

The main focus of this section is the welfare implications of Bayesian offers and buyer uncertainty. When buyers' beliefs are distributed across heterogeneous forecasting models they will exit the centralized market with different expectations and, hence, nominal money balances. Under the buyer-takes-all bargaining arrangement, heterogeneous buyers will make different offers for two distinct reasons: they have different expectations about the value of money and they carry distinct money balances. Thus, heterogeneity affects the intensive margin of trade. Once matched with a seller (who has perfect foresight), there will be a fraction of the matches that occur between buyers and sellers with different beliefs. When the buyer makes a take-it-or-leave-it offer there is the possibility, depending on the dispersion of beliefs, that the seller will reject the offer. Thus, the heterogeneous belief friction can lead to an extensive margin of trade through the endogenous partial acceptability of money that magnifies search frictions. These frictions only matter along a dynamic equilibrium path as we have restricted all buyers and sellers to have the same beliefs within a steady-state monetary equilibrium. However, the possibility that a seller may not accept the buyer's offer, motivates buyers to alter their take-it-or-leave-it offer by taking into account their subjective acceptance probability: this was the main insight of the previous discussion on Bayesian offers.²⁸

In this section, we study in greater detail how heterogeneity impacts the nature of trade and the welfare properties of monetary equilibria. The focus of this section is on Bayesian offers where buyers learn about the CM price of money by endogenously selecting between perfect foresight, with a cost $C = 1$, and a costless adaptive predictor. Those buyers who forecast with the adaptive predictor have imperfect knowledge about the seller's participation constraint and so, when making a take-it-or-leave-it offer to firms, they place a subjective prior on sellers' beliefs that is normally distributed, centered on the buyer's belief with a variance Σ . By explicitly modeling the subjective acceptance probability, buyers make cautious offers with higher money in payments in exchange for lower quantities of the good, all else equal. This section studies the properties of monetary equilibria in this environment.

In the analysis that follows, we treat Σ , the subjective uncertainty about sellers' beliefs, as a bifurcating parameter. In a sense, buyers are assumed to update the first moment of their subjective priors while remaining dogmatic about the second moment. We make this

²⁸The Appendix presents additional results on the welfare properties of dynamic equilibria when buyers have dogmatic beliefs.

assumption for tractability, though, the Appendix shows that qualitative properties of monetary equilibria with heterogeneous beliefs are robust to an environment where buyers update their beliefs about the first and second moments. The other key bifurcating parameters are γ , the intensity of choice in predictor selection, θ , the gain in the adaptive forecasting rule, and α , the curvature parameter in the DM utility function. We consider the comparative effect of different values for Σ across a range of parameterizations.

The previous section demonstrates that heterogeneous beliefs can lead to cycling and aperiodic fluctuations around a monetary steady-state. In principle, price fluctuations can ease and tighten liquidity constraints and, therefore, impact the welfare of monetary equilibria. We study this issue in this section by computing welfare along a dynamic monetary equilibrium path with heterogeneous beliefs and by studying the different behavior depending on expectations-type. We measure welfare as the average (over time) realized surplus in the DM: i.e. $\mathcal{W} = T^{-1} \sum_{t=0}^T n_t [u(q_t^R) - c(q_t^R)] + (1 - n_t) [u(q_t^A) - c(q_t^A)]$, where q^R, q^A are rational and adaptive buyer bargaining outcomes, respectively, and T is large. Throughout, the point of comparison is \mathcal{W} relative to the average steady-state welfare in a rational expectations equilibrium. In the analysis that follows, we fix the intensity of choice parameter and study the properties of the equilibrium paths for varying degrees of uncertainty Σ .

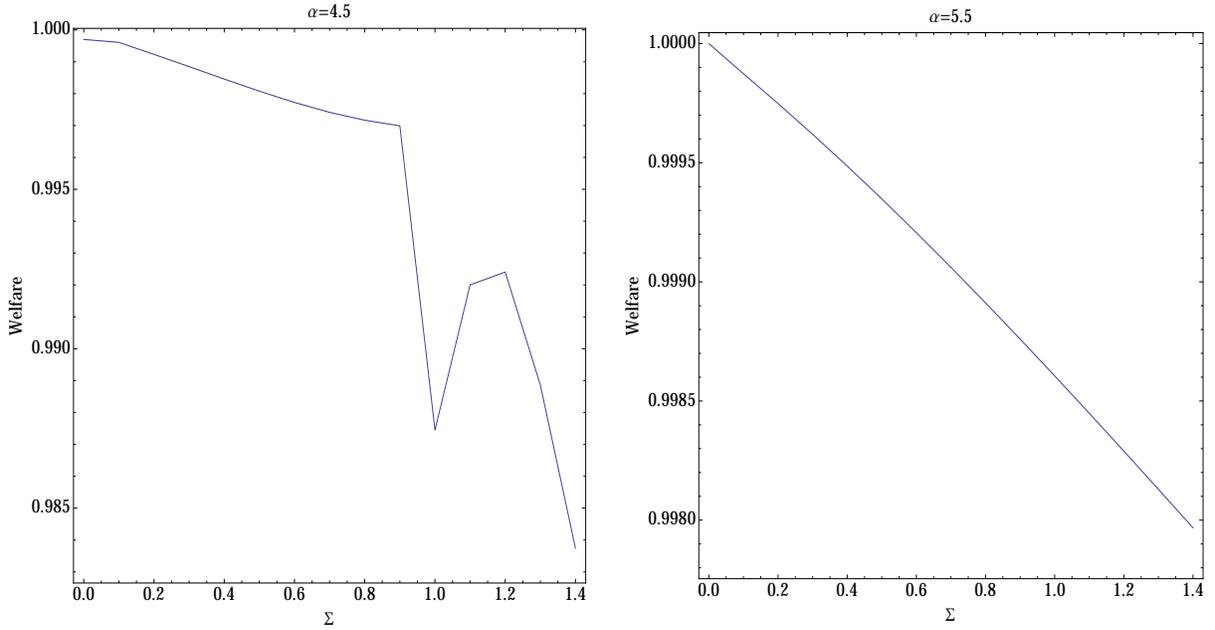
Figures 10 - 11 plots the results when $\gamma = 0.3$, which coincides with a stable steady-state in the dogmatic priors case, so that for low values of Σ the steady state remains stable; however as Σ rises, the system bifurcates with potentially large welfare effects. Figure 10 shows that for moderate degrees of uncertainty, $0 \leq \Sigma \leq 1.4$, in the case of both a low (left plot) and high value (right plot) of α , welfare is decreasing in the degree of uncertainty. When $\alpha = 5.5$, i.e. with a high degree of curvature in DM utility, welfare is monotonically decreasing in uncertainty. For $\alpha = 4.5$, though, the effect of uncertainty is non-monotonic, with a sharp drop in welfare near $\Sigma \approx 0.80$, though, on average, there is a negative relationship between welfare and uncertainty. As we will see below, the non-monotonicities are associated with bifurcations that result from changing the degree of uncertainty.

To better understand the welfare properties in Figure 10, Figure 11 plots the equilibrium state variables as a function of the uncertainty parameter Σ . This figure demonstrates that uncertainty can destabilize a steady-state equilibrium and lead to a series of period-doubling bifurcations. Begin with the middle-left panel, which plots ϕ against Σ . For low values of uncertainty the steady state is stable. As the uncertainty increases past $\Sigma \approx .3$, the steady state destabilizes and a period-doubling cascade ensues.

The downward-sloping nature of the graphs in both panels in Figure 10 is suggested by Figure 11. As uncertainty about sellers' beliefs increase, adaptive buyers offer higher payments of money in exchange for lower quantities of the DM good, and lower consumption in the DM by the adaptive learning buyers leads to the lower welfare.

The dramatic fall in welfare seen in Figure 10, corresponding to $\Sigma \approx 0.8$, is also explained by Figure 11. As the primary bifurcation is crossed, the qualitative features of the offer made by adaptive agents does not change dramatically, and the acceptance rate remains at one; however, at the secondary bifurcation corresponding to the emergence of the 4-cycle, there is a sharp rise in the highest quantity offered (see upper left panel): this offer is rejected by the rational agents, which lowers the acceptance rate and dramatically affects welfare.

