Dampening General Equilibrium: From Micro to Macro *

George-Marios Angeletos† **Chen Lian**‡

April 7, 2017

Abstract

We argue that standard practice exaggerates the potency of general-equilibrium (GE) mechanisms by combining rational expectations with common knowledge of the state of the economy. Conversely, we formalize the notion that GE adjustment is weak, or that it takes time, by modifying an elementary Walrasian economy in two alternatives ways. In one, we replace rational-expectations equilibrium with cognitive processes that mimic Tâtonnement dynamics, Cobweb dynamics, or Level-k Thinking. In the other, we maintain rational expectations but remove common knowledge of aggregate shocks. This permits us, not only to illustrate the broader plausibility of the sought-after notion, but also to elaborate on the sense in which our preferred approach—removing common knowledge—can be seen as a disciplined substitute to relaxations of rational expectations. We discuss possible applications, including how our results help reduce the gap between the macroeconomic effects of interest and the micro or local elasticities that a growing empirical literature estimates in the cross section.

Keywords: Partial Equilibrium, Tâtonnement, Cobweb, Level-k Thinking, Reflective Equilibrium, Incomplete Information, Coordination, Beauty Contests.

---

*This paper subsumes an earlier draft that contained the same message but used a different framework. For helpful comments, we thank Adam Guren, Martin Beraja, and seminar participants at MIT, Duke, the EFCE group at the 2016 NBER Summer Institute, and the 2016 ESEM in Edinburgh.

†MIT and NBER; angelet@mit.edu.

‡MIT; lianchen@mit.edu.

1 Introduction

General-equilibrium (GE) effects are key to understanding how economic outcomes respond to aggregate shocks or to policy changes. In some contexts, GE effects reinforce partial-equilibrium (PE) effects, acting as "macroeconomic multipliers" that amplify exogenous shocks or raise the effectiveness of certain policies. In other contexts, GE effects counterbalance PE effects, helping stabilize aggregate outcomes and curtailing policy effectiveness. Either way, GE effects limit the usefulness of otherwise attractive PE intuitions. For instance, although the notion that "a reduction in demand can trigger a recession" appears to be self-evident from a PE perspective, it is inconsistent with the neoclassical/RBC framework because of countervailing GE effects.

GE effects also limit the usefulness of a growing empirical literature that seeks to quantify the macroeconomic effects of aggregate shocks by exploiting the cross-sectional heterogeneity in the exposure to such shocks. For instance, consider Mian and Sufi (2014). This work offers compelling evidence that US regions that experienced steeper drops in consumer credit during the recent recession also experienced larger employment losses. This evidence is consistent with the notion of demand-driven recessions; it is also suggestive of the role that consumer deleveraging may have played at the aggregate level. There is, however, an important limitation. GE effects such as the adjustment of prices, wages, and income at the national level are absorbed by the time fixed effect of the considered regressions. As a result, there is a gap—of unknown magnitude and sign—between the "micro" or "local" elasticities that are estimated in this kind of work and the macroeconomic effects that are ultimately of interest.

In this paper, we argue that the standard practice exaggerates the potency of GE mechanisms by combining a strong solution concept with strong informational assumptions. We do so by exploring multiple formalizations of the notion that GE adjustment may be "weak" or "slow". In all but one of them, we replace Rational Expectations Equilibrium (REE) with certain kind of cognitive procedures, which can be thought of as forms of bounded rationality; these procedures mimic Tâtonnement or Cobweb dynamics, Level-k Thinking, or Reflective Equilibrium (Garcıa-Schmidt and Woodford, 2015). In the remaining formalization, which we marginally favor, we maintain the REE concept but remove common knowledge of aggregate shocks.

This multifaceted approach permits us to attain two sets of goals. On the applied front, we corroborate the notion of weak or slow GE adjustment, thus also helping reduce the aforementioned gap between micro and macro elasticities. On the methodological front, we build a certain bridge between three literatures: the growing macroeconomic literature on higher-order uncertainty;[1] an older tradition that sought to capture the *off-equilibrium* price-adjustment process in Walrasian economies; and a more recent literature that replaces rational expectations with Level-k Thinking, or certain variants of it. We are thus able to elaborate on why some of the considered alternatives are better suited for capturing the sought-after notion than others, as well as on the sense in which our preferred route—that of removing common knowledge—can be seen as a useful substitute to "bounded rationality".

Framework. We consider an elementary Walrasian economy, featuring decentralized and sequential trading. There are two periods, which can be thought as the "present" and the "future", and a large number of "marketplaces", which define the boundaries of market interactions. Every agent—be it a firm or a household—can trade in a single marketplace in each period, but can be randomly be relocated from one marketplace to another as time passes. These are stark assumptions, which help us capture two realistic features: that most trading is decentralized; and that agents face uncertainty about future trades because they participate in different markets, and interact with different sets of agents, as time passes.[2]

[1] Higher-order uncertainty refers to the uncertainty that agents may face about the beliefs of other agents. This is distinct from the (first-order) uncertainty that agents may face about the underlying payoff-relevant fundamentals. See Morris and Shin (2001) and Angeletos and Lian (2016b) for discussions of the importance of this difference for applications.

[2] In these respects, our modeling approach is reminiscent of those used, inter alia, in the literatures on decentralized trading and OTC markets (e.g., Lagos and Wright, 2005; Duffie, Garleanu, and Pedersen, 2005). Also, when mapping the theory to the data, the relocation of agents from one market to another does not have to be interpreted literally as migration across geographic regions.

Starting from a status quo, we let an exogenous shock hit the economy's fundamentals (the cross-sectional profile of preferences, technologies, and endowments) and ask how the economy's outcomes (the various quantities and prices) adjust to the shock. The answer to this question depends on two sets of assumptions, one that governs the PE effect *within* each marketplace, and another that governs the GE effects *across* marketplaces.[3]

The first set contains the following assumptions: agents know their *own* fundamentals (i.e., consumers know their own preferences and endowments, and firms know their own technologies), both before and after the shock; agents respond optimally to the change in their own fundamentals, as well as to any price adjustment that takes place in the markets they currently participate in; and markets clear at all times. These elementary assumptions help characterize the demand and supply schedules and the market-clearing outcomes in any given marketplace, taking as given the following: the conjectures the agents in this marketplace form about the concurrent outcomes in *other* marketplaces and, thereby, about the prices they *themselves* may face in future trades.

Modeling the adjustment in these conjectures is synonymous to modeling the extent to which agents internalize the GE spillovers across the different marketplaces—and this is precisely where the second set of assumptions comes into the picture. What exactly this set is depends on whether one follow standard practice or depart from it. Standard practice imposes, not only (a.1) that agents form rational expectations, but also (a.2) that they have common knowledge of one another's rationality, as well as (b) common knowledge of the underlying aggregate shock and/or of the realized outcomes in the entire economy. The first two assumptions (a.1 and a.2) are embedded in the REE concept; the third assumption (b) is separate.

In our analysis, we maintain the first set of assumptions, but relax the second set of assumptions. In so doing, we are able to vary the potency of the GE effects that operate at the level of the economy (or the "macro level"), while holding constant the PE effects that operate at the level of a marketplace (or the "micro level").

Frictionless Benchmark. Our benchmark scenario follows—and indeed epitomizes—the standard practice by imposing the REE concept together with common knowledge of the aggregate shock.[4] As it turns out, this benchmark also replicates the outcomes of an Arrow-Debreu variant that allows the agents to trade a complete set of date- and state-continent trades in a single centralized market that operates only once, at the beginning of time.

This property is useful for two reasons. First, it clarifies that our choice to let trading be sequential and decentralized does not introduce a friction *by itself*. Second, it explains the sense in which GE adjustment is "perfect" and "instantaneous" in our benchmark: it is *as if* the agents can condition their behavior on the outcomes of *other* marketplaces and can therefore internalize all the GE spillovers to a full extent and with no delay.

With this benchmark as a reference point, we proceed to show how some—but not necessarily all—of the considered modifications help capture the notion that GE adjustment is "imperfect" or "slow".

Tâtonnement. As already noted, the degree by which the agents internalize the relevant GE effects are tied to the degree by which they adjust their conjectures about the outcomes of other marketplaces. Our first modification allows these conjectures to be determined according to an algorithm that mimics Tâtonnement dynamics.

Unlike the textbook version of Tâtonnement dynamics, there is no Walrasian auctioneer and no actual dynamics. Instead, the algorithm is instantaneous and takes place inside the agents' minds; that is, it is only a cognitive process. Nevertheless, this process mimics a Walrasian auctioneer who bases her initial price conjecture on the pre-shock equilibrium point; computes the implied gap between demand and supply; adjusts the price conjecture accordingly; and iterates. The number of iterations is interpreted as the "depth" of the cognitive process.

[3]By the PE effect of an aggregate shock we henceforth refer to the adjustment that takes place in the market-clearing outcomes of the typical marketplace, under the hypothesis that the outcomes in all other marketplaces stay constant. Because each marketplace is infinitesimal vis-a-vis the entire economy, the PE effect of an aggregate shock coincides with the overall equilibrium response of the typical marketplace to an idiosyncratic shock that hits only that particular marketplace.

[4]In our benchmark, Rational-Expectations Equilibrium (REE) coincides with Perfect Foresight Equilibrium (PFE), because of the assumption that all uncertainty gets realized in the first period. As explained in due course, this is only an innocuous simplification.

In our framework, this algorithm is equivalent to iterating on a contraction mapping, whose fixed point pins down the REE price conjectures. It follows that, in the limit as the depth of the process becomes infinite, the post-shock adjustment in the relevant price conjectures coincide with their frictionless counterparts. But as long as the agents are "bounded rational" in the sense that the depth of the process is finite, this process generates a *weaker* adjustment than in the frictionless benchmark. By the same token, it is *as if* the GE spillovers have been reduced. Indeed, in the opposite limit where cognitive depth is zero, each marketplace responds to the aggregate shock *as if* it has been hit by an idiosyncratic shock.

This result offers our first formalization of the sought-after notion that GE adjustment is "attenuated" relative to the frictionless benchmark. Because the PE adjustment that takes place at the marketplace level remains the same as in that benchmark, the gap between the micro and macro elasticities of interest is reduced. Importantly, these properties hold true regardless of whether the GE spillovers reinforce or offset the underlying PE effects, and therefore regardless of whether the frictionless macro elasticity is higher or lower than the micro one.

Lack of Common Knowledge. Our second modification maintains the REE concept but allows the agents to lack common knowledge of the aggregate shock and of the concurrent outcomes in marketplaces other than the ones they currently reside. While conceptually distinct, this modification is shown to have similar observable implications as the one based on Tâtonnement dynamics. In particular, we establish the following equivalence result: for any depth of the cognitive process in Tâtonnement economy, there is a level of common knowledge in the new variant such that the irrational conjectures in the former are recast as the average rational expectations in the latter, and the two economies generate the same observables at the aggregate level; and vice versa.

Let us sketch the argument behind this result. The contraction mapping whose fixed point pins down the REE price conjectures remains the same regardless of whether agents have common knowledge or not. There is, however, a key difference. When the agents have common knowledge of the aggregate shock, the fixed point pins down the adjustment in the REE conjectures as a function of the shock itself. When, instead, agents lack common knowledge, the fixed point pins down the adjustment in REE conjectures as function of the entire hierarchy of beliefs of the shock. Because higher-order beliefs move less than lower-order beliefs, the adjustment is necessarily weaker in the latter case than in the former—which is the same as saying that the GE effect of the aggregate shock is attenuated relative to the frictionless benchmark. Finally, by varying the magnitude of the informational friction, we can vary the degree of attenuation in a similar fashion as when varying the number of iteration in the Tâtonnement variant.