Figure 10: Welfare: $\gamma = 0.3, \theta = 0.2$

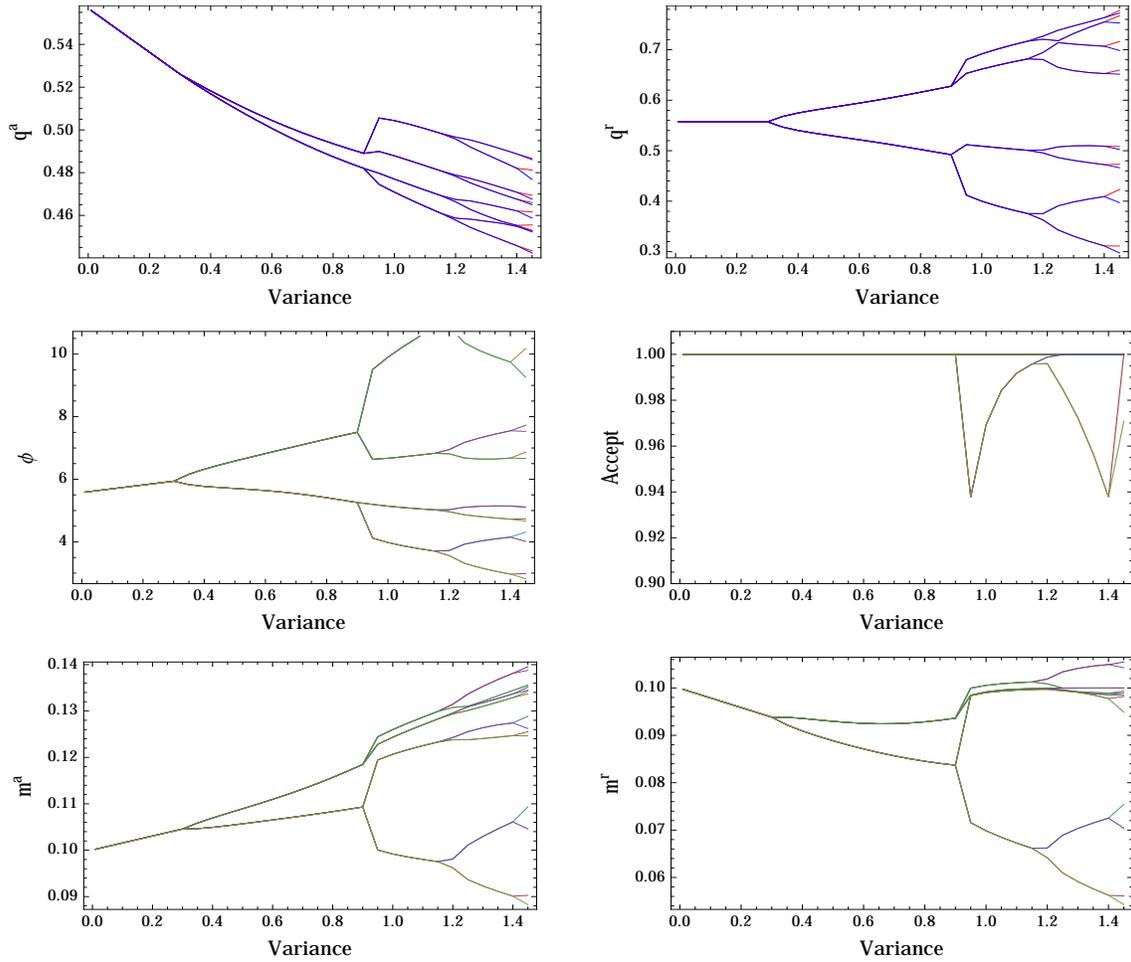


The lower welfare from these adaptive buyers is tempered by periods of higher welfare for the rational buyers. The upper and lower right plots in Figure 11 illustrate the behavior of rational buyers along these equilibrium paths. The model parameterization leads to fluctuations in the fraction of rational that alternates between $n = 1$ and $n < 1/2$, i.e. a majority of adaptive agents. The periods of a high value of money leads to higher DM consumption by rational agents. Though, these buyers hold less money than adaptive the higher values of money relax their liquidity constraints and facilitates greater consumption than they would have in a rational expectations equilibrium. There are periods, however, where rational buyers experience tighter liquidity constraints, lower DM consumption, and so overall welfare is lower along these equilibrium paths.

Figure 12 sets $\gamma = 0.75$, a value for which dogmatic priors lead to a two-cycle when $\sigma = 4.5$, and a stable steady state when $\sigma = 5.5$. Figure 12 plots the welfare as a function of uncertainty, Σ , again for two different values of σ . The key distinction between the two plots is the left plot features a stable two-cycle for small values of Σ , whereas the steady-state is stable on the right. Welfare remains bounded below the value in the rational expectations equilibrium in both cases. A moderate value for uncertainty Σ brings welfare close to the REE value. In the left plot, there is a discrete jump in welfare at $\Sigma \approx 0.30$, as will be seen below, this accords with a bifurcation in the acceptance rate: for values of $.3 \leq \Sigma < .6$ the acceptance rate is 1. Then conditional on an acceptance rate of 1, welfare is decreasing in uncertainty in both the left and right plots. As we see below, this is because higher values of uncertainty have no effect on the extensive margin (acceptance rate), but push down the offers made by buyers. This has a negative impact on welfare.

Figure 13 plots the values of the state variables as a function of uncertainty, again giving

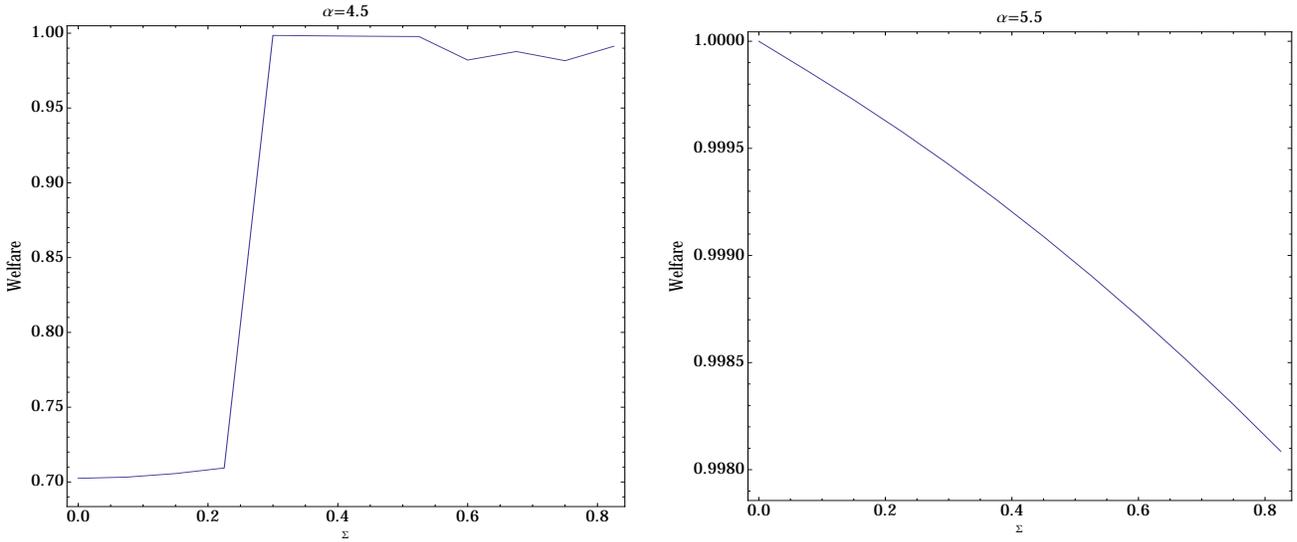
Figure 11: Heterogeneous Beliefs and Acceptance Rates: $\sigma = 4.5, \gamma = 0.3, \theta = 0.2$



insights into the results illustrated in the left plot of Figure 12. For small values of Σ , there is a stable two cycle, with the price ϕ_t alternative between high and low values around the steady-state. Accordingly, for sufficiently small uncertainty, i.e. $\Sigma < .30$, the acceptance rate also follows a two-cycle alternating between an acceptance rate of 1 and approximately 0.4. As we saw in Figure 11, higher uncertainty leads to adaptive buyers making offers with more money for lower quantities of the DM good. At $\Sigma \approx 0.3$, the offers and money holdings by adaptive buyers crosses a threshold where now all offers are accepted by sellers. Meanwhile, the rational agents alternate between higher and lower consumption as variations in ϕ_t relax and tighten liquidity constraints. Interestingly, adaptive buyers' offers do not fluctuate as widely with variations in price compared to the rational buyers. This is because the uncertainty inherent to the Bayesian take-it-or-leave-it offers makes the offer less sensitive to variations in the price.

The acceptance rates plotted in Figure 13 also provide useful insights into the role uncertainty plays in the welfare properties of monetary equilibria with heterogeneous beliefs. The degree of uncertainty about sellers' beliefs is a bifurcation parameter that affects the nature of the equilibrium dynamics. Conditional on a particular qualitative equilibrium

Figure 12: Welfare: $\gamma = 0.75, \theta = 0.2$

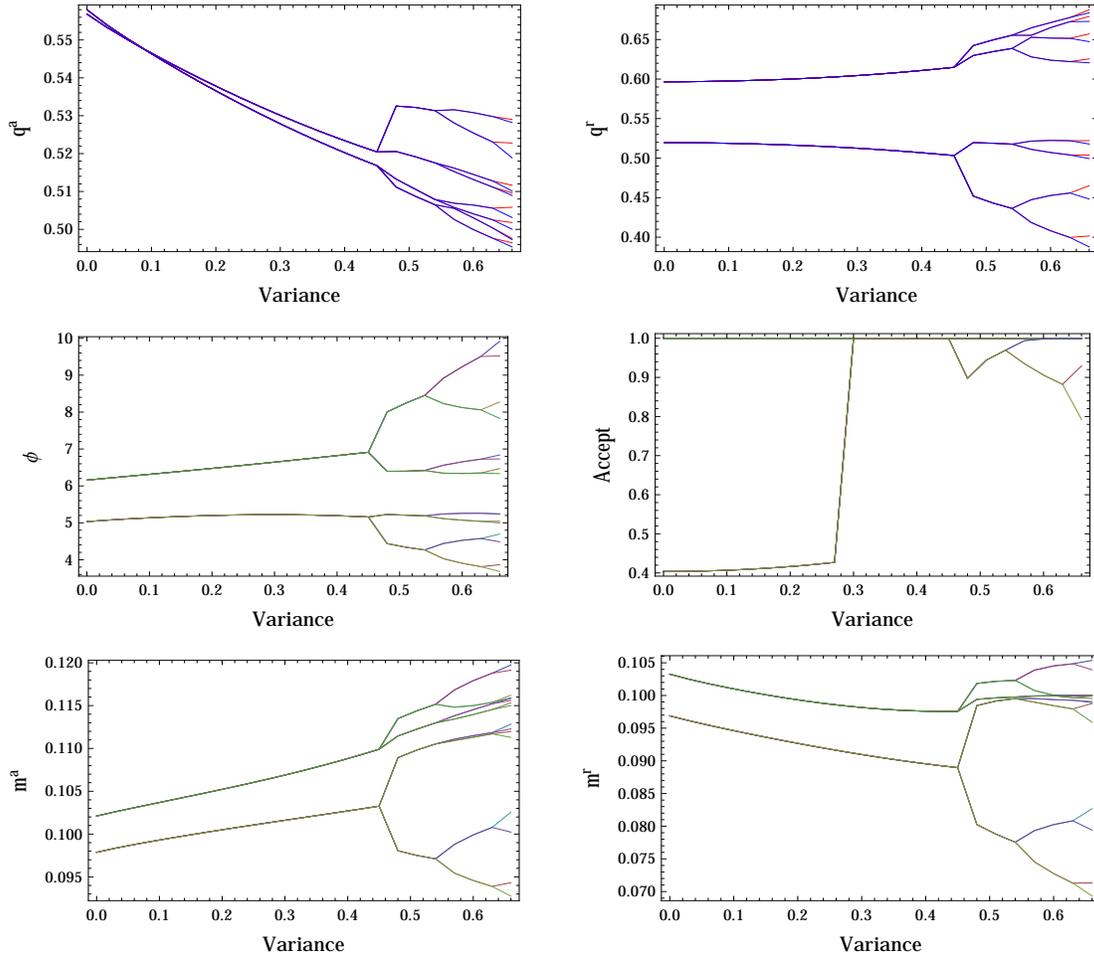


path – that is, hold the dimension of the cycle fixed – higher values of Σ increase the average acceptance rate until the acceptance rate becomes 1 in all periods. For these values of Σ , more uncertainty is welfare improving. Then, still holding the cycle-dimension fixed, subsequent increases in uncertainty are welfare decreasing since there is no impact on the extensive margin, but the caution induced by the Bayesian offers reduces welfare. There is then some critical value of Σ that then bifurcates the system increasing the cycle dimension by two, which again leads to volatility in the acceptance rate. Then the effect of uncertainty on welfare repeats, holding the dimension of the cycle fixed at the new higher value. Thus, the effect of uncertainty on welfare is a complicated balancing of three forces:

1. uncertainty tends to lead to more cautious offers (lower welfare);
2. uncertainty tends to increase the acceptance rate (higher welfare);
3. small changes in uncertainty can bifurcate the equilibrium and, therefore, can impart sharp decreases in the mean acceptance rate, and hence welfare.

The previous section demonstrated that for other values of the intensity of choice γ the equilibrium can exhibit complex dynamics and strange attractors. We also examined the case where $\gamma = 1.3$, a value that in the dogmatic case featured complex dynamics and plots the welfare again as a function of uncertainty Σ . Here we briefly describe the welfare implications of strange attractors. For $\sigma = 4.5$, uncertainty has a non-monotonic effect on welfare as small changes in Σ frequently bifurcate the strange attractors altering the nature of the complex dynamics. The same behavioral patterns emerge as the previous figures, except now there can be substantially more variation in acceptance rates for a broad range of uncertainty parameters. In particular, high values of Σ lead to substantial variation in the CM price ϕ_t and, correspondingly, consumption for the rational agents that at times approach the efficient level of output in the DM.

Figure 13: Heterogeneous Beliefs and Acceptance Rates: $\sigma = 4.5, \gamma = 0.75, \theta = 0.2$



In this simple example of two types, with all sellers having perfect foresight, it may seem obvious that buyers will learn the sellers' acceptance rule. However, in a richer setting there will be many different classes of buyers and sellers, each differing by their forecasting rule. Each buyer of a certain type will only learn about a particular seller-type's acceptance rule after enough matches have taken place. It is reasonable to expect then that such a learning process would be slow and that the demonstrated dynamics could persist for finite, and possibly long, stretches of time.²⁹

In summary, heterogeneous beliefs can lead to monetary equilibria that are welfare reducing, relative to a rational expectations equilibrium. Although, the effect on individual agent behavior differs by expectations-type, we find that welfare is generally reduced because fluctuating beliefs can lead to a partial acceptability of money in some matches. These results, though, depend on the uncertainty in buyers' prior beliefs about sellers' acceptance rule. The effect of uncertainty on welfare depends on a delicate balance between its effect on the intensive margin, the acceptance rate, and the stability of the monetary equilibrium. The next section shows that these insights can be useful for interpreting experimental evidence.

²⁹We leave such a learning model to future research.

5 Implications for experimental evidence

Besides being of theoretical interest, the results presented above can provide a means to interpret key experimental findings in Duffy and Puzzello (2013). In particular, the results from a Lagos-Wright laboratory experiment in Duffy and Puzzello (2013) include the following:

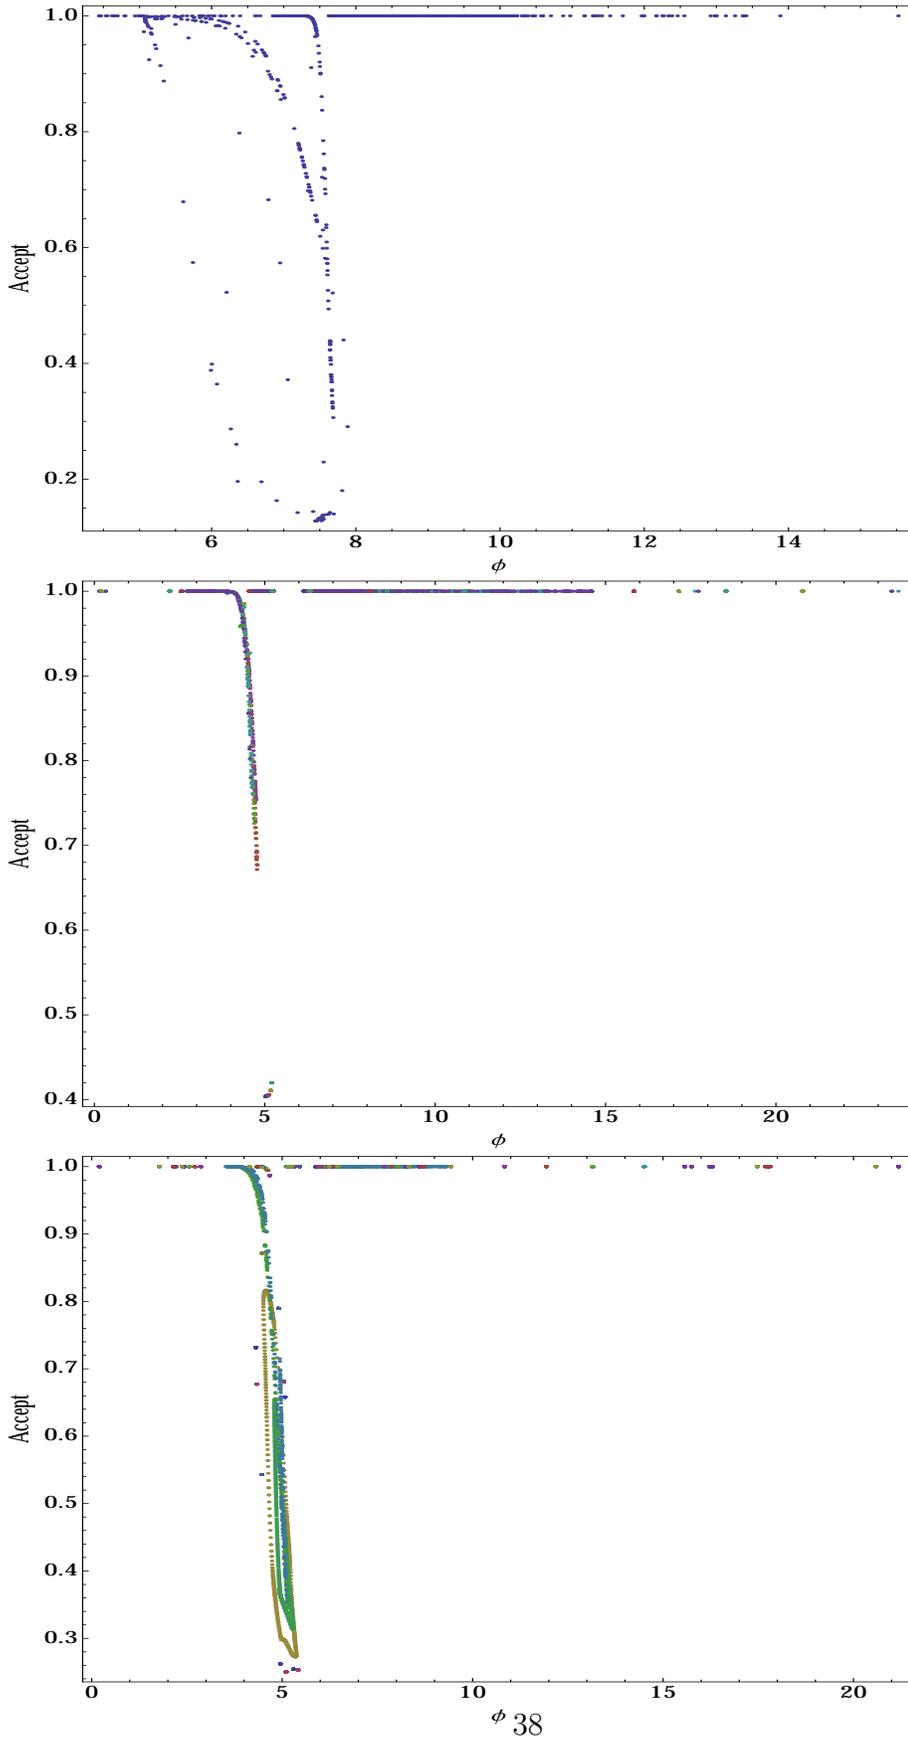
1. Over 95% of bilateral trades involve fiat money.
2. Approximately 40-60% of all buyer offers are rejected by sellers.
3. The likelihood that a buyer's offer will be accepted decreases as the quantity of the good requested increases.
4. There is a non-degenerate distribution of money holdings at the end of centralized market rounds despite the fact that subjects rebalance their portfolios in the centralized market.

These experimental results provide support in favor of monetary equilibria but cannot be completely explained by the benchmark Lagos-Wright model. This section explores one possible avenue to account for the experimental findings: the Lagos-Wright model extended to include heterogeneous beliefs. On the surface, the model can explain the low acceptance rate via the extensive margin of trade that arises from trade between buyers and sellers with heterogeneous beliefs. Furthermore, it was demonstrated that as buyer's perceived value of money and/or belief uncertainty is higher than seller's then (1.) the quantity traded decreases (intensive margin) and (2.) the seller may reject some offers (extensive margin). We also showed that when beliefs are heterogeneous across buyers, then there is a non-degenerate distribution of money holdings at the end of each centralized market meeting. Importantly, it is along a non-stationary dynamic equilibrium path in which the model is consistent with the experimental findings. The results in Section 3 characterize how heterogeneous beliefs can destabilize the perfect foresight steady-state and give rise to non-stationary paths.