Notwithstanding this equivalence result, we show that the variant based on lack of common knowledge makes the following intriguing prediction: the larger the gap between the micro and the macro elasticity in the frictionless benchmark, the larger the fraction of this gap that gets "erased" by any given level of informational friction. In this sense, the remedy we offer to the disconnect between the micro and the macro elasticities of interest appears to work better the more severe the problem is to start with.

Cobweb, Level-k, and more. Under the lenses of the aforementioned equivalence result, the approach that relaxes the common-knowledge foundations of the standard practice emerges as a modeling substitute to a certain kind of "bounded rationality". Of course, this property does not extend to *every* other plausible relaxation of the REE concept. For instance, suppose we were to adapt the near-rationality approach of Akerlof and Yellen (1985a,b) or the closely related concept of epsilon-equilibrium in games. This modification would let the macro elasticity of the modified economy to vary within a neighborhood of its frictionless counterpart, but would not restrict the relation between the two—the GE effect could by either attenuated or amplified.

With this point in mind, we explore whether the sought-after notion of GE attenuation can be captured by three additional approaches, which also relax the REE concept. The first lets the relevant price conjectures to be determined by an algorithm that mimics Cobweb dynamics instead of Tâtonnement dynamics. The second utilizes an auxiliary result, which permits us to represent our Walrasian economy as a game, and applies the concept of Level-k Thinking

(Nagel, 1995; Stahl, 1996). The third builds on the reflective-equilibrium concept of Garcıa-Schmidt and Woodford (2015).

All these variants predict outcomes that converge to their frictionless counterparts as the "depth of reasoning" increases without bound. They also preserve the underlying PE effect. But do they help accommodate the sought-after notion of GE attenuation? We show the answer is yes for the variant that builds on Garcıa-Schmidt and Woodford (2015), but not necessarily for those based on Cobweb dynamics and Level-k Thinking.

In our framework, the variant based on Cobweb dynamics is equivalent to the one based on Level-k Thinking, in a similar manner as the variant baed on Tâtonnement dynamics is always equivalent to the one based on lack of common knowledge. However, the first two variants are *not* always equivalent to the latter two. Consider, in particular, environments in which the GE effect offsets the PE effect, such as neoclassical settings in which agents compete for fixed resources; these environments are akin to games of strategic substitutability. In such environments, Cobweb dynamics and Level-k Thinking open the door to the opposite conclusion than our previous analysis: they allow the relevant price or quantity conjectures to *overshoot* relative to the frictionless benchmark.

We view such overshooting to be a "pathological" prediction and show how it is avoided by the reflective-equilibrium concept of Garcıa-Schmidt and Woodford (2015). By design, this solution concept has the same flavor and the same conceptual underpinnings as Level-k Thinking. Yet, it entails a certain technical change, which eliminates the aforementioned overshooting possibility.

Dynamic Extension. Our baseline framework help formalize the notion that "GE adjustment is weak". In a dynamic extension, we capture the complementary notion that "GE adjustment takes time" by allowing the relevant conjectures to adjust slowly from the pre-shock frictionless level to the post-shock one. In the variants that rest on relaxing the REE concept, this requires the ad hoc assumption that the "depth of reasoning" increases with time. In the variant with lack of common knowledge, instead, it obtains naturally from the property that the level of common knowledge increases *endogenously* with time as agents observe past market outcomes.

Take-home lessons. The combination of our results yields GE attenuation as a robust prediction of two kinds of frictions: lack of common knowledge of the state of the economy; and certain forms of bounded rationality. In this regard, the two approaches appear to be close substitutes to each other. Nevertheless, the variants that drop the REE concept face few conceptual and practical challenges. These include the obvious vulnerability to Lucas's critique; the aforementioned question of why the depth of reasoning may, or may not, increase with time; and a certain difficulty in accommodating the sought-after notion in stationary environments with recurring shocks.

In our view, this tilts the balance in favor of the approach that maintains rational expectations but removes common knowledge of the state of the economy. For our applied purposes, however, the key lesson remains that GE mechanisms may be less potent in the short run than what the standard practice considers.

The following point is also worth emphasizing. Lack of common knowledge attenuates GE effects regardless of how these effect relate to PE effects. But whether this manifests as under-reaction or over-reaction to the underlying shocks depends, of course, on whether the GE effects were amplifying or offsetting the PE effects. Therefore, our preferred approach does not provide a theory of under-reaction *per se*; rather, it provides a theory of GE attenuation.

Applications. Our framework is too abstract to permit a careful consideration of any particular application. We view this abstraction as a strength, because it allows us to deliver the key insights in a flexible manner and to connect various strands of the literature. We study few applications in companion work, helping translate the abstract insights of this paper to concrete lessons about the sources of the business cycle (Angeletos and Lian, 2016c), the power of forward guidance (Angeletos and Lian, 2016a), and the macroeconomic effects of taxes and government spending (Angeletos and Lian, 2017). These applications differ from one another in terms of the context of interest, the relevant micro-foundations, and the mapping from the theory to the data; they nevertheless all share the theme that removing common knowledge upsets standard predictions by attenuating GE effects.

Related Literature. Our paper builds heavily on the macroeconomic literature on higher-order uncertainty that followed the influential contributions of Morris and Shin (1998, 2002) and Woodford (2003).[5] We borrow from this literature the key observation that higher-order uncertainty impedes coordination and slows down belief adjustment,[6] but add the following elements. First, we relate this property to attenuation of general-equilibrium mechanisms. Second, we highlight that this attenuation is robust to whether the environment features strategic complementarity (GE amplifies PE) or strategic substitutability (GE offsets PE). Third, we clarify that the same attenuation manifests as *under-reaction* to shocks in the one case and as *over-reaction* in the other case. Fourth, we offer this attenuation as a potential remedy to the gap between the macro effects of interest and the kind of micro or local elasticities estimated in, inter alia, Mian, Rao, and Sufi (2013); Mian and Sufi (2012, 2014); Beraja, Hurst, and Ospina (2016); Nakamura and Steinsson (2014). Last but not least, we build a certain bridge between the aforementioned literature, the older tradition on Tâtonnement and Cobweb dynamics, and another literature that drops the REE concept in favor of Level-k Thinking or other related concepts.

At first glance, our work may appear to relate to the literature on micro vs macro rigidity Caplin and Spulber (1987); Caballero and Engel (1999); Golosov and Lucas Jr (2007) or the literature on labor-supply elasticities (Chetty et al., 2011, 2013; Keane, Rogerson et al., 2012). This is not the case, because neither of these literatures is about GE effects. The former is about aggregation in settings with rich heterogeneity and non-linearity in PE effects.[7] The latter is about the appropriate calibration of the key preference or others parameter that determines the PE response of labor supply to variation in wages.

By studying certain alternatives to the REE concept, we connect, not only to the aforementioned work by Garcıa-Schmidt and Woodford (2015), but also to Guesnerie (1992, 2008); Evans and Ramey (1992, 1995); Evans and Honkapohja (2001), and the literature surveyed in Woodford (2013). These prior works are often concerned with how the limits of the certain learning or cognitive algorithms compare to the set of REE. Our paper bypasses these issues by working with a model in which the REE is the unique fixed point of a contraction mapping. It then proceeds to emphasize the GE attenuation that obtains either by considering certain algorithms away from the limit or, in our preferred approach, by maintaining the REE concept but removing common knowledge.

Finally, our paper relates to Gabaix (2016) in the following regard. That paper departs from the REE concept by assuming that the *perceived* law of motion of the aggregate state of the economy is less responsive to the underlying aggregate shocks than the *actual* law of motion. This assumption amounts to attenuating the relevant GE effects. Note, however, that one can obtain the opposite result (amplification of the GE effects) simply by assuming the opposite belief distortion (agents believe that the endogenous state over-reacts). For better or worse, our preferred approach does not provide this degree freedom: because the variation in higher-order beliefs is necessarily bounded by the variation in first-order beliefs, our approach predicts the potency of the GE effects in the absence of common knowledge is *necessarily* lower than in the frictionless benchmark.

Layout. Section 2 introduces the framework. Section 3 studies the frictionless benchmark. Sections 4 and 5 explore our two leading variants: the one that builds on Tâtonnement dynamics; and the one that removes common knowledge of the aggregate shock. Section 6 studies the variants that build on Cobweb dynamics, Level-k Thinking, and Garcıa-Schmidt and Woodford (2015). Section 7 contains take-home lessons and discusses some merits of the approach that retains the REE concept but removes common knowledge. Section 8 develops an extension that captures the notion that GE adjustment takes time. Section 9 concludes.

---

[5] See Angeletos and Lian (2016b) for a survey and additional references. Related are also the literatures on rational inattention (Sims 2003; Mackowiak and Wiederholt 2009) and sticky information (Mankiw and Reis, 2002), but *only* insofar as these literatures open the door to higher-order uncertainty by introducing idiosyncratic noise in the observation of the underlying fundamentals.

[6] Applications of this general idea include Abreu and Brunnermeier (2003); Allen, Morris, and Shin (2006); Angeletos and La'O (2010, 2013); Bacchetta and van Wincoop (2006); Morris and Shin (2006); Nimark (2008, 2017); Venkateswaran (2014) and Venkateswaran (2014).

[7] This literature studies models with fixed costs of adjustment. In such models, the relevant PE effect is zero for firms that stay inside the region of inaction after the shock, but non-zero (and potentially very large) for the firms that are pushed outside that region.

2 Framework

In this section, we introduce our baseline framework. We first spell out the micro-foundations of the framework, namely the market structure and the specification of preferences, technologies, and endowments. We next derive the associated demand and supply functions. From that point and on, we find it useful to work with the log-linearized version of the model. Apart from making the analysis tractable, this permits us to treat the relevant micro and macro elasticities as constants.

2.1 Microfoundations

Marketplaces, goods, and agents. Economic decisions take place over two periods, which we call "morning" and "afternoon"; these can be thought of as proxies for "present" and "future" in contexts with more than two periods, such as the one considered in Section 8. In each period, trade takes place in a continuum of segmented locations, which we call "marketplaces" and index by $m \in [0, 1]$.

There is a double continuum of firms and households, each indexed by $i \in [0, 1] \times [0, 1]$, and three goods. One of these goods is storable and can be used for consumption and/or production purposes in both periods. Possible examples include land and capital, or leisure in model with linear utility in leisure, such as Lagos and Wright (2005). We let this good serve as the numeraire. The remaining two goods, which we call the "morning good" and the "afternoon good", are produced and consumed only in the respective period.

Each marketplace operates two markets: one in the morning, where the morning good is traded with the numeraire; and another in the afternoon, where the afternoon good is traded with the numeraire. In each period, an agent can trade only in the marketplace she is currently in. As time passes, agents can move from one marketplace to another: after the morning markets have closed but before the afternoon ones have opened, each agent receives an idiosyncratic shock that determines whether she stays in her original marketplace or whether she gets relocated to a different marketplace. The probability of an agent's staying in her original marketplace in the "afternoon" is $\rho \in [0, 1)$ and that of relocating is $1 - \rho$. Conditional on reallocation, all other marketplaces are equally likely.

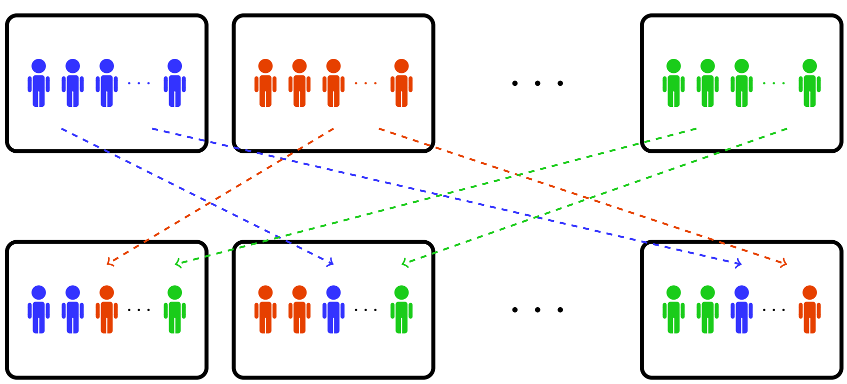

Figure 1: Marketplaces and Trading

These assumptions are illustrated in Figure 1. The marketplaces are represented by the boxes in the figure. In the morning, an equal mass of agents is located to each of the marketplaces. After the morning market has cleared, relocation takes places: while a fraction $\rho \in [0, 1)$ of the agents from each marketplace stay put, the remaining fraction is randomly relocated to all other marketplaces. In the afternoon, each marketplace is therefore populated by two types of agents: a mass ρ of the agents that were originally located in that market; and a representative sample of the agents that where originally spread in the rest of the economy. As mentioned in the Introduction, these stark assumptions help capture two realistic features: that agents participate in a limited number of markets at any given point of time; and that agents nevertheless care about what's going on in the rest of the economy because this affects their future trading opportunities and thereby also their current incentives.

Preferences, technologies and endowments (the "fundamentals") differ across agents. The initial distribution of agents across marketplaces is such that every marketplace receives equal masses of firms and households, but these masses are *not* representative samples of the entire population in the economy. This is illustrated in the figure by the fact that different boxes contain agents of different color. Note then that, whether they relocate or not, agents maintain their original "color" (preferences, technologies, and endowments) as time passes. It follows that the average fundamental in a marketplace changes because, and only because, of the relocation of agents.

Without any loss of generality, we finally let the initial distribution of agents be such that the following is true: in the morning, all the agents in any given marketplace share the same preferences, technologies and endowments. (In the figure, the original agents in any given box have exactly the same color, as opposed to different shades of the same color.) This assumption is completely innocuous for our results; its sole role is to economize on the notation by letting us equate the idiosyncratic fundamental of any given agent with the average fundamental of the marketplace in which that agent resides during the morning. With this in mind, we henceforth let $\theta_m \in \mathbb{R}$ be a variable that parameterizes the fundamentals of the firms and the consumers who trade in marketplace m during the morning. (Think of cross-sectional variation in θ_m as the different colors in the figure.)

Remark. Marketplaces do not have to coincide with geographic regions. Accordingly, the assumption that agents move from one marketplace to another should not be interpreted literally as labor mobility. Instead, this assumption is only meant to capture the fact that every agent typically trades with different sets of other agents at different points of time. That said, when seeking to map the theory to the data, it seems realistic to assume that agents who are closer to each other in terms of geographic distance are *more* likely to participate in the same marketplace than agents that are further away from each other. We use this assumption in Section 3.5—and only there—so as to relate our theoretical contribution to the empirical literature we mentioned in the introduction.

Firms. Consider a firm i that trades in marketplace m during the morning and in marketplace m' during the afternoon. Let q denote its net supply of the morning good, q_i^* its net supply of afternoon good, and q_i^n its net supply of the numeraire. Its technology is represented by a function $K : \mathbb{R}^3 \to \mathbb{R}$ such that

$$q_i^n = -K(q_i, q_i^*; \theta_m) \tag{1}$$

where K is strictly increasing and convex in (q_i, q_i^*), satisfies $K(0, 0; \theta_m) = 0$, and is differentiable in all its arguments. In simpler words, $K(q_i, q_i^*; \theta_m)$ is the real cost of producing the pair (q_i, q_i^*) : it is the amount of the numeraire that the firm must employ as input in order to produce the pair (q_i, q_i^*) as output. It follows that the realized profit of the firm is given by

$$\pi_i = p_m q_i + p_{m'}^* q_i^* + q_i^n = p_m q_i + p_{m'}^* q_i^* - K(q_i, q_i^*; \theta_m),$$

where p_m is the price of morning good in market i and $p_{m'}^*$ is the price of afternoon goods in market m'.

Households. Consider a household i that trades in marketplace m in the morning and marketplace m' in the afternoon. Her preferences are represented by the following utility function:

$$U(c_i, c_i^*; \theta_m) + c_i^n, \tag{2}$$

where U is twice differentiable, strictly increasing, and strictly concave, c_i denotes consumption of the morning goods, c_i^* denotes consumption of the afternoon goods, c_i^n denotes consumption of the numeraire, and θ_m is the aforementioned exogenous fundamental.

The household's endowment of the numeraire is allowed to depend on the local fundamental and is denoted by $e^n(\theta_m)$. Without loss of generality, the endowments of the other two goods are set to zero. The budget constraint of the household is thus given by

$$p_m c_i + p^*_{m'} c^*_i + c^n_i = y_i \equiv e^n(\theta_m) + d_i, \tag{3}$$

where d_i are the dividends the household receives from owning shares on the firms. Because preferences are quasi-linear in the numeraire, the answer to this question is irrelevant for our purposes because any variation in income is absorbed solely by the consumption of the numeraire. With this in mind, we let $d_i = \int_{[0,1]^2} \pi_j dj$, which means that each household owns a fully diversified portfolio of all the firms in the economy.

Remark. We have allowed the technology and the preferences to be non-separable across the two goods so as to introduce interdependence between the morning and the afternoon markets. Such interdependence emerges naturally in applications that feature capital or storable goods. More generally, it means that behavior is forward looking. Without this feature, there would have been no GE spillover across the different marketplaces.

2.2 Demand, Supply, and Market Clearing

Throughout, we impose the following minimal requirements on what agents know and do.

Assumption 1. *Every firm knows her technology; every household knows her preferences and endowments; and every agent knows the price at which she transacts in any given market. Furthermore, every household (respectively, firm) is individually rational in the sense that her chosen quantities maximize her utility (respectively, profits) given the aforementioned knowledge and her subjective belief of either future variables that have not yet being determined or past and current variables that she has not herself observed.*

Assumption 2. *For any m, the agents that are currently in marketplace m share the same subjective beliefs, and this fact is itself known to all of the agents.*

These assumptions are "minimal" in two complementary senses. First, it is implied by the REE concept, which is what the standard practice imposes, but is considerably weaker than it because it allows the aforementioned subjective beliefs to differ from the "objective truth". And second, it suffices for obtaining the demand and the supply schedules in each marketplace under arbitrary subjective beliefs of the outcomes in other marketplaces.

To ease the analysis, we henceforth re-interpret all the variables as log-deviations from a symmetric steady state in which all marketplaces have the same fundamentals, and we work with the log-linearized demand and supply system. With potential abuse of notation, we also let $c_m,\ q_m,\ c^*_m,\ q^*_m$, etc., denote the average (log) consumption and the average (log) output in a marketplace m. We next let $\bar{c} \equiv \int \bar{c}_m dm$, $\bar{q} \equiv \int \bar{q}_m dm$, etc., denote the economy-wide aggregates. We finally let $\hat{E}_m[p^*_{m'}]$ denote the subjective—potentially irrational—belief that the typical agent in marketplace m holds about the price she is likely to face in the afternoon.[8]

Lemma 1. *There exist linear functions D, S, D^*, and S^* such that the following hold for every marketplace m: the morning and afternoon demands are given by, respectively,*

$$c_m = D\left(p_m, \hat{E}_m[p^*_{m'}], \theta_m\right) \qquad \textit{and} \qquad c^*_m = \rho D^*\left(c_m, p^*_m, \theta_m\right) + (1-\rho) D^*\left(\bar{c}, p^*_m, \bar{\theta}\right); \tag{4}$$

and the corresponding supplies are given by, respectively,

$$q_m = S\left(p_m, \hat{E}_m[p^*_{m'}], \theta_m\right) \qquad \textit{and} \qquad q^*_m = \rho S^*\left(q_m, p^*_m, \theta_m\right) + (1-\rho) S^*\left(\bar{q}, p^*_m, \bar{\theta}\right). \tag{5}$$

[8] Keep in mind that, from the perspective of the individual agent, m' is a random variable that is revealed in the afternoon.

This result characterizes the demand-and-supply structure of the economy. Let us explain where it comes from. Thanks to the log-linearization, the morning demand of each household, and similarly the morning supply of each firm, can be expressed as a *linear* function of the local morning price, the local fundamental, and the local subjective belief of the afternoon prices. 1, the left-hand sides of conditions (4) and (5) are self-explanatory. To understand the right-hand sides, note that the demand in any given afternoon market has two components: one reflecting the agents who were in this market from the morning; and another reflecting the agents who were relocated from other markets. The former have mass ρ and their demand is given by $D^*\left(c_m, p_m^*, \theta_m\right)$; the latter have mass $1-\rho$ and their average demand is given $\int D^*\left(c_{m'}, p_m^*, \theta_{m'}\right) dm' = D^*\left(\bar{c}, p_m^*, \bar{\theta}\right)$. The same logic applies on the supply side.

We next impose market clearing.

Assumption 3. *Markets clear:* $c_m = q_m$ *and* $c_m^* = q_m^*$*, for all* m.

How much can we tell about the behavior of the economy on the basis of Assumptions 1, 2, and 3 alone? Unfortunately, they are not enough to predict how the economy responds to the aggregate shock. The reason is that these assumptions do *not* pin down the subjective beliefs that agents may hold in the morning about their trades in the afternoon. Nevertheless, these assumptions permit us to express the market-clearing outcomes of each marketplace as functions of the fundamentals and of these beliefs. We show this in the sequel.

2.3 The Mapping from Fundamentals and Subjective Beliefs to Observables

Consider the morning markets. Using Lemma 1, the net demand in marketplace m can be expressed as follows:

$$n_m \equiv c_m - q_m = D\left(p_m, \hat{E}_m[p_{m'}^*], \theta_m\right) - S\left(p_m, \hat{E}_m[p_{m'}^*], \theta_m\right) = N\left(p_m, \hat{E}_m[p_{m'}^*], \theta_m\right), \tag{6}$$

where $N \equiv D - S$ is the net (or excess) demand function. We henceforth assume that a higher price reduces the excess demand.

Assumption 4. $N(p, \cdot, \cdot)$ *is decreasing in* p.

It follows that, for every marketplace m and every possible value of the pair $(\hat{E}_m[p_{m'}^*], \theta_m)$, there exists a unique p_m that clears the morning market. We therefore obtain the following characterization of the morning outcomes.

Lemma 2. *There exist linear functions* $\tilde{Q}$ *and* $\tilde{P}$ *such that, for every* m *and every* $(\hat{E}_m[p_{m'}^*], \theta_m)$*, the morning outcomes are given by*

$$c_m = q_m = \tilde{Q}\left(\hat{E}_m[p_{m'}^*], \theta_m\right) \qquad \text{and} \qquad p_m = \tilde{P}\left(\hat{E}_m[p_{m'}^*], \theta_m\right).$$

By expressing the market-clearing outcomes in the morning as functions of arbitrary subjective beliefs of the afternoon prices, we build a bridge to the literature on "temporary equilibrium" (Grandmont, 1977). The twist is that, by introducing segmented marketplaces, we have equated the "temporary equilibrium" of the entire economy with the "partial equilibrium" of each marketplace in isolation.

Consider now the afternoon markets. Using Lemma 1 and the fact that $c_m = q_m$ for all m (by market clearing of the morning markets), the net afternoon demand in marketplace m can be expressed as follows:

$$\begin{aligned} n_m^* &\equiv \left\{\rho D^*\left(c_m, p_m^*, \theta_m\right) + (1-\rho) D^*\left(\bar{c}, p_m^*, \bar{\theta}\right)\right\} - \left\{\rho S^*\left(q_m, p_m^*, \theta_m\right) + (1-\rho) S^*\left(\bar{q}, p_m^*, \bar{\theta}\right)\right\} \\ &= \rho N^*\left(q_m, p_m^*, \theta_m\right) + (1-\rho) N^*\left(\bar{q}, p_m^*, \bar{\theta}\right), \end{aligned} \tag{7}$$

where $N^* \equiv D^* - S^*$. To guarantee the existence of a unique p_m^* that clears the afternoon market for every realization of the fundamentals and the morning quantities, we henceforth impose the following variant of Assumption 4.