Monetary theory has proposed a variety of extensions to the benchmark Lagos-Wright model that can explain many of these findings. The model with heterogeneous beliefs, however, makes distinct theoretical predictions that can be verified in experimental data. In particular, because the distribution of heterogeneity is an endogenous and time-varying object this implies that the partial acceptability of money is also endogenous and time-varying. The nature of heterogeneous beliefs depends on the equilibrium CM price process, among other structural features of the model, and so the model, moreover, predicts a very particular non-linear relationship between the acceptance rate and the CM money price. This relationship is plotted in Figure 14.

The top plot scatters acceptance rates and CM prices along a particular (strange) attractor when the model is parameterized with $\gamma = 1.9, \Sigma = 1.2$, which has a steady-state price of approximately 8. The middle plot scatters acceptance rates and CM prices across simulations when $\gamma = 0.75$ and $0 \leq \Sigma \leq 1.2$, while the bottom plot fixes $\gamma = 1.3$ and varies $0 \leq \Sigma \leq 0.75$, these cases have steady-state CM prices of approximately 5. The top plots provides the relationship between acceptance and CM price along a particular equilibrium

Figure 14: Acceptance Rates and Prices in Model.



path. However, this is a general feature of the model as the bottom two plots illustrate the same relationship across equilibria.

There are two striking features in Figure 14.

1. When the CM price ϕ is above its steady-state value then the acceptance rate is always equal to 1.
2. When the CM price ϕ is below its steady-state value then there is a non-linear relationship between the acceptance rate and ϕ .

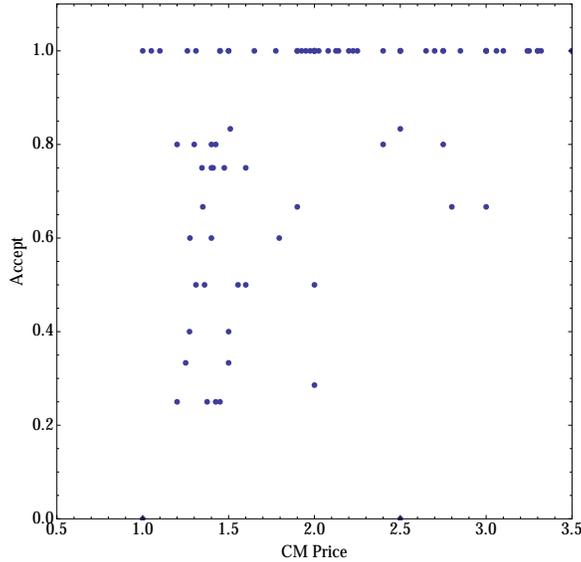
The adaptive forecasting model always forecasts price to revert back to its steady-state value, thus when ϕ is above steady-state seller beliefs about the value of money will be greater than buyers, regardless of Σ : the sellers accept buyers' offers. The further the price is from steady-state, the higher the acceptance rate. This second result is intuitive as price is further away from steady-state the fraction of buyers that are rational is higher, there is less heterogeneity, and so the acceptance rate is higher. However, close to, but below the steady-state, the fraction of adaptive learning agents is higher as the benefits to the perfect foresight predictor are outweighed by the costs, and now the sellers' beliefs about the value of money are lower than buyers. Thus, we find lower acceptance rates.

We now check for evidence of this theoretical channel in the Duffy and Puzzello (2013) experimental data.³⁰ Figure 15 plots the relationship in acceptance rates and prices in the Duffy and Puzzello (2013) experimental data. To generate this figure we took the data from all of the experimental treatments in Duffy and Puzzello (2013) and calculated the acceptance rate in each period of the experiments. In their parameterization, the steady-state CM price is 2, the number of periods in the each experiment varied from 29 to 48. Their framework is exactly the benchmark Lagos-Wright model with buyer take-it-or-leave-it offers. In Figure 15, there is the same general relationship between acceptance rates and the CM price of money. Except for a few outliers, when price is above steady-state the acceptance rate is 1. In periods where the price is below 2, there is a negative relationship between the price and the acceptance rate, with very low prices having an acceptance rate of 1. Thus, the experimental data has the same general features of a non-monotonic relationship between the acceptance rates and prices as predicted by the model with heterogeneous beliefs.

The results in this paper are also complementary to recent experimental findings reported in Berentsen, McBride, and Rocheteau (2014). Berentsen, McBride, and Rocheteau (2014) conduct related experiments in the laboratory but where there is uncertainty, and possibly private information, about the CM value of the assets. They find that uncertainty greatly decreases the extensive of margin, in take-it-or-leave-it offers made by buyers, when there are information frictions and no search frictions. In their paper, the unexpectedly low acceptance rates are attributed to fairness considerations in bargaining. The results in this paper are complementary to their findings, especially the results when experimental subjects are symmetrically uninformed about the CM value of the asset. However, we offer a different interpretation. In the present paper, agents hold different beliefs, arising from different forecasting models, about the value of the asset and the distribution of agents across predictors

³⁰We are grateful to John Duffy for sharing the data that facilitated this analysis.

Figure 15: Acceptance Rates and Prices in Experimental Data. Source: Duffy-Puzzello



is endogenous. We find, like Berentsen, McBride, and Rocheteau (2014) that informational frictions can lead to trading inefficiencies, along both the intensive and extensive margins, even when the search friction is shut down. Unlike fairness explanations, we also document a non-linear relationship between the acceptance rate and the CM price, with the distance from the steady-state being an important predictor of acceptance rates.

Another distinguishing implication of our framework is that sellers who accept Bayesian offers will realize a surplus that results from the cautious offers induced by a buyer’s uncertainty about the seller’s acceptance rule. The model predicts a non-linear relationship between that surplus and the distance the economy is from steady-state. In the model with heterogeneous beliefs, the seller will receive a surplus from the mistaken beliefs by adaptive buyers about the seller’s incentive compatibility constraint. When CM prices are higher, the buyer underestimates the seller’s value of money and, as a result, sellers realize a surplus. Figure 16 plots the realized seller surplus across simulations, where γ is fixed (at 0.75 on the left, and 1.3 on the right) and uncertainty Σ varies. This figure demonstrates a strong positive relationship between realized seller surplus when CM prices are above steady-state.

Figure 17 confirms this implication in the experimental data. To generate this figure, we compute realized seller surplus across all experimental treatments. It is evident from this figure that for CM prices below steady-state there is not a strong relationship between seller surplus and CM price. However, for prices above steady-state, there is a strong positive relationship. One would not expect to see this non-linear relationship between seller surplus and CM prices, and the non-monotonic relationship between acceptance rates and seller surplus, in alternative explanations for the anomalous experimental findings.

Figure 16: Seller Surplus and Prices in Model.

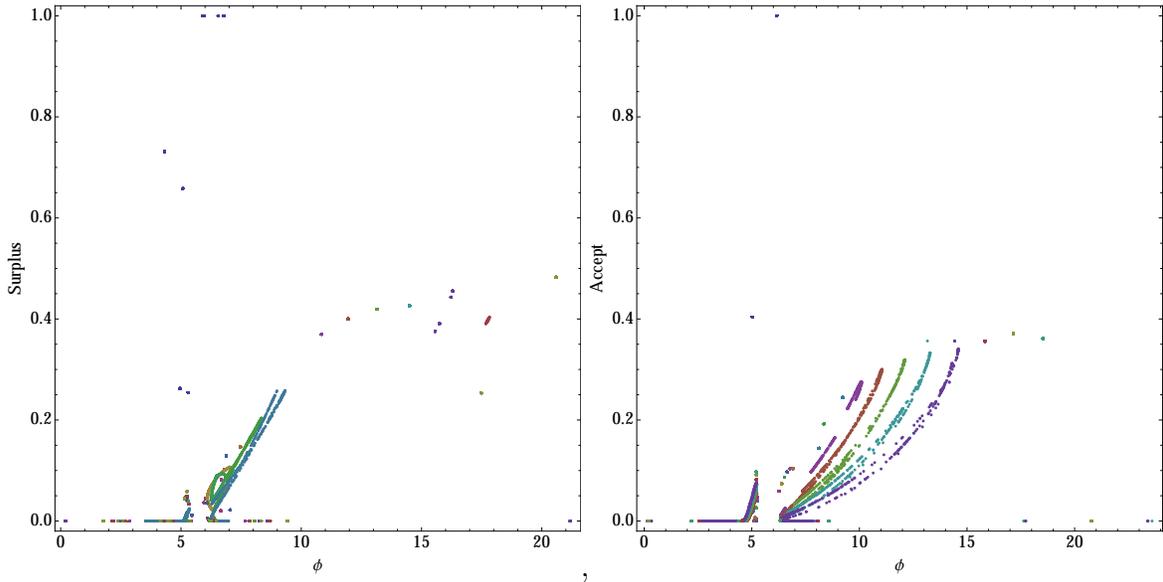
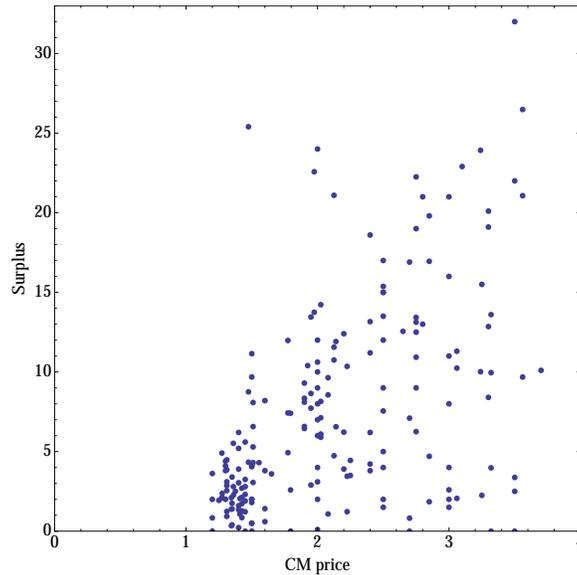


Figure 17: Seller Surplus and Prices in Experimental Data. Source: Duffy-Puzzello



6 Conclusion

Monetary economics seeks theories that adhere to the “Wallace dictum” that the use of fiat money should arise as an equilibrium phenomenon from the physical environment, preferences, technologies, and other fundamental factors. Monetary search theory gives a role to money because of search and bargaining frictions in bilateral, or over-the-counter, markets where limited commitment or incomplete record keeping preclude credit arrangements between buyers and sellers. Of course, search models for money feature multiple stationary

equilibria and so which equilibrium is actually realized depends on agents' beliefs. This paper focuses on beliefs as another fundamental in monetary search models and introduces heterogeneous beliefs by relaxing the (strong) requirement that all buyers in decentralized markets must hold rational expectations. We introduce heterogeneous beliefs by assuming that there are a fraction of agents who hold rational expectations (and must pay a computational cost for formulating RE) and the remaining agents formulate expectations from an adaptive learning rule inspired by Evans and Honkapohja (2001). The distribution of agents across these models is an endogenous object of the model so that beliefs and prices are determined simultaneously in an equilibrium.