Assumption 5. *$N^*(\cdot, p^*, \cdot)$ is decreasing in p^*.*

To characterize the market-clearing outcomes, we find it convenient to introduce two auxiliary variables: the average net demand, $\bar{n}^* \equiv \int n_m dm$, and the average price, $\bar{p}^* \equiv \int p^*_m dm$. Aggregating condition (7) gives $\bar{n}^* = N^*\left(\bar{q}, \bar{p}^*, \bar{\theta}\right)$, which in turn means that

$$\bar{n}^* = 0 \qquad \text{if and only if} \qquad \bar{p}^* = P^*(\bar{q}, \bar{\theta}),$$

where P^* is defined so that

$$N^*(q, P^*(q, \theta), \theta) = 0, \tag{8}$$

for all (q, θ). Similarly, using (7), we have that

$$n^*_m = 0 \qquad \text{if and only if} \qquad p^*_m = \rho P^*\left(q_m, \theta_m\right) + (1 - \rho) P^*\left(\bar{q}, \bar{\theta}\right).$$

We therefore obtain the following characterization of the afternoon outcomes.

Lemma 3. *There exist linear functions Q^* and P^* such that, for all realizations of the fundamentals and the morning quantities, the afternoon outcomes are given by*

$$\bar{q}^* = Q^*(\bar{q}, \bar{\theta}), \quad \bar{p}^* = P^*(\bar{q}, \bar{\theta}), \quad q^*_m = \rho Q^*(q_m, \theta_m) + (1 - \rho)\, \bar{q}^* \;\forall m, \quad \textit{and} \quad p^*_m = \rho P^*\left(q_m, \theta_m\right) + (1 - \rho)\, \bar{p}^* \forall m.$$

To sum up, we have shown how the morning outcomes can be expressed as functions of the fundamentals and of certain conjectures, and how the afternoon outcomes can in turn be expressed as functions of the fundamentals and the morning outcomes. To reach this point, we have only relied on two elementary assumptions, namely individual optimality and market clearing, plus the auxiliary assumption that excess demands are downward slopping. To complete the characterization of the observables of the economy, what remains to be done is to determine how the relevant conjectures are formed and how they respond to shocks. This requires additional, and more delicate, assumptions regarding the solution concept and the information structure. We elaborate on this issue in the subsequent sections. We close the present section by specifying the exogenous shock that that triggers a change in both the fundamentals and the relevant beliefs.

2.4 The Exogenous Shock

We conclude the description of our framework by specifying the aggregate shock to which the economy's response we study in the sequel. Denote with $\bar{\theta} = \int \theta_m dm$ the aggregate, or average, fundamental. Throughout, we let $\bar{\theta}$ be a single-dimensional variable, $\bar{\theta} \in \mathbb{R}$, and focus on a once-and-for-all change in it, from some initial level, $\bar{\theta} = \bar{\theta}_{old}$, to some new level, $\bar{\theta} = \bar{\theta}_{new} \neq \bar{\theta}_{old}$. We treat the initial level, $\bar{\theta}_{old}$, as a fixed parameter (and hence as a commonly known object) and the change, $\Delta\bar{\theta} \equiv \bar{\theta}_{new} - \bar{\theta}_{old}$, as a random variable drawn from a Normal distribution centered around 0. We next specify the corresponding changes in the "local" fundamentals as follows:

$$\Delta\theta_m = \delta_m \Delta\bar{\theta} + z_m \tag{9}$$

where $\Delta\theta_m$ is the change in the fundamental if marketplace m in the morning, $\Delta\bar{\theta}$ is the aggregate shock, δ_m is a fixed parameter that captures the exposure of m to the aggregate shock, with $\int \delta_m dm = 1$, and z_m is a purely idiosyncratic shock. The latter is independent of $\Delta\bar{\theta}$, is drawn from a Normal distribution whose mean is zero and whose p.d.f. is henceforth denoted by φ, and is such that $\int z_m dm = \int z\varphi(z)dz = 0$ for all realizations of uncertainty.

The applied question of interest can then be posed as follows: how do the observables of the economy (quantities or prices) respond to the aforementioned shock? Answering this question boils down to answering two subquestions: how the shock shifts demand and supply in each marketplace for given conjectures of the outcomes in other marketplaces; and how it shifts the conjectures themselves. The answer to the first subquestion can readily be obtained from Lemmas 2 and 3 and relies only on the elementary assumptions that markets clear and that agents optimize under local knowledge of their fundamentals and of the prices at which they currently trade. By contrast, the answer to the second subquestion requires a choice of a solution concept and a specification of the information that the agents have about markets they do *not* currently participate in.

3 The Frictionless Benchmark

Our benchmark is defined by imposing the Rational Expectations Equilibrium (REE) concept along with the assumption that, at any given period, the cross-sectional profile of the fundamentals is common knowledge. More specifically, we assume that all agents share the same information in the morning and this information contains the entire profile of $(\theta_m, p_m)_{m\in[0,1]}$ in the economy; all agents know this fact; all agents know that all agents know this fact; and so on. Admittedly, this is a strong assumption; but it is the standard one.

Remark 1. By itself, the assumption that agents share the same information in the morning does not pin down their subjective beliefs about the afternoon prices. In combination with the REE concept, however, this assumption guarantees that all agents share the same subjective beliefs and that these beliefs indeed coincide with the objective, rational, expectation of the afternoon prices. In what follows, we prove this property and proceed to characterize the observables.

Remark 2. In our setting, the rational expectation of the afternoon prices happen to coincide with the actual realizations of these prices, due to our simplifying assumption that there are not shocks to fundamentals between the morning and the afternoon. However, as explained below, our characterization of the relevant expectations and of the associated morning outcomes is robust to dropping this simplifying assumption.

3.1 From Subjective Conjectures to Rational Expectations

Consider an arbitrary market m in the morning. Lemma 2 together with the assumption that agents know the local θ_m imply that, in any REE, the agents also know q_m. By Lemma 3, we then have that

$$\hat{E}_m[p^*_{m'}] = \mathbb{E}\left[\rho p^*_m + (1-\rho)\bar{p}^*\right] = \rho^2 P^*\left(q_m, \theta_m\right) + \left(1-\rho^2\right)\mathbb{E}\left[\bar{p}^*\right], \tag{10}$$

where $\mathbb{E}$ henceforth denotes the rational expectation operator conditional on the information that is available in the morning. By Lemma 2, on the other hand,

$$q_m = \tilde{Q}\left(\hat{E}_m[p^*_{m'}], \theta_m\right).$$

Let $\tilde{\alpha} \equiv \frac{\partial P^*}{\partial q}\frac{\partial \tilde{Q}}{\partial p^*}$. Provided that $\rho^2\tilde{\alpha} \neq 1$, we can combine the above two equations and solve for q_m as a function of θ_m and $\bar{p}^*$, and can then obtain the corresponding aggregate as a function of $\bar{\theta}$ and of the rational expectation of $\bar{p}^*$. We thus have the following result.

Lemma 4. *Suppose $\rho^2\tilde{\alpha} \neq 1$. There exist linear functions Q and P such that*

$$\bar{q} = Q(\mathbb{E}\left[\bar{p}^*\right], \bar{\theta}), \quad \bar{p} = P(\mathbb{E}\left[\bar{p}^*\right], \bar{\theta}), \quad q_m = Q(\mathbb{E}\left[\bar{p}^*\right], \theta_m)\ \forall m, \quad \textit{and} \quad p_m = P(\mathbb{E}\left[\bar{p}^*\right], \theta_m)\ \forall m.$$

What remains is to pin down $\mathbb{E}[\bar{p}^*]$. We now derive the fixed-point relation that determines $\mathbb{E}[\bar{p}^*]$ under the REE concept. This concept implies that the agents inside the model—just like us, the theorists—know that equilibrium outcomes satisfy Lemmas 3 and 4 (and they know that others know, ad infinitum). From the former lemma, $\bar{p}^* = P^*(\bar{q}, \bar{\theta})$; from the latter, $\bar{q} = Q\left(\mathbb{E}[\bar{p}^*], \bar{\theta}\right)$. Combining the two, the agents can infer that the realized value of $\bar{p}^*$ and its rational expectation must satisfy the following restriction:

$$\bar{p}^* = \mathcal{T}\left(\mathbb{E}[\bar{p}^*], \bar{\theta}\right), \tag{11}$$

where $\mathcal{T}(p^*, \theta) \equiv P^*(Q(p^*, \theta), \theta)$. Taking expectations on both sides gives

$$\mathbb{E}[\bar{p}^*] = \mathcal{T}\left(\mathbb{E}[\bar{p}^*], \bar{\theta}\right). \tag{12}$$

It follows that we and the agents know the following two facts: that $\bar{p}^* = \mathbb{E}[\bar{p}^*]$ always; and that $\mathbb{E}[\bar{p}^*]$ is itself pinned down by the fixed points of the mapping $\mathcal{T}$.

As noted earlier, the exact coincidence between $\bar{p}^*$ and $\mathbb{E}[\bar{p}^*]$ is due to the assumption that no innovation in the aggregate fundamentals is possible between the morning and the afternoon.[9] If we were to relax this assumption, the morning quantities would still satisfy $\bar{q} = Q\left(\mathbb{E}[\bar{p}^*], \bar{\theta}\right)$, but now the afternoon prices would be given by $\bar{p}^* = P^*(\bar{q}, \bar{\theta}) + \epsilon$, where ϵ is a term that captures the effect of the realized innovation. As a result, condition (11) would now have to be replaced by $\bar{p}^* = \mathcal{T}(\mathbb{E}[\bar{p}^*], \theta) + \epsilon$. Nevertheless, because ϵ is unpredictable in the morning, condition (12) would still hold. We conclude that, even when the realized $\bar{p}^*$ varies around $\mathbb{E}[\bar{p}^*]$ due to innovations in fundamentals, $\mathcal{T}$ is the mapping whose fixed points pin down the values of $\mathbb{E}[\bar{p}^*]$ that are consistent with REE.

Note that the slope of $\mathcal{T}(p^*, \cdot)$ with respect to p^* is given by

$$\alpha \equiv \frac{\partial \mathcal{T}}{\partial p^*} = \frac{1-\rho^2}{1-\tilde{\alpha}\rho^2}\tilde{\alpha} \tag{13}$$

where $\tilde{\alpha} \equiv \frac{\partial P^*}{\partial q}\frac{\partial \tilde{Q}}{\partial p^*}$. To guarantee that $\mathcal{T}$ is contraction mapping for every ρ, we make the following assumption.

Assumption 6. $\tilde{\alpha} \in (-1, 1)$.

This assumption is, not only sufficient for $\mathcal{T}$ to be a contraction mapping, but also sufficient for Lemma 4 to hold, which itself is needed for the construction of $\mathcal{T}$ in the first place. Given this assumption, *we* can solve (12) for $\mathbb{E}[\bar{p}^*]$ as a function of $\bar{\theta}$—and so can the agents inside the model. We thus reach at the following result, which completes the equilibrium characterization of our benchmark.

Proposition 1. *The equilibrium exists, is unique, and is such the following is true:*

(i) the rational expectation of $\bar{p}^*$ is given by

$$\mathbb{E}[\bar{p}^*] = \mathcal{P}(\bar{\theta}) \tag{14}$$

where the function $\mathcal{P}$ is defined by the fixed point of $\mathcal{T}$;

(ii) the realized prices and quantities in the morning are given as in Lemma 4, with $\mathbb{E}[\bar{p}^*]$ determined as above.

Remark. In the preceding analysis, we have assumed that agents have common knowledge of $(\theta_m, p_m)_{m\in[0,1]}$, the entire cross-sectional profile of the fundamentals and the prices. In our setting, the same outcomes can be attained under the weaker assumption that agents have common knowledge of merely $\bar{\theta}$ and $\bar{p}$. However, in settings with richer forms of matching, common knowledge of the average fundamental and/or the average price is generally not enough

[9] By an "innovation" we mean a change that is unpredictable in the morning. Predictable changes (such as "news shocks") are already nested, because we have not restricted how θ enters preferences and technology.

for replicating the frictionless outcomes. In practical terms this means that, in a complex economy, outcomes can depart in a significant manner from the frictionless benchmark even if there is readily available information about macroeconomic statistics and even if agents pay full attention to them.[10]

3.2 Micro vs Macro, and PE vs GE

We are now ready to characterize how the morning outcomes respond to aggregate shocks. We focus on morning outcomes because we think of them as better proxies for the kind of outcomes that may be observable to an "econometrician". This is because "afternoon" in our model proxies the role of all the relevant future market interactions, which may have not even be realized by the time the econometrician makes her measurements. To simplify the exposition, we also focus on quantities; similar results apply for prices as well.

Combining the two parts of Proposition 1, we have that the realized quantities at the local and aggregate level are given by, respectively,

$$q_m = Q(\mathcal{P}(\bar{\theta}), \theta_m) \qquad \text{and} \qquad \bar{q} = Q(\mathcal{P}(\bar{\theta}), \bar{\theta}). \tag{15}$$

Letting Δ denote the change in a variable relative to its pre-shock value, we reach the following characterization:

Proposition 2. *There exist scalars ϵ^{micro} and ϵ^{Macro} such that, for all realizations of the underlying aggregate and idiosyncratic shocks, the corresponding changes in the equilibrium quantities are given by*

$$\Delta q_m = \Delta\bar{q} + \epsilon^{micro}\left(\Delta\bar{\theta}_m - \Delta\bar{\theta}\right) \qquad \textit{and} \qquad \Delta\bar{q} = \epsilon^{Macro}\Delta\bar{\theta}.$$

The scalar ϵ^{Macro} is given by

$$\epsilon^{Macro} \equiv \frac{\partial Q}{\partial \theta} + \frac{\partial Q}{\partial p^*}\frac{\partial \mathcal{P}}{\partial \theta}$$

and measures the change in aggregate activity relative to the change in aggregate fundamentals. This is the type of object that macroeconomists are primarily interested in. The scalar ϵ^{micro}, on the other hand, is given by

$$\epsilon^{micro} \equiv \frac{\partial Q}{\partial \theta}$$

and measures the change in local activity relative to the change in local fundamentals, holding aggregate conditions constant. We henceforth refer to ϵ^{Macro} as the "macro elasticity" and to ϵ^{micro} as the "micro elasticity".[11]

The precise values of ϵ^{micro} and ϵ^{Macro} depend on the underlying micro-foundations and vary from application to application. For our purposes, neither the precise values of these elasticities nor the underlying micro-foundations are essential. Instead, the key is only that ϵ^{Macro} typically differs from ϵ^{micro}, due to a general-equilibrium effect.

Let us elaborate by considering three scenarios. In first scenario, we consider an aggregate shock and let $\mathbb{E}[\bar{p}^*]$ adjust accordingly. In the second, we consider the same aggregate shock but fix the subjective belief of $\bar{p}^*$ to its pre-shock value. In the last scenario, we consider an idiosyncratic shock that affects only a particular marketplace. Because each marketplace is of infinitesimal size, such a shock triggers an adjustment in the local outcomes without triggering any adjustment in any other market.

The first scenario identifies ϵ^{Macro}. The second scenario isolates the partial-equilibrium effect of the aggregate shock: it captures the effect that the shock has on the typical marketplace under the assumption that this marketplace expects the other marketplaces not to adjust to the shock. The difference between the scenarios identifies the kind of

[10]Allowing agents to be "rationally inattentive" a la Sims (2003) reinforces this point, to the extent that rational inattention boils down to each agent observing a more noisy private signal of aggregate economic conditions than of idiosyncratic or local economic conditions.

[11]Had we allowed for richer heterogeneity, the sensitivity of q_m to θ_m could differ across marketplaces. In that case, we would define ϵ^{micro} as a weighted average of the m-specific elasticities.

general-equilibrium adjustment we are after in this paper. Finally, the second and the third scenario result in the same outcomes, meaning that the partial-equilibrium effect of an aggregate shock on the aggregate quantity coincides with the average equilibrium effect of an idiosyncratic shock on local quantity.

We formalize these ideas in the following corollary.

Corollary 1. *The macro elasticity can be decomposed to a partial-equilibrium and a general-equilibrium component as follows:*

$$\epsilon^{Macro} = \mathtt{PE} + \mathtt{GE},$$

where

$$\mathtt{PE} \equiv \int \left(\frac{\partial Q}{\partial \theta} \delta_m \right) dm = \epsilon^{micro} \qquad \textit{and} \qquad \mathtt{GE} \equiv \frac{\partial Q}{\partial p^*} \frac{\partial \mathcal{P}}{\partial \theta}.$$

The term denoted by PE captures the aggregate response that obtains when we hit all markets at once with an aggregate shock and nevertheless have each market respond *as if* the shock was specific to that market or, equivalently, *as if* the agents expected $\bar{p}^*$ to stay constant. The term denoted by GE, instead, captures the *additional* effect that obtains once we, and the agents inside the model, take into account that $\bar{p}^*$ itself adjusts in response to the aggregate shock so as to clear the afternoon markets.[12]

As evident from the above corollary, the aforementioned GE effect identifies the gap between ϵ^{Macro} and ϵ^{micro}. To make the analysis interesting, we henceforth impose the following.

Assumption 7. *The GE effect is non-zero:* $\frac{\partial Q}{\partial p^*} \frac{\partial \mathcal{P}}{\partial \theta} \neq 0$.

We conclude this section by noting that this gap can be either negative or positive. The case in which $\epsilon^{Macro} > \epsilon^{micro}$ captures settings in which the GE effect works in the same direction as the PE effect, acting as an amplification mechanism. The alternative case, $\epsilon^{Macro} < \epsilon^{micro}$, captures settings in which the GE effect works in the opposite direction than the PE effect. We illustrate both scenarios in Figure 2, using the example of a demand shock. The top panel in the figure shows the PE adjustment of the market of morning goods, that is, the adjustment that would have taken place if the shock had been idiosyncratic. The middle panel shows the GE adjustment for a case in which the GE amplifies the PE effect. The bottom panel illustrates a case in which the GE effect negates the PE effect *exactly*.

The results of our paper hold true regardless of whether the GE effect amplifies or offsets the PE effect. What is key to note is only the following. So far, the theory predicts that the aggregate shock causes the economy to jump from point X to point Z, regardless of where Z is located relative to Y. This underscores that the decomposition between PE and GE effect is *irrelevant* for the observables of the economy at the aggregate level: all that matters is the *total* macroeconomic effect, not its decomposition to PE and GE components. By contrast, the modifications we study in Sections 4 and 5 will try to make sense of why this decomposition may be relevant and why the GE adjustment may be "partial" in the sense that economy moves to a point in the middle of the segment between Y and Z rather than all the way to Z.

3.3 Connection to Arrow-Debreu and Additional Clarifications

We now show that the outcomes of our frictionless benchmark coincide with those of an appropriate Arrow-Debreu variant, which lets all markets operate at once and recasts the trading risk of our framework as idiosyncratic technology and preference shocks.[13] This helps clarify some basic ideas and refine the context of our contribution.

[12] A few clarifications about terminology. First, the *overall* general-equilibrium response of $\bar{q}$ to the aggregate shock is of course given by the entire ϵ^{Macro}; what we henceforth call GE effect is the component of the overall effect that is due to the adjustment of p^*, that is, the effect that occurs beyond the PE effect. Second, what we call PE effect in this paper is sometimes referred to as *local* GE effect in applied work. Finally, whereas we use the term "effect" to refer to the *marg*inal effect of a shock, in applied work the same term is often used to refer to the marginal effect times the size of the shock..

[13] For the purposes of this subsection, we momentarily revert to the interpretation of the variables prior to the log-linearization we have employed in the preceding equilibrium derivations.

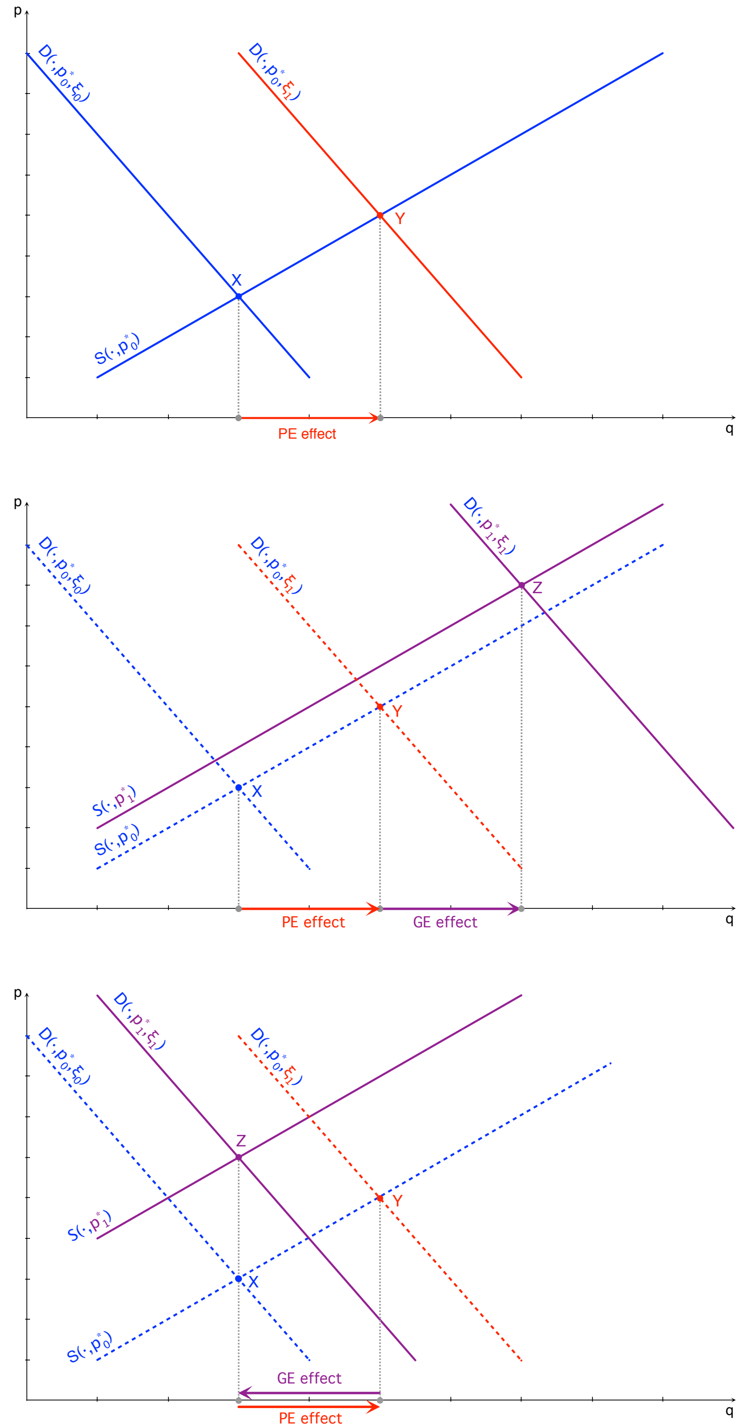

Figure 2: PE and GE effects

In the considered variant, there is neither sequential trading nor segmented marketplaces. Instead, there is a single centralized market, which operates only once, but allows agents to trade over a sufficiently rich set of commodities. This set includes a numeraire and multiple varieties of "morning" and "afternoon" goods. These varieties are indexed by $(m, m') \in [0,1]^2$. Finally, there are a double continuum of households and a double continuum of firms, each indexed by $i = (i_1, i_2) \in [0,1]^2$.