The primary contribution of this paper is to demonstrate how endogenously determined heterogeneous beliefs affect the existence, and dynamic properties, of monetary equilibria. These dynamics yield distinct theoretical predictions that are consistent with the results from Lagos-Wright laboratory experiments by Duffy and Puzzello (2013). The primary results are as follows. Heterogeneous beliefs can alter the number and nature of monetary equilibria by destabilizing the steady state and leading to a variety of periodic and aperiodic cycles, including complicated dynamics, that remain bounded around the monetary equilibrium. Three key implications follow from the non-linear dynamics induced by heterogeneous beliefs. First, welfare is reduced, relative to the rational expectations equilibrium; when buyers and sellers have heterogeneous beliefs, the level of trade within a match may be too low or too high relative to the rational expectations outcome. The reduced-welfare is partly the result of an endogenous extensive margin of trade, a partial acceptability of money in some matches that magnifies search frictions. The partial acceptability arises along a dynamic path when a buyer and seller with different expectations about the value of money are matched, they may be unable to reach agreement on the terms of trade. Second, buyers who are uncertain about their beliefs behave like a Bayesian by placing a prior on a sellers' participation constraint and, as a result, tend to make more cautious offers. We identify three forces through which buyer uncertainty impacts welfare: (1.) uncertainty induces cautious offers whereby the buyer offers a higher payment for a lower quantity; (2.) all else equal, uncertainty increases the acceptance rate of offers made by buyers; (3.) small changes in uncertainty can bifurcate the equilibrium and decrease the mean acceptance rate. Third, the model with heterogeneous beliefs make unique testable implications about acceptance rates and gains from trade that are consistent with experimental evidence.

What is the intuition for why heterogeneous beliefs can destabilize the stationary monetary equilibrium? In a stationary equilibrium, perfect foresight and the adaptive predictor deliver the same forecast. Depending on the parameterization, the monetary equilibrium may be determinate or indeterminate. When it is indeterminate, then the monetary equilibrium is a sink under perfect foresight dynamics and a source under adaptive expectations. In a neighborhood of the stationary equilibrium, perfect foresight and the adaptive predictor will both forecast similarly and agents will be unwilling to pay the computational costs necessary to form rational expectations. As the fraction of adaptive agents increases, the dynamics push the economy away from the stationary equilibrium until it reaches a point where the accuracy gains to rational expectations outweigh its costs, the fraction of rational agents increases and the economy moves towards the stationary equilibrium. The resulting tension between the stabilizing (or, attracting) and repelling dynamics can yield periodic

orbits and complex dynamics.

Since the heterogeneous belief friction can introduce new inefficiencies into the economy, a natural question is whether a monetary policy can be designed to implement the efficient equilibrium for any specification of heterogeneous beliefs. In an extension to this paper, we consider such a policy implemented via an interest on currency policy proposed by Friedman (1959), and studied in Sargent and Wallace (1985) and Andolfatto (2010). In this policy, the policy maker pays interest on money holdings, which can alternatively be interpreted as interest on reserves. We consider two cases: first, a Friedman rule where an interest peg is designed to implement the efficient outcome in a stationary equilibrium; second, an interest rate feedback rule, similar in spirit to the Taylor rule that is a staple of monetary policy analysis in New Keynesian models, that adjusts interest paid whenever inflation deviates from its efficient stationary value. We show that the Friedman rule is not necessarily the optimal rule as implementing the efficient outcome can be hindered by destabilizing heterogeneous beliefs. A Taylor-type rule, however, can implement the efficient equilibrium, by stabilizing the steady state, and overcome the heterogeneous belief friction provided that policy reacts strongly when inflation is away from its efficient stationary value. While a detailed analysis is beyond the scope of the present study, it is worth briefly mentioning that the size of the reaction coefficient in the Taylor rule that can implement stability of the efficient steady-state depends on the details of the heterogeneous beliefs and, interestingly, on the uncertainty of buyers about sellers' beliefs. A reaction coefficient greater than one – i.e. a policy rule that satisfies the Taylor principle – is not a necessary condition for stability, a result that differs from the standard policy advice in New Keynesian models. However, greater degrees of uncertainty in buyers' Bayesian offers imply that the policy rule needs to react stronger to inflation innovations to stabilize the efficient steady-state. The economic justification for a Taylor rule is different as well since the result does not depend on the presence of nominal rigidities. These interesting policy implications are explored in greater detail in on going research.

Appendix

Proof of Proposition 1. The first order condition corresponding to the problem (5) is given by

$$LHS(q) \equiv m \frac{u'(q)}{c'(q)} = (u(q) - \phi^b m) h_\varepsilon \left(\frac{c(q)}{m} \right) \equiv RHS(q). \quad (18)$$

Let q_{\min} satisfy $u(q_{\min}) = \phi^b m$. Since $q < q_{\min}$ implies $RHS(q) < 0 < LHS(q)$, there can be no solution to (5) in the interval $(0, q_{\min})$. Next, notice $RHS(q_{\min}) = 0 < LHS(q_{\min})$ and $RHS(q) \rightarrow \infty$ as $q \rightarrow \infty$. Also, $q > q_{\min}$ implies $LHS'(q) < 0$ and $RHS'(q) > 0$. We conclude that there is a unique solution to (18).

To show that this solution is a maximum, it suffices to show that

$$\lim_{q \rightarrow \infty} (u(q) - \phi^b m) \left(1 - F_\varepsilon \left(\frac{c(q)}{m} \right) \right) = 0,$$

which, by concavity of u , itself is implied by showing that $\lim_{q \rightarrow \infty} q(1 - F_\varepsilon(q)) = 0$. By l'Hopital's rule, it suffices to show that $\lim_{q \rightarrow \infty} q^2 dF_\varepsilon(q) = 0$. Suppose to the contrary that this limit is equal to $M > 0$. Then there is some \hat{q} so that $q > \hat{q}$ implies $q^2 dF_\varepsilon(q) > \frac{M}{2}$. Thus

$$\int_{-\infty}^{\infty} q^2 dF_\varepsilon(q) \geq \int_{\hat{q}}^{\infty} q^2 dF_\varepsilon(q) \geq \int_{\hat{q}}^{\infty} \frac{M}{2} dq = \infty.$$

But this contradicts the existence of finite second moments. ■

Proof of Corollary 4.

For $\alpha > \alpha_c$, the steady state is dynamically unstable, and near α_c , the behavior of price paths may be determined by analyzing the nature of the associated bifurcation. Because f is sufficiently smooth at $\bar{\phi}$, and since $f_\phi(\bar{\phi}, \alpha_c) = -1$, a flip bifurcation obtains provided that

$$\begin{aligned} 0 &\neq 1/2 f_{\phi\phi}(\bar{\phi}, \alpha_c)^2 + 1/3 f_{\phi\phi\phi}(\bar{\phi}, \alpha_c) \\ 0 &\neq f_{\phi\alpha}(\bar{\phi}, \alpha_c). \end{aligned}$$

Straightforward computation yields $1/2 f_{\phi\phi}(\bar{\phi}, \alpha_c)^2 + 1/3 f_{\phi\phi\phi}(\bar{\phi}, \alpha_c)$

$$\begin{aligned} &= -\frac{1}{6} \beta (\theta^2 + 1) M^{2 - \frac{2}{\theta^2}} \\ &\times \left[2(2\theta^2 + 1) M^{\frac{2}{\theta^2}} \left(\theta^2 \beta^{\frac{1}{\alpha}} - (\theta^2 - 1) \beta^{\frac{\theta^2}{\theta^2+1}} \right)^{-\frac{1}{\theta^2} - 3} - 3(\theta^2 + 1) \beta^{\frac{4}{\theta^2+1} - 3} \left(\frac{\beta^{\frac{\theta^2}{\theta^2+1}}}{M} \right)^{-\frac{2}{\theta^2}} \right], \end{aligned}$$

and $f_{\phi\alpha}(\bar{\phi}, \alpha_c)$

$$\begin{aligned} &= \frac{\beta \left(\theta^2 \beta^{\frac{1}{\alpha}} - (\theta^2 - 1) \beta^{\frac{\theta^2}{\theta^2+1}} \right)^{-\frac{1}{\theta^2} - 2}}{\theta^2 + 1} \\ &\times \left[(\theta^2 + 1) \left((\theta^2 - 1) \beta^{\frac{\theta^2}{\theta^2+1}} - \theta^2 \beta^{\frac{1}{\alpha}} \right) \left(\theta^2 - \log \left(\theta^2 \beta^{\frac{1}{\alpha}} - (\theta^2 - 1) \beta^{\frac{\theta^2}{\theta^2+1}} \right) \right) \right. \\ &\left. + \theta^2 (\theta^2 - 1) \beta^{\frac{\theta^2}{\theta^2+1}} \log(\beta) \right]. \end{aligned}$$

For $\theta = 1$, we find that $1/2f_{\phi\phi}(\bar{\phi}, \alpha_c)^2 + 1/3f_{\phi\phi\phi}(\bar{\phi}, \alpha_c) = 0$; and again, analogous to the perfect foresight case, numerical analysis indicates that all non-steady state adaptive price paths are explosive for $\alpha > \alpha_c$; and otherwise, the conditions for a flip bifurcation are generically satisfied, and so as α crosses α_c from below, the steady state destabilizes and the a two-cycle emerges. \square

Proof of Proposition 5.