Every household likes to consume a single variety of the morning good and a single variety of the afternoon good. She knows a priori the variety of the morning good she likes, but faces idiosyncratic uncertainty about the likable variety of the afternoon good. Fix an m and consider any household i who likes variety m of the morning good. The probability that she likes variety $m' = m$ of the afternoon good is given by $\rho \in [0,1)$. With probability $1-\rho$, the variety she likes is drawn from a uniform distribution over $[0,1]$. Finally, her preferences (expected utility) are given by

$$\mathcal{U}_i = \rho \left[U\left(c_i, c^*_{i,m}, \theta_m\right) + c^n_{i,m} \right] + (1-\rho) \int \left[U\left(c_i, c^*_{i,m'}, \theta_m\right) + c^n_{i,m'} \right] dm',$$

where $c^*_{i,m'}$ and $c^n_{i,m'}$ is household i's consumption of afternoon goods and numeraire if she likes variety m' in the afternoon. As standard in the Arrow-Debreu paradigm, the household's uncertainty is subsumed in her preferences: it is *as if* the household faces no uncertainty, likes all the varieties of the afternoon good, and happens to have the preferences defined above. By the same token, because the Arrow-Debreu structure permits each agent to make her purchase of each variety of the afternoon good contingent on the realization of her idiosyncratic uncertainty, her budget can be expressed as follows:

$$p_m c_i + \rho \left(p^*_m c^*_{i,m} + c^n_{i,m} \right) + (1-\rho) \int \left(p^*_{m'} c^*_{i,m'} + c^n_{i,m'} \right) dm' = y_i,$$

where $y_i = e_i(\theta) + \Pi$ and where Π are the total profits in the economy.

We make similar assumptions on the production side. Each firm can produce a single variety of the morning good and a single variety of the afternoon good; knows a priori the morning variety; and faces idiosyncratic uncertainty about the afternoon variety. Consider a firm i that produces variety m of the morning good. It produces variety $m' = m$ of the afternoon good with probability $\rho \in [0,1)$; and with the remaining probability, the variety she can produce is drawn from uniform distribution over $[0,1]$. It follows that the firm's expected profit is given by

$$p_m q_i + \rho \left\{ p^*_m q^*_{i,m} - K(q_i, q^*_{i,m}) \right\} + (1-\rho) \left\{ \int p^*_{m'} q^*_{i,m'} - K(q_i, q^*_{i,m'}) \right\} dm',$$

where $q^*_{i,m'}$ and $c^n_{i,m'}$ is firm i's production of afternoon goods and numeraire if she likes variety m' in the afternoon.

It is straightforward to show that the variant described above gives rise to the same equilibrium prices and quantities as our frictionless benchmark. The only difference is the following: what used to be forward-looking price conjectures in that benchmark have been recast as actual prices in the present variant. By the same token, what was a *dynamic* GE effect in the former has been recast as a *static* GE effect in the latter.

3.4 Remarks

Remark 1. If one takes the view that any given agent can engage in a *single* market interaction at any given instant of time, then all GE effects are of a dynamic nature and are tied to expectations. This perspective explains our choice to equate the GE effects that operate across marketplaces to a certain kind of expectations. But as we have just illustrated, these expectations can play an essential role only once we depart from our frictionless benchmark: otherwise, it is *as if* all GE effects are static and expectations are inactive.

Remark 2. Our frictionless benchmark is akin to assuming that all agents can get together in the same room and can *perfectly* and *instantaneously* coordinate both their current and their future reaction to the underlying aggregate shock. Here, this point was formalized by showing that our benchmark attains the same outcomes as an Arrow-Debreu variant that leaves essentially no room for either dynamics or expectations. Later on, we will further corroborate this point by showing that our frictionless benchmark is akin to a static, complete-information, game in which players face no uncertainty about one another's actions.

Remark 3. Applied research in macroeconomics (and elsewhere) often departs from the Arrow-Debreu framework by allowing for incomplete markets, monopoly power, and the like. Yet, by imposing the REE concept along with public information about aggregate shocks, standard practice preserves the aforementioned kind of frictionless coordination in the adjustment of beliefs, behavior, and prices to aggregate shocks. By contrast, what we are after in this paper is precisely the introduction of a certain friction in this kind of adjustment—a friction that can be interpreted as attenuating the relevant GE effects.

Remark 4. For applied purposes, it is often appealing to abstract from the fact that most real-world trading is sequential and, instead, represent the economy with a static model. In such models, one can still tie GE effects to expectations by letting each household or firm have multiple selves, each of which engages in one market interaction and must forecast the choices of other shelves. Think of this as a "big family" or a "team problem": every member of the family/team makes a choice that contributes towards a common goal, without necessarily having knowledge of the choices of other members.[14]

3.5 Connection to Empirical Work

Recently, there has been a boom of empirical research trying to gauge the macroeconomic effects of aggregate shocks by exploiting the cross-sectional heterogeneity in the exposure of different geographical regions to these shocks. Important examples include Mian and Sufi (2012, 2014); Mian, Rao, and Sufi (2013); Beraja, Hurst, and Ospina (2016) and Beraja, Hurst, and Ospina (2016) in the context of the Great Recession; Nakamura and Steinsson (2014) in the context of fiscal multipliers; and Autor, Dorn, and Hanson (2013) and Acemoglu et al. (2016) in the context of how import competition from China has effected US labor market outcomes.

The type of empirical exercises conducted in these works can be represented in our framework as follows. Suppose that the data contain observations of $\Delta\bar{\theta}$, the aggregate shock of interest, and of $\Delta\bar{q}$, the corresponding change in the outcome of interest.[15] Suppose further that, apart from the shock of interest, there are other shocks that are neither of interest to nor observed by the econometrician. It follows that

$$\Delta\bar{q} = \epsilon^{Macro}\Delta\bar{\theta} + \varepsilon, \tag{16}$$

where ϵ^{Macro} is the macro elasticity of interest and ε is residual that captures the other shocks.

Clearly, an unbiased estimate of ϵ^{Macro} can be obtained on the basis of aggregate times series only if ε is uncorrelated with $\Delta\bar{\theta}$ or the econometrician has an instrument for $\Delta\bar{\theta}$ that is itself uncorrelated with ε. In practice, these conditions are rarely met. To overcome this limitation, the aforementioned works shift focus to the cross section and offer a credible instrument for the differential exposure of different regions to the shock. There is, however, an important caveat: what is actually estimated is a certain kind of micro elasticity, rather than the macro elasticity of interest.

[14] See Angeletos and Lian (2016c) for an application along these lines.

[15] For the purposes of the present discussion, one can think of the econometrician observing $\bar{\theta}$ and $\bar{q}$ at multiple points of time, each of which correspond to a *different* morning in an appropriate multi-period version of our framework. See Section 8 for such an extension.

To see this more clearly, suppose that a marketplace in the theory corresponds to a "region" (the level of observation) in the data. The change in the regional outcome of interest can then be expressed as follows:

$$\Delta q_m = \frac{\partial Q}{\partial \theta}(\delta_m \Delta\bar{\theta} + z_m) + \frac{\partial Q}{\partial p^*}\frac{\partial \mathcal{P}}{\partial \theta}\Delta\bar{\theta} + (\varepsilon + \xi_m)$$

where ε and ξ_m capture, respectively, the aggregate and the idiosyncratic effects of the other, unobserved, shocks. Using the facts that $\frac{\partial Q}{\partial \theta} = \texttt{PE} = \epsilon^{micro}$ and $\frac{\partial Q}{\partial p^*}\frac{\partial \mathcal{P}}{\partial \theta} = \texttt{GE} = \epsilon^{Macro} - \epsilon^{micro}$, we can restate the above as his can be

$$\Delta q_m = \epsilon^{micro}\delta_m \Delta\bar{\theta} + \eta + v_m$$

where

$$\eta \equiv \varepsilon + (\epsilon^{Macro} - \epsilon^{micro})\Delta\bar{\theta} \qquad \text{and} \qquad v_m \equiv \xi_m + \epsilon^{micro} z_m.$$

Note that η is the same across all regions is thus subsumed by the constant of a cross-sectional regression or, if we have longer data, by the time-fixed effect of a panel regression. Having a credible instrument for the differential exposure to the aggregate shock of interest means that, once the time-fixed effect has been partial out, the available instrument covaries with $\delta_m \Delta\bar{\theta}$ but not with v_m, the residual of the aforementioned regression. This permits the econometrician to obtain an unbiased estimate of ϵ^{micro}, which is valuable—but unless the gap between ϵ^{Macro} and ϵ^{micro} happens to be small, this estimate gives little information about ultimate object of interest, namely about the macroeconomic effect of $\Delta\bar{\theta}$.

This epitomizes the conundrum faced by the aforementioned line of empirical work: allowing for a time fixed effect in the relevant regressions partials out, not only the concurrent shocks that contaminate the aggregate time series, but also the GE effects of the shock of interest. In the sequel, we offer a partial resolution to this conundrum: by a offering a rationale for why at least some of these GE effect may be impotent in the short run, we reduce the gap between the theoretical object that is of interest and the one that is actually estimated in the aforementioned line of work.

Remark 1. The above discussion has assumed that a "marketplace" in our theory coincides with a "region" in the data (that is, with the level of observation in the available cross-sectional data). In the absence of such a coincidence, the mapping between the theory and the data is more nuanced. However, provided that any two agents who live in the same region are more likely to participate in the same markets than any two agents from different regions, the essence of the conveyed message is likely to survive.

Remark 2. The above discussion has also assumed that the micro elasticity is the same across all regions. In empirical work, this assumption is often relaxed. Our point remains valid provided that one reinterprets ϵ^{micro} as the appropriate cross-sectional average of the regional micro elasticities.

Remark 3. The aforementioned empirical literature often emphasizes differential effects of the same shock on different kinds of economic outcomes, such as employment in tradable versus non-tradable sectors in the case of Mian and Sufi (2014). The points made above apply regardless: for each of the considered outcomes, the aforementioned work estimates, at best, a local elasticity.

4 Cognitive Tâtonnement

Our first modification lets the agents act on the basis of *irrational* conjectures of future prices. These conjectures, however, are not entirely arbitrary. Instead, they satisfy two key properties. The first is that the agents understand fully how local outcomes are affected by local shocks; in effect, this means that we retain the REE concept at the marketplace/PE level, even though we relax it at the aggregate/GE level. The second is that agents adjust their conjectures of aggregate outcomes using an algorithm that resembles Tâtonnement dynamics. Relative to the traditional definition of Tâtonnement dynamics, the twist here is that we are not introducing any *actual* dynamics. Instead, we are merely replacing rational expectations with another *instantaneous* cognitive process, which happens to be defined by the solution to a certain differential equation as opposed to the REE fixed point.

Let us elaborate. The assumed cognitive procedure consists of multiple rounds of making an initial conjecture about the average afternoon price, $\bar{p}^*$, calculating the implied imbalance in the market for afternoon goods, and subsequently updating the original conjecture in the direction that helps reduce the imbalance. We index the rounds by t, treat t as a continuous variable, and denote with $T \in (0, \infty)$ the *final* round, that is, the point at which the guess-and-update iterations stop and actual behavior gets determined. Although it may be more natural to think of "rounds" as a discrete variable, letting t be a continuous variable is consistent with existing treatments of Tâtonnement dynamics;[16] for our purposes, the advantage of treating t as a continuous variable is to let the price conjecture $\hat{p}^*$ be a continuous function of T.