Writing $\Phi_t = (\phi_t, \phi_{t-1}, \phi_{t-2})$ and defining h in the obvious way, we may rewrite the dynamic system as $\Phi_t = h(\Phi_{t-1})$. The eigenvalues of Dh evaluated at the steady state are given by 0 and

$$\frac{-e^{\gamma C} - 1 \pm \sqrt{e^{\gamma C} (-4(\alpha - 1)^2\theta + e^{\gamma C} + 2) + 1}}{2(1 - \alpha)}.$$

While not particularly complicated, precise analytic results can only be obtained for particular restrictions on θ .

Since we require $\gamma > 0$, it follows that for stability, $\alpha > 3$.

Note that $h_{\phi}(\bar{\phi}, \gamma_c) = -1$; further, we may compute

$$\begin{aligned} h_{\phi\gamma}(\bar{\phi}, \gamma_c) &= \left(\frac{1}{\alpha - 1} - 1 \right) c \\ 1/2h_{\phi\phi}(\bar{\phi}, \gamma_c)^2 + 1/3h_{\phi\phi\phi}(\bar{\phi}, \gamma_c) &= \frac{1}{6}(\alpha - 2) \left(\frac{12 \log(\alpha - 2)}{(\alpha - 1)c} - \frac{M^2\beta^{-2/\alpha}}{\alpha} \right), \end{aligned}$$

revealing that as γ crosses γ_c from below, a flip bifurcation obtains, and locally, as the steady state destabilizes, a stable two cycle emerges. \square

Nash Bargaining

The quasi-linearity of the W (for both the buyer and seller) allows for a particularly simple formulation of the generalized Nash solution:³¹

$$(q(m, \phi^b, \phi^s), d(m, \phi^b, \phi^s)) = \arg \max_{q, d \leq m} (u(q) - \phi^b d)^\lambda (-c(q) + \phi^s d)^{1-\lambda}.$$

The parameter λ captures the bargaining power of the buyer. When there is an interior solution (i.e. the $d \leq m$ does not bind), then the solution solves

$$\frac{u'(q)}{c'(q)} = \frac{\phi^b}{\phi^s} \tag{19}$$

$$d = \frac{\lambda c(q)\phi^b + (1 - \lambda)u(q)\phi^s}{\phi^b\phi^s}. \tag{20}$$

When the constraint $d \leq m$ binds, then the solution is $d = m$ and q solves

$$m = \frac{\lambda u'(q)c(q) + (1 - \lambda)c'(q)u(q)}{\lambda \phi^s u'(q) + (1 - \lambda)\phi^b c'(q)}. \tag{21}$$

³¹As in Gu, Mattesini, Monnet, and Wright (2012) we take the Generalized Nash solution as a behavioral primitive and abstract from specifying a formal strategic bargaining game.

Equations (19)–(21) characterize the bargaining outcome (q, d) given money stock m and beliefs (ϕ^b, ϕ^s) .

Having specified Nash bargaining with heterogeneous beliefs, we now direct attention to trades in one given period, and so we again omit time subscripts. Also, in this portion only, we allow for the possibility that the seller and/or the buyer may hold more money than is needed for/gained from DM trades.

The unique, pareto-efficient quantity of trade, q^* , satisfies $u'(q^*) = c'(q^*)$. To see this, assume $u'(q) > c'(q)$. Then a planner could require the corresponding seller to increase his specialized-goods production by a small quantity dq and similarly require the buyer to increase his general-goods production by a small quantity dx . If the planner transfers dq to the buyer and dx to the seller then the utility differentials are given by

$$dU^s = dx - c'(q)dq \text{ and } dU^b = u'(q)dq - dx.$$

Provided $c'(q)dq < dx < u'(q)dq$, both of these differentials are positive, and such choice for dx is possible since $u'(q) > c'(q)$. A similar argument holds in case $u'(q) < c'(q)$.

Here we assume that the terms of DM trades are arranged through Nash bargaining, and thus characterized by (19)–(21). If the liquidity constraint binds, that is, if $d = m$, then the quantity traded is generically inefficient, even if $\phi^b = \phi^s$: this follows from (21) and the fact that the equilibrium quantity of money is determined exogenously. If buyers hold excess money balances after the trade, that is, if $d < m$, then, by (19), efficiency obtains precisely when $\phi^b = \phi^s$. In this case, heterogeneous expectations implies an inefficient quantity traded in the DM. We summarize these observations in the following proposition:

Proposition 6 *If the Nash bargaining solution results in efficient trade then expectations are homogeneous.*

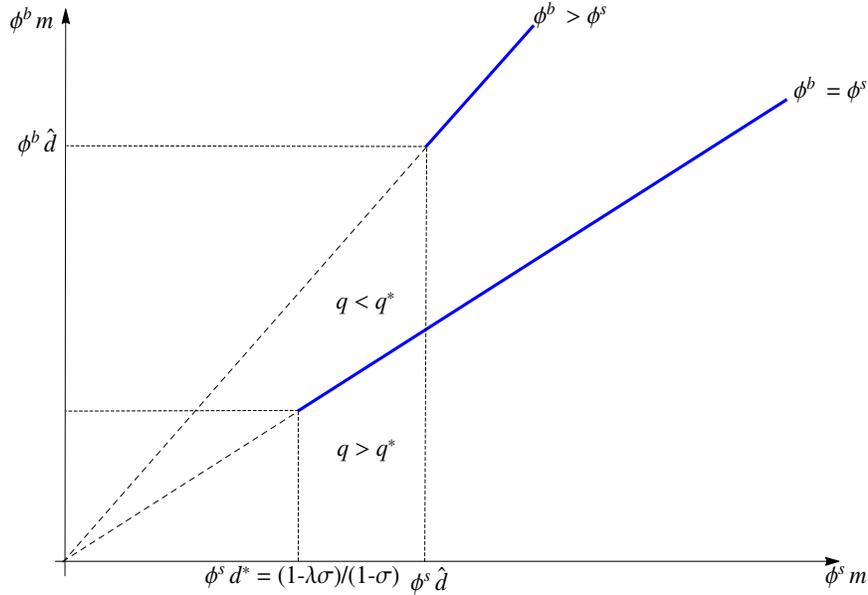
The bargaining solution evidently depends on the heterogeneity in beliefs between buyers and sellers. Proposition 6 shows that, even in the unconstrained case, the efficient level of trade will only arise when buyers and sellers have identical beliefs about the value of money. Notice that when $\phi^b/\phi^s > 1$ then $q < q^*$ and when $\phi^b/\phi^s < 1$ then $q > q^*$. Thus, even without any liquidity premium, i.e. scarcity of money, Nash bargaining with heterogeneous beliefs can lead to trade below or above the inefficient level of trade.

Figure 18 illustrates the possible outcomes from the Nash bargaining solution. The axes measure the perceived value of money, in terms of the numeraire good, to sellers and buyers, respectively.

Consider first the 45-degree line, consisting of all of the possible values for which buyers and sellers have identical beliefs. The point ϕd^* is the efficient outcome that arises from the Nash bargaining solution with homogeneous beliefs. The dashed part of the line, then, represents values of m where buyers are liquidity constrained. The solid line is for the situation where buyers have more money on hand than needed to purchase the efficient level of trade and so $q = q^*$ along this segment.

Figure 18 can be used to identify regions of inefficient trade when buyers are not liquidity constrained. Suppose that $\phi^b > \phi^s$, and consider the line with slope ϕ^b/ϕ^s . Associated with

Figure 18: Bargaining and Heterogeneous Beliefs



ϕ^b, ϕ^s is a quantity \hat{q} satisfying equation (19) where $\hat{q} < q^*$. Also, there is an unconstrained value of $d = \hat{d}$ determined by \hat{q} : it satisfies (20), and thus depends on ϕ^b, ϕ^s as well as the curvature of the utility and cost functions and the relative bargaining strength λ . Figure 18 labels the corresponding point $(\phi^s \hat{d}, \phi^b \hat{d})$. Figure 1 illustrates the dependence of the Nash bargaining solution on the value of m . The dashed segment of the line correspond to outputs for which buyers are liquidity constrained: $d = m$ and $q < \hat{q}$; along the solid line, have $q = \hat{q}$ and $d < m$.³²

These results illustrate how heterogeneous beliefs can lead to trading inefficiencies. Under homogeneous, rational expectations, a trading inefficiency arises whenever buyers do not hold enough money to facilitate the efficient level of consumption in the DM. The results in this section show that whenever beliefs differ between buyers and sellers, the level of trade will be inefficient. This occurs even when buyers are not constrained by their money holdings. If buyers place a higher value on money than sellers, in terms of the numeraire good, then bargaining will lead to an outcome with less trade as sellers will not want to produce as much as buyers want. Similarly, when sellers value money more than buyers they will produce a greater quantity than buyers will be willing to consume.

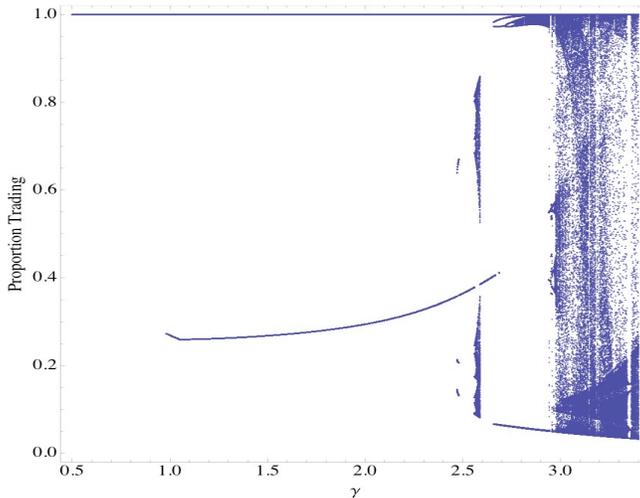
Welfare properties with dogmatic beliefs

A non-trivial extensive margin of trade, or partial acceptability of money in trade, is a

³²We note that if $\phi^b/\phi^s > 1$ then $\hat{q} < q^*$, and if $\phi^b < \phi^s$ then $\hat{q} > q^*$. In particular, the non-constrained outcome may be above or below the efficient level of trade.

robust feature of non-stationary equilibria under the take-it-or-leave-it protocol when buyers have dogmatic beliefs, that is, when their subjective priors over sellers' beliefs is a unit mass on the buyer's own beliefs. Since all agents hold homogeneous beliefs in a steady state, in a stationary monetary equilibrium there will be trade, possibly at an inefficient level, as the buyer makes an offer that the seller is certain to accept. However, in periodic and other complicated equilibria, there may be bilateral meetings where trade does not occur. For example, take the case where buyers are distributed between perfect foresight, with proportion n_t , and the adaptive learning rule with a fraction $1 - n_t$. Now consider a two-cycle ("near" the bifurcation point) where the value of money alternates between $\phi_h > \phi_l$. Buyers with adaptive beliefs will always expect $\phi_{t+1}^e < \phi_h = \phi_{t+1}$ when $\phi_t = \phi_l$ and $\phi_{t+1}^e > \phi_l = \phi_{t+1}$ when $\phi_t = \phi_h$. A seller, who is assumed to have perfect foresight, accepts an offer whenever $\phi_{t+1} \geq \phi_{t+1}^e$. The participation constraint of sellers matched with "adaptive" buyers will be violated whenever $\phi_t = \phi_h$. Thus, every other period $(1 - n_t)$ of the sellers will reject the buyers' offers and there is an extensive margin of trade. There is a partial acceptability that magnifies search frictions in the sense that every other period buyers and sellers with different beliefs do not trade.

Figure 19: Acceptability Bifurcation Diagram: $\theta = 0.1$

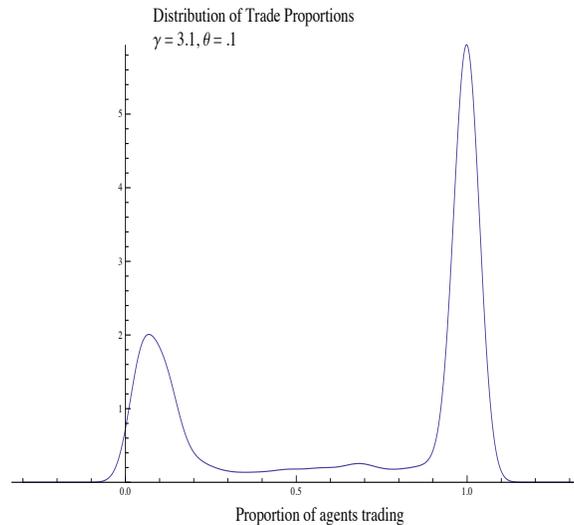


We can illustrate this insight more generally. Figure 19 plots the bifurcation diagram for the acceptance rate (i.e. the fraction of buyers and sellers that trade in the decentralized market) for the same parameterization in Figure 6. For each value of γ the model is simulated, after a long transient period, and in each period the proportion of matches that result in trade are recorded. The realization of all of these acceptance rates along the length of the attractor are plotted against the value of γ . Recall from Figure 6 that for small values of γ there is a stable steady state with all agents trading in every period. Around $\gamma \approx 1$ the steady state bifurcates into a two-cycle. Figure 19 confirms the intuition above that, along

the two cycle, in every other period a fraction of the buyers (those with “adaptive” beliefs) do not trade. As γ continues to increase the two-cycle bifurcates again, and along with it the acceptance rate varies over time as well. For larger values of γ there are complicated dynamics and the pattern of trade becomes complicated along these stable attractors.

Figure 20 provides further insight by plotting results when $\gamma = 3.1$ which is one particular slice of the bifurcation diagram in Figure 19. The bifurcation figure makes it clear that the proportion of buyers that actually trade with a seller depends on the nature of the equilibrium dynamics. In general, if the economy is on a stable attractor, the proportion of agents trading will vary with time, and it is natural, in this way, to interpret the attractor as identifying an asymptotic stationary distribution. Figure 20 plots the smoothed kernel estimates of this distribution using 1,000,000 data points (following a long transient period). This result shows that heterogeneity leads to a non-trivial monetary distribution and an acceptance rate that is endogenously time-varying even when the search friction is shut down ($\sigma = 1$). Thus, this figure provides a clear illustration of the heterogeneous belief friction when buyers make take-it-or-leave-it offers with dogmatic subjective beliefs.

Figure 20: Acceptance Rate Asymptotic Stationary Distribution: $\theta = 0.1$

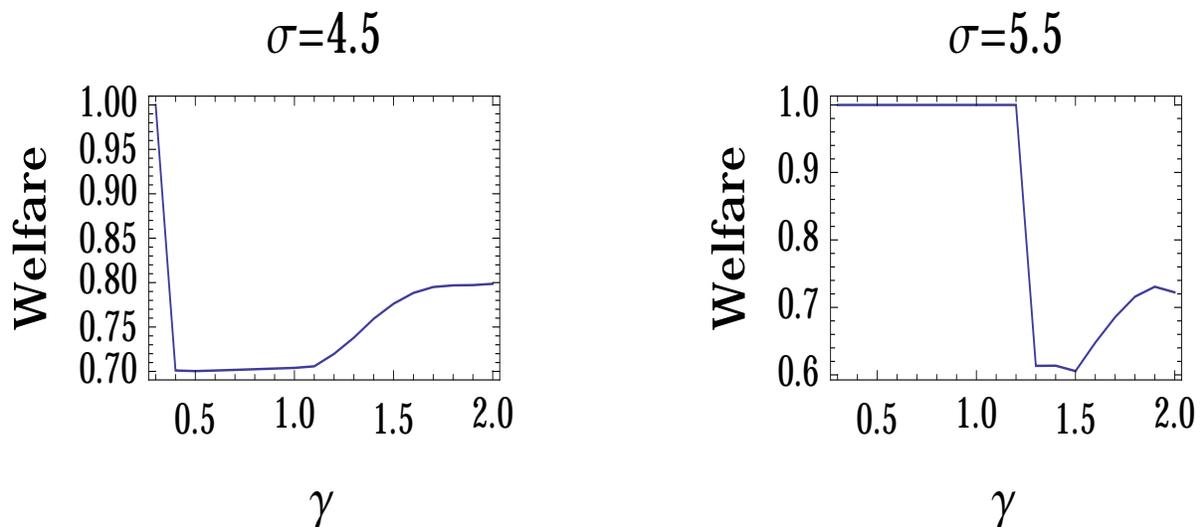


Heterogeneous beliefs alters the precise details of trade in over-the-counter markets, i.e. the intensive and extensive margins. Whether an extensive margin of trade will emerge from a setting with heterogeneous beliefs depends, in a subtle way, on the particular bargaining protocol and the assumption about common knowledge of beliefs within a match. The results above show that with buyer-takes-all offers, and beliefs that are not common knowledge within a match, that it is possible along a dynamic equilibrium path for some buyers’ offers to be rejected. Under Nash bargaining, though, trade will always occur so long as there is a positive surplus to split between buyer and seller. Our formulation of Nash bargaining has common knowledge of beliefs within a match, so that one can extend this logic to the case of buyer-takes-all and common knowledge of beliefs. Thus, it is clear that whether an extensive margin will arise or not depends a lot about the details of the bargaining process

and whether beliefs are common knowledge. The robustness and generality of these results to alternative bargaining/information assumptions is a topic left for future research.

As background for the results in the main text, we also include here an examination of welfare in the case where buyers have dogmatic beliefs about seller's participation constraint. In the case where buyers have dogmatic beliefs about the seller's participation constraint there is, of course, no uncertainty so Figure 21 plots welfare across a range of monetary equilibria. In the plot, welfare is measured as the unconditional average surplus in the DM $u(q) - c(q)$.³³ The horizontal axis is the intensity of choice parameter γ , which is the bifurcation parameter of interest in the dogmatic priors case. The vertical axis measures welfare relative to the steady-state monetary equilibrium under rational expectations. The inefficiency of heterogeneous beliefs is measured relative to the stationary monetary equilibrium which already features an inefficiently low level of DM output (provided the liquidity constraint binds). Thus, Figure 21 measures the inefficiency that results from heterogeneous beliefs with dogmatic priors.

Figure 21: Welfare: $\theta = 0.2$



While in principle welfare could be higher as a result of the dynamic fluctuations induced by rationally heterogeneous expectations. However, Figure 21 demonstrates that is not the case. The fluctuations in price above its steady-state value leads to more DM consumption which has a positive effect on welfare. Any positive effects, though, are outweighed by the partial acceptability in matches between buyers and sellers with heterogeneous beliefs, in particular when buyers expect the value of money to be higher than sellers. This effect is most clearly seen in the right plot of Figure 21, which plots welfare against the intensity of choice when there is a high degree of curvature in DM utility ($\sigma = 5.5$). For low values of γ , the steady-state is stable and so the welfare of the model with heterogeneous beliefs is the same as the rational expectations model. For sufficiently high γ , the steady-state becomes

³³This welfare measure is equivalent to the discounted sum of DM surpluses when the initial condition is drawn with equal probability from any of the points along the cycle.

unstable and there is a stable two-cycle. Every other period, all adaptive buyers' offers are rejected and, as a consequence, there is a sharp drop in welfare. A similar pattern emerges in the left panel with a lower value for σ . Notice in the figure that as γ increases further welfare increases somewhat, though still substantially lower than in the rational expectations steady-state. This occurs as additional period-doubling bifurcations lead to cycles which feature some periods with higher acceptance rates which raises the average acceptance rate over time.