Definition 1. *Fix a $T \in (0, \infty)$. The Tâtonnement solution is given by a conjecture $\hat{p}^*$ and a pair (q_m, p_m) for every m such that the following hold:*

(i) For all m, the pair (q_m, p_m) is consistent with household optimality, firm optimality, and market clearing in the local market for morning goods, under the following conjectures for the afternoon prices:

$$\hat{E}_m[p^*_{m'}] = \rho^2 P^*\left(q_m, \theta_m\right) + \left(1 - \rho^2\right) \hat{E}_m[\bar{p}^*] \qquad \text{and} \qquad \hat{E}_m[\bar{p}^*] = \hat{p}^* \tag{17}$$

(ii) $\hat{p}^*$ is given by $\hat{p}^* = \hat{P}^*(T)$, where the function $\hat{P}^*$ is obtained by solving the following ODE:

$$\frac{d\hat{P}^*(t)}{dt} = \mathcal{N}\left(\hat{P}^*(t), \bar{\theta}_{new}\right) \quad \forall t \geq 0 \tag{18}$$

with initial condition $\hat{P}^*(0) = \bar{p}^*_{old} \equiv \mathcal{P}(\bar{\theta}_{old})$, where $\mathcal{N}\left(\hat{p}^*, \bar{\theta}\right) = N^*\left(Q\left(\hat{p}^*, \bar{\theta}\right), \hat{p}^*, \bar{\theta}\right)$ is the net aggregate demand for afternoon goods if its average price is $\hat{p}^*$.[17]

Recall that the combination of household optimality, firm optimality, and market clearing imply that the morning outcomes satisfy Lemma 2 for *some* subjective beliefs about the afternoon prices. Part (i) of the above definition requires that these conjectures satisfy condition (17). This condition imposes the aforementioned property that agents understand how *local* afternoon prices are affected by current local conditions, but takes as given the conjecture $\hat{p}^*$ that agents form about *average* afternoon prices.

Part (ii) then completes the definition of the solution concept by specifying the conjecture $\hat{p}^*$ as the product of the following iterative procedure. An initial conjecture is formed at round $t = 0$ by letting $\hat{p}^*(0)$ coincide with $\bar{p}^*_{old}$, the *pre-shock* equilibrium price. Using this conjecture, the "Walrasian auctioneer" (herein interpreted as a fictional entity residing inside the minds of the typical agent) computes the implied aggregate demand and supply for afternoon goods and subsequently adjusts her conjecture depending on whether the implied aggregate demand of afternoon goods was in excess of aggregate supply. That is, if $\mathcal{N}\left(\hat{p}^*(0), \bar{\theta}_{new}\right) > 0$, the conjecture is adjusted upwards; if, instead, $\mathcal{N}\left(\hat{p}^*(0), \bar{\theta}_{new}\right) < 0$, the conjecture is adjusted downwards. The same updating procedure is applied at all $t \in (0, T)$. At $t = T$, the updating terminates, generating the final conjecture, $\hat{p}^* = \hat{P}^*(T)$, upon which actual behavior is based.[18]

[16] See, e.g., Section 17.H in Mas-Collell, Whinston, and Green (1995).

[17] The function $\mathcal{P}$ is the same as that obtained in condition 14.

[18] Our results extend directly to a generalized version of (18) that lets $\frac{d\hat{P}^*(t)}{dt} = b(t)\mathcal{N}\left(\hat{P}^*(t), \bar{\theta}_{new}\right)$, where $b : \mathbb{R}_+ \to \mathbb{R}_+$ is any function that is bounded away from zero. Insofar as t is continuous, varying the function b is the same as varying T, and therefore letting $b(t) = 1$ for all t is merely a simplification. However, if we had treated t as a discrete variable, allowing for a more flexible b could be useful for the reasons explained shortly.

In Section 8, we will develop a variant that allows the relevant economic decisions to be repeated over many periods. In that variant, we will tie the passage of time with an increase in T and, in this sense, give a real-time interpretation to the assumed cognitive procedure. For the time being, however, we interpret a higher T as a higher "depth of reasoning", in the sense that of a larger number of updates of the initial price conjecture.[19]

In the rest of this section, we characterize the price conjectures and the outcomes that obtain when the REE concept is replaced with the solution concept in Definition 1. We then show how this attenuates the GE adjustment relative to our frictionless benchmark, thus also helping reduce the gap between the micro and the macro elasticities.

Consider first the price conjecture $\hat{p}^*$. By Assumptions 5 and 6, $\mathcal{N}(p^*, \cdot)$ is decreasing in p^*, implying that the following is true.

Lemma 5. *There exists a continuous and strictly increasing function* $w : \mathbb{R}_+ \to [0,1]$, *with* $w(0) = 0$ *and* $\lim_{T\to\infty} w(T) = 1$, *such that, for any* T *and any* $\Delta\bar{\theta}$, *the conjecture* $\hat{p}^*$ *satisfies*

$$\hat{p}^* = \bar{p}^*_{old} + w(T)\left(\bar{p}^*_{new} - \bar{p}^*_{old}\right) \tag{19}$$

where $\bar{p}^*_{old} \equiv \mathcal{P}(\bar{\theta}_{old})$, $\bar{p}^*_{new} \equiv \mathcal{P}(\bar{\theta}_{new})$, *and* $\bar{\theta}_{new} \equiv \bar{\theta}_{old} + \Delta\bar{\theta}$.

By varying T, we can thus let the conjecture $\hat{p}^*$ span the entire interval between $\bar{p}^*_{old}$ and $\bar{p}^*_{new}$, that is, we can span the entire adjustment that happens in $\bar{p}^*$ in the REE benchmark.[20] Furthermore, the higher T is, the closer the conjecture is to the post-equilibrium price. Given the proposed interpretation of T as "depth of reasoning", we have that deeper reasoning maps to a smaller error relative to the REE benchmark.

This result may not be surprising. It is, however, important to recognize that it relies on the property that the net demand for the afternoon good is decreasing in its price. Without this property, the distance between the conjecture $\hat{p}^*$ and its REE counterpart would actually increase without bound as T increases.[21]

Furthermore, note that the result depends on treating the rounds of the Tâtonnement process as a continuous rather than a discrete variable in the following regard. Under the adopted specification, the following two properties hold: first, by varying T in $\mathbb{R}_+$, we can let $\hat{p}^*$ span the entire interval between $\bar{p}^*_{old}$ and $\bar{p}^*_{new}$; second, $\hat{p}^*$ can never fall outside this interval; and third, $\hat{p}^*$ converges monotonically to $\bar{p}^*_{new}$ as $T \to \infty$. If instead we had let t be discrete, these properties would not be true—unless we were to modify the Tâtonnement process so as to let the speed of adjustment itself very with t.[22]

Consider now the behavior that obtains on the basis of the aforementioned conjecture. Combining Lemma 2 with condition (17), we obtain the following characterization of the morning outcomes as functions of the local fundamentals and the conjecture $\hat{p}^*$.

Because the demand and supply functions for morning goods remain the same as those in the original economy, the following is trivially true.

Lemma 6. *For all* m,

$$q_m = Q\left(\hat{p}^*, \theta_m\right) \qquad \text{and} \qquad p_m = P\left(\hat{p}^*, \theta_m\right), \tag{20}$$

where P *and* Q *are the same functions as those obtained in Lemma 4.*

[19] This interpretation may suggest a connection to the concept of level-k reasoning in games; we explore this connection is Section 6.

[20] The following qualification is due. Recall our earlier remark about the possibility of allowing for innovations in the fundamentals to occur between the morning and the afternoon. In such an extension, $\mathcal{P}(\bar{\theta})$ gives, not the average price that gets realized in the afternoon, but rather its rational expectation in the morning. With this in mind, it is best to interpret $\bar{p}^*_{old}$ and $\bar{p}^*_{new}$ as the frictionless values of the morning *belief* of $\bar{p}^*$, as opposed of the realized $\bar{p}^*$.

[21] In the literature, this issue is known as "Tâtonnement stability"; see, e.g., Proposition 17.H.1 in Mas-Collell, Whinston, and Green (1995).

[22] By this we mean letting $\hat{P}^*(t+1) - \hat{P}^*(t) = b(t)\mathcal{N}(\hat{P}^*(t), \bar{\theta}_{new})$ for an appropriately chosen $b : \mathbb{N} \to \mathbb{R}_+$.

That is, the morning outcomes are determined in the same fashion as in our frictionless REE benchmark, except for the following change: the rational expectation $\mathbb{E}[\bar{p}^*]$ has been replaced with the Tâtonnement conjecture $\hat{p}^*$. By direct implication, the aggregate output of morning goods can be expressed as $\bar{q} \equiv \int q_m dm = Q\left(\hat{p}^*, \bar{\theta}\right)$; and because of the linearity of Q, the change in $\bar{q}$ triggered by the aggregate shock can be expressed as

$$\Delta\bar{q} \quad = \quad \frac{\partial Q}{\partial \theta}\Delta\bar{\theta} + \frac{\partial Q}{\partial p^*}\left(\hat{p}^* - \bar{p}^*_{old}\right) \tag{21}$$

Using this result together with Lemma 5, we reach at the following characterization of the response of the aggregate quantity to the aggregate shock.

Proposition 3 (Tatonnement)**.** *For every T, there exists a scalar $\epsilon_{T\hat{a}t}(T)$ such that, for any realization $\Delta\bar{\theta}$ of the aggregate shock, the corresponding change in $\bar{q}$ that obtains along the Tâtonnement(T) solution is given by*

$$\Delta\bar{q} = \epsilon_{T\hat{a}t}(T)\Delta\bar{\theta}.$$

Furthermore,

$$\epsilon_{T\hat{a}t}(T) = \epsilon^{micro} + w(T)\left(\epsilon^{Macro} - \epsilon^{micro}\right), \tag{22}$$

where ϵ^{micro} and ϵ^{Macro} are the same objects as in Section 3 and where w is the same function as the one in Lemma 5 (hence, w is continuous and strictly increasing in T, with $w(0) = 0$ and $w(\infty) \equiv \lim_{T\to\infty} w(T) = 1$).

In short, $\epsilon_{T\hat{a}t}(T)$ identifies the macro elasticity of the Tâtonnement economy in which the depth of reasoning is T. This elasticity is arbitrarily close to the underlying *micro* elasticity when T is low enough,[23] but gets closer and closer to the *macro* elasticity of the frictionless REE benchmark as T increases. This reflects the fact that, by design, the Tâtonnement process helps arrest the GE adjustment that is present in that benchmark. Formally, using Lemma 5, the GE effect of the modified economy can itself be expressed as

$$\mathtt{GE}_{T\hat{a}t}(T) \equiv \frac{\partial Q}{\partial p^*}\frac{\hat{p}^* - \bar{p}^*_{old}}{\Delta\bar{\theta}} = w(T)\mathtt{GE},$$

where $\mathtt{GE} \equiv \frac{\partial Q}{\partial p^*}\frac{\partial \mathcal{P}}{\partial \theta}\left(= \epsilon^{Macro} - \epsilon^{micro}\right)$ is the GE effect of the frictionless benchmark and where $w(T)$ is the aforementioned function. Accordingly, the fact that $w(T) \in (0, 1)$ for all $T > 0$ formalizes the notion that GE adjustment is "incomplete", or that it is "weakened" by the considered relaxation of the solution concept; and the fact that $w(T)$ increases with T can be interpreted as the property that a larger GE adjustment requires "deeper reasoning".