Learning about uncertainty

In Section 4, results are presented on the dynamics of Bayesian offers where buyer uncertainty is taken as a fixed parameter of buyers' priors about sellers' beliefs. Ideally, buyers would update their both the mean and variance of their subjective beliefs when formulating offers. The main text considers the case where the buyers learn about the first moment of the distribution as a benchmark case. This section presents a formulation where they also learn about the second moment. The dynamics with learning about first and second moments are qualitatively similar, though we do not focus on this case because the numerical analysis is computationally burdensome for all adaptive learning rules where $\theta \neq 1$.

In this extension to the basic model, the adaptive learning rule is given by the following pair of equations:

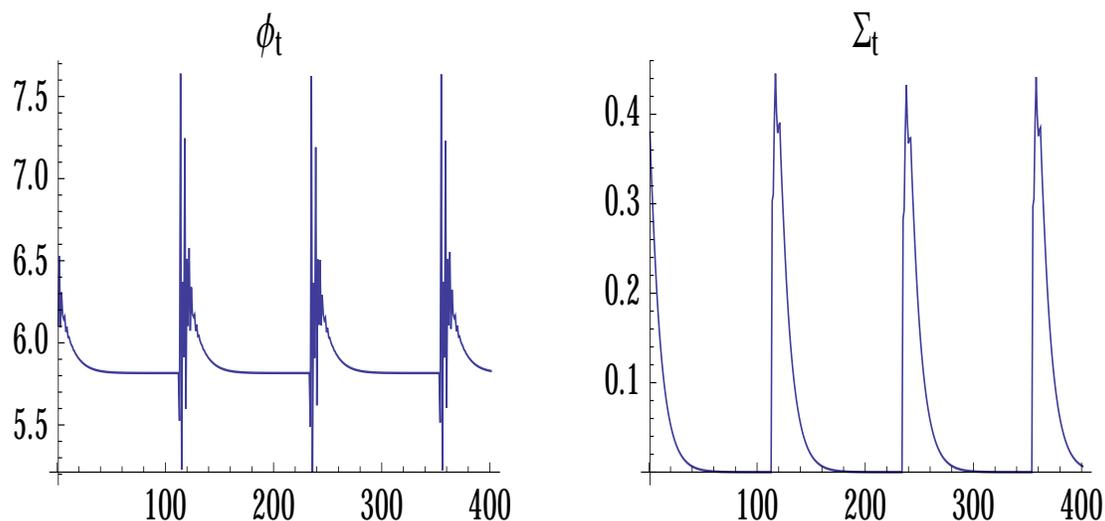
$$\begin{aligned}\phi_t^e &= \bar{\phi} + \theta (\phi_{t-1} - \bar{\phi}) \\ \Sigma_t &= \Sigma_{t-1} + \theta_v \left((1 - \theta)^2 (\phi_t - \bar{\phi})^2 - \Sigma_{t-1} \right)\end{aligned}$$

where Σ_t is a recursive estimate of the conditional variance for ϕ_t when the conditional mean is given by ϕ_t^e . In this formulation, the gain in the estimating equation for Σ_t , $\theta_v \in (0, 1)$, is allowed to differ from the gain for the conditional mean θ . These beliefs are substituted into the equations for Bayesian offers in Section 2.4.3. Figure 22 plots the price dynamics and path for Σ_t in one representative simulation. Heterogeneous beliefs generate large swings in prices. These dynamics are reminiscent of the learning dynamics in Branch and Evans (2011).

References

- ANDOLFATTO, D. (2010): "Essential Interest Bearing Money," *Journal of Economic Theory*, 145(4), 1319–602.
- ANUFRIEV, M., AND C. H. HOMMES (2012): "Evolutionary Selection of Individual Expectations and Aggregate Outcomes," *American Economic Journal: Microeconomics*, 4(4), 35–64.
- ASSENZA, T., P. HEEMEIJER, C. HOMMES, AND D. MASSARO (2011): "Individual Expectations and Aggregate Macro Behavior," working paper.
- BARANOWSKI, R. (2012): "Adaptive Learning and Monetary Exchange," working paper.

Figure 22: Learning about uncertainty



BERENTSEN, A., G. CAMERA, AND C. WALLER (2005): “The Distribution of Money Balances and the Nonneutrality of Money,” *International Economic Review*, 46, 465–487.

BERENTSEN, A., M. MCBRIDE, AND G. ROCHETEAU (2014): “Limelight on Dark Markets: An Experimental Study of Liquidity and Information,” Working paper.

BERENTSEN, A., AND G. ROCHETEAU (2003): “Money and Gains from Trade,” *International Economic Review*, 44(1), 263–297.

BRANCH, W. A. (2004): “The Theory of Rationally Heterogeneous Expectations: Evidence from Survey Data on Inflation Expectations,” *Economic Journal*, 114, 592–621.

BRANCH, W. A., AND G. W. EVANS (2006): “Intrinsic Heterogeneity in Expectation Formation,” *Journal of Economic Theory*, 127, 264–295.

——— (2011): “Learning about Risk and Return: A Simple Model of Bubbles and Crashes,” *American Economic Journal: Macroeconomics*, 3(3), 159–191.

BROCK, W. A., AND C. H. HOMMES (1997): “A Rational Route to Randomness,” *Econometrica*, 65, 1059–1095.

COGLEY, T., AND T. J. SARGENT (2008): “Anticipated Utility and Rational Expectations as Approximations of Bayesian Decision Making,” *International Economic Review*, 49.

DIAMOND, P. A. (1982): “Aggregate Demand Management in Search Equilibrium,” *Journal of Political Economy*, 90, 881–894.

DUFFY, J., AND D. PUZZELLO (2013): “Gift Exchange versus Monetary Exchange: Theory and Evidence,” *forthcoming in American Economic Review*.

- ENGINEER, M., AND S. SHI (1998): “Asymmetry, Imperfectly Transferable Utility, and the Role of Fiat Money in Improving Terms of Trade,” *Journal of Monetary Economics*, 41, 153–183.
- (2001): “Bargains, Barter and Money,” *Review of Economic Dynamics*, 4, 188–209.
- EVANS, G. W., AND S. HONKAPOHJA (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton, New Jersey.
- FRIEDMAN, M. (1959): *A Program for Monetary Stability*. Fordham University Press, New York.
- GOLOSOV, M., G. LORENZONI, AND A. TYSVINSKI (2014): “Decentralized Trading with Private Information,” *Econometrica*, 82(3), 1055–1091.
- GRANDMONT, J.-M. (1977): “Temporary General Equilibrium Theory,” *Econometrica*, 45(3), 535–572.
- GU, C., F. MATTESINI, C. MONNET, AND R. WRIGHT (2012): “Endogenous Credit Cycles,” *Journal of Political Economy*, forthcoming.
- HOMMES, C. H. (2013): *Behavioral Rationality and Heterogeneous Expectations in Complex Economic Systems*. Cambridge University Press.
- JACQUET, N. L., AND S. TAN (2011): “Money, Bargaining, and Risk Sharing,” *Journal of Money, Credit and Banking*, 43(7), 419–442.
- KIYOTAKI, N., AND R. WRIGHT (1989): “On Money as a Medium of Exchange,” *Journal of Political Economy*, 97, 927–954.
- KREPS, D. M. (1998): “Anticipated Utility and Dynamic Choice,” in *Frontiers of Research in Economic Theory*, ed. by D. Jacobs, E. Kalai, and M. Kamien, pp. 242–274. Cambridge University Press.
- LAGOS, R., AND G. ROCHETEAU (2005): “Inflation, Output and Welfare,” *International Economic Review*, 46, 495–522.
- LAGOS, R., AND R. WRIGHT (2003): “Dynamics, Cycles, and Sunspot Equilibria in ‘Genuinely Dynamic, Fundamentally Disaggregative’ Models of Money,” *Journal of Economic Theory*, 109, 156–171.
- (2005): “A Unified Framework for Monetary Theory and Policy Analysis,” *Journal of Political Economy*, 113, 463–484.
- LESTER, B., A. POSTLEWAITE, AND R. WRIGHT (2012): “Information, Liquidity, Asset Prices, and Monetary Policy,” *Review of Economic Studies*, forthcoming.
- LI, Y., G. ROCHETEAU, AND P.-O. WEILL (2012): “Liquidity and the Threat of Fraudulent Assets,” working paper.

- LJUNGQVIST, L., AND T. J. SARGENT (2004): *Recursive Macroeconomic Theory, Second edition*. MIT Press, Cambridge, MA.
- MARCET, A., AND T. J. SARGENT (1989): “Convergence of Least-Squares Learning Mechanisms in Self-Referential Linear Stochastic Models,” *Journal of Economic Theory*, 48, 337–368.
- MOLICO, M. (2006): “The Distribution of Money and Prices in Search Equilibrium,” *International Economic Review*, 47, 701–722.
- NOSAL, E., AND G. ROCHETEAU (2011): *Money, Payments, and Liquidity*. MIT Press.
- NOSAL, E., AND N. WALLACE (2007): “A Model of (the threat of) Counterfeiting,” *Journal Of Monetary Economics*, 54, 229–246.
- PETERSON, B., AND S. SHI (2004): “Money, price dispersion and welfare,” *Economic Theory*, 24, 907–932.
- ROCHETEAU, G. (2011): “Payments and Liquidity under Adverse Selection,” *Journal Of Monetary Economics*, 58, 191–205.
- ROCHETEAU, G., AND R. WRIGHT (2005): “Money in Search Equilibrium, in Competitive Equilibrium, and in Competitive Search Equilibrium,” *Econometrica*, 73, 175–202.
- SARGENT, T. J. (1999): *The Conquest of American Inflation*. Princeton University Press, Princeton NJ.
- SARGENT, T. J., AND N. WALLACE (1985): “Interest on Reserves,” *Journal Of Monetary Economics*, 15, 279–290.
- SHI, S. (1997a): “A Divisible Model of Fiat Money,” *Econometrica*, 65, 75–102.
- (1997b): “A Divisible Search Model of Fiat Money,” *Econometrica*, 65, 75–102.
- YOUNG, P. H. (2004): *Strategic Learning and its Limits*. Oxford University Press.