We illustrate these points in Figure 3 for the same kind of demand shock as the one considered in Figure 2. We consider the scenario in which the GE effect happens to amplify the PE effect; the same logic applies in the opposite scenario as well. As explained before, the macro response of the original economy is represented by the shift from point X to point Z and this shift can be decomposed to a PE adjustment, captured by the shift from X to Y, and a GE adjustment, captured by the shift from Y to Z. In the original economy, the latter shift is complete and instantaneous. Under the present modification, it is partial: the economy moves to an intermediate point along the segment between Y and Z. The larger T, the close to Z the economy gets.

Remark. By treating the Tâtonnement dynamics as a cognitive process that takes place within a fixed moment in time (namely, within the "morning" of our two-period economy), we have formalized the notion of "incomplete" GE adjustment while abstracting from dynamics. In Section 8, we describe a multi-period extension in which additional rounds require more time. Under this interpretation, the economy moves slowly from Y to Z as time passes, helping accommodate the complementary notion that "GE adjustment takes time".

[23] The micro elasticity is the same in the two economies, for they both impose individual rationality and market clearing at the local level.

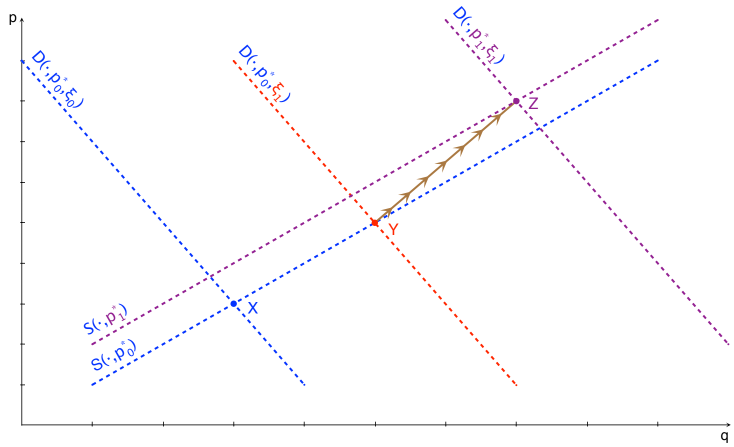

Figure 3: Tâtonnement

5 Removing Common Knowledge

The preceding analysis offered a simple, and logically coherent, formalization of the notion that GE adjustment may be weaker than what predicted by standard modeling practice. This formalization, however, is borrowed from old ideas about "off equilibrium" adjustment, which find little place in modern methodology.[24] It also requires a violation of rational expectations, which raises the usual delicate questions.

We now show how informational frictions, and in particular lack of common knowledge of the aggregate shocks, can help deliver the same take-home message without the aforementioned caveats. That is, we show how the notion of "incomplete GE adjustment" can be recast in a different modification of the frictionless benchmark, which lets agents form rational expectations and play equilibrium but modifies the information structure upon which the relevant expectations are based. The key friction is no more a departure from the REE concept, but rather a certain kind of coordination friction that operates *in* equilibrium: it is the uncertainty the each agent faces about what other agents know and, thereby, about how they act along the equilibrium.

5.1 Set up

The economy considered in this section shares the same primitives (preferences and technologies) as either the original economy studied in Section 3 or the Tâtonnement variant studied in Section 4. Unlike the latter case, there is no departure from REE. Unlike the former case, the agents do not have common knowledge of the aggregate shock to the underlying fundamentals; instead, they have dispersed, private, information about it.

Assumption 8. *For each marketplace m, the information that the agents have in the morning about all aggregate variables is summarized in a sufficient statistic s_m given by*

$$s_m = \Delta\bar{\theta} + v_m,$$

where $\Delta\bar{\theta}$ is the underlying change in the fundamentals, itself modeled as a random variable drawn from $N(0, \sigma_\theta^2)$, and v_m is an idiosyncratic noise term, drawn from $N(0, \sigma_v^2)$, i.i.d. across markets, and independent of $\Delta\bar{\theta}$.

[24] This statement needs to be qualified in the following regard: Level-k Thinking and the related solution concepts of Evans and Ramey (1992) and Garcıa-Schmidt and Woodford (2015).

Note that, because the aforementioned signal differs across marketplaces, the agents in one marketplace do not know what the agents in other marketplaces know. In other words, agents face, not only uncertainty about the underlying $\Delta\bar{\theta}$, but also uncertainty about what others believe, about what others belief that others believe, and so on. As it will become clear, this kind of *higher-order* uncertainty causes GE attenuation.

Also note that the observation of the local fundamental contains information about the aggregate shock $\Delta\bar{\theta}$. The signal s_m described above is meant to be a sufficient statistic for *all* the information that marketplace m possesses in the morning about $\bar{\theta}$, and therefore encompasses any information contained in the observation of the local θ_m.

Finally note that the variation in the aggregate fundamental is payoff-irrelevant for any given agent as long as we fix her *own* fundamental and the price at which she trades in the afternoon. Removing common knowledge of $\Delta\bar{\theta}$ therefore serves one and only one modeling role: it influences the expectations that agents form during the morning about the concurrent activity in other markets and the prices that will be determined in the afternoon.

We henceforth denote with $E_m[\cdot]$ the rational expectation that agents in marketplace m hold in the morning. We also parameterize the information structure by the scalar

$$\lambda \equiv \frac{1}{1+(\sigma_v^2/\sigma_\theta^2)} \in (0,1). \tag{23}$$

We finally define the solution concept as follows.

Definition 2. *Fix any $\lambda \in (0,1)$. The incomplete-information solution is given by the REE of the variant economy in which the information structure is described by Assumption 8 together with condition (23).*

In the rest of this section, we characterize the incomplete-information equilibrium and spell out its implications regarding the response of aggregate outcomes to aggregate shocks. Before doing that, however, is worth making the following point regarding the interpretation of λ. By construction, the scalar λ captures the signal-to-noise ratio in the available information about the underlying aggregate shock. But because the information is private, this scalar also controls the degree of common knowledge.

To see what we mean by this, let $\bar{E}^h[.]$ denote that h-th order average belief operator, defined recursively by $\bar{E}^1[\cdot] \equiv \bar{E}[\cdot] \equiv \int E[\cdot]dm$ and $\bar{E}^h[\cdot] \equiv \int E[\bar{E}^{h-1}[\cdot]]dm$ for all $h \geq 2$. Under the assumed information structure, the first- and higher-order beliefs satisfy, for all $h \geq 1$, $\bar{E}^h[\Delta\bar{\theta}] = \lambda^h \Delta\bar{\theta}$; equivalently,

$$\bar{E}^h[\bar{\theta}] = \bar{\theta}_{old} + \lambda^h \Delta\bar{\theta}. \tag{24}$$

It follows that λ controls, not only how precise the available information about the underlying shock is (equivalently, how close the first-order beliefs of $\bar{\theta}$ are to the true realized value of $\bar{\theta}_{new}$), but also how anchored higher-order beliefs are to the common prior (which is herein centered around $\bar{\theta}_{old}$). As $\lambda \to 1$, the friction vanishes: the underlying value of $\bar{\theta}$ becomes, not only perfectly known to each market, but also common knowledge; that is, all markets know that all markets know $\bar{\theta}$, all markets know that all markets know that all markets know $\bar{\theta}$, and so on. At the other extreme, as $\lambda \to 0$, the friction is maximal. Finally, for any $\lambda \in (0,1)$, we have that $\bar{E}^h[\bar{\theta}] \to \bar{\theta}_{old}$ as $h \to \infty$ regardless of how close λ is to 1. That is, beliefs of sufficiently high order remain arbitrarily close to the pre-shock situation even when the informational friction, as measured by the distance of λ from 1, is arbitrarily small.

Why is this point important? Because, as we show next, the rational expectations that agents form about future prices turn out to depend, not only on their first-order beliefs of $\bar{\theta}$, but also, and quite importantly, on their higher-order beliefs of $\bar{\theta}$. It is through this channel that lack of common knowledge attenuates the GE effects of the frictionless benchmark, thus also reducing the gap between the micro and macro elasticities of interest.

5.2 Equilibrium Characterization

Let $E_m[\bar{p}^*]$ denote the rational expectation that agents in marketplace m form about $\bar{p}^*$ in the morning, based on the local knowledge of s_i.

Lemma 7. *For all m, the local outcomes satisfy*

$$q_m = Q\left(E_m[\bar{p}^*], \theta_m\right) \qquad \textit{and} \qquad p_m = P\left(E_m[\bar{p}^*], \theta_m\right), \tag{25}$$

where Q and P are the same functions as those obtained in Lemma 4.

That is, the morning outcomes are determined in the same fashion as in our frictionless benchmark, except that the *common* rational expectation $\mathbb{E}[\bar{p}^*]$ is replaced by the *local* rational expectation $E_i[\bar{p}^*]$. At first glance, this may appear to be an innocuous change. This is indeed the case insofar as one is concerned with the *partial-equilibrium* predictions of the theory: as we move between the frictionless benchmark and the present variant, we vary the information upon which the rational expectations of $\bar{p}^*$ are formed, but do not vary how q_m and p_m respond either to these expectations or to the local fundamentals. By the same token, the two economies feature the same *micro* elasticity, that is, the same sensitivity of local activity to local shocks. Yet, as we show next, the two economies make different general-equilibrium predictions and, as a result, feature different *macro* elasticities.

The aggregate quantity of the morning goods can now be expressed as follows:

$$\bar{q} \equiv \int q_m dm = Q(\bar{E}[\bar{p}^*], \bar{\theta}) \tag{26}$$

where $\bar{E}[\bar{p}^*] \equiv \int E_m[\bar{p}^*]\, dm$ is the average belief of $\bar{p}^*$ in the cross-section of markets. To characterize how $\bar{q}$ responds to the change in the fundamentals, we therefore need to characterize how this average belief responds. And because agents have rational expectations, this means that we need to characterize the fixed-point relation between the average belief $\bar{E}[\bar{p}^*]$ and the actual $\bar{p}^*$. This is similar to our frictionless benchmark, except for the following difference: because agents do not share the same information about aggregate economic conditions, this fixed-point relation turns out to be more "delicate" than before.

Let us elaborate. First, consider how demand and supply are determined in the afternoon markets. At this stage, the quantities of the morning goods are predetermined. By Lemma 3, the afternoon prices satisfy

$$\bar{p}^* = P^*(\bar{q}, \bar{\theta}), \tag{27}$$

where P^* is defined as in condition (8). Replacing $\bar{q}$ by (26), we get the following fixed-point relation between the realized value of $\bar{p}^*$ and the average expectation of it in the morning:

$$\bar{p}^* = \mathcal{T}\left(\bar{E}[\bar{p}^*], \bar{\theta}\right), \tag{28}$$

where $\mathcal{T}$ is the same mapping as in condition (12). Finally, taking expectations on both sides, we reach the following result, which is robust to the introduction of unpredictable shocks to the fundamentals between the morning and the afternoon.[25]

[25] This echoes the related point we made in the frictionless benchmark. In the presence of the aforementioned kind of shocks, condition (28) has to be replaced with $\bar{p}^* = \mathcal{T}\left(\bar{E}[\bar{p}^*], \bar{\theta}\right) + \varepsilon$, where ε captures the effect of these shocks. But since these shocks are unpredictable in the morning, Lemma 8 remains valid.

Lemma 8. *The following fixed-point relation holds in equilibrium:*

$$\bar{E}[\bar{p}^*] = \mathcal{T}\left(\bar{E}^2[\bar{p}^*], \bar{E}[\bar{\theta}]\right). \tag{29}$$

We can now see how the lack of common knowledge changes the nature of the fixed-point problem that pins down the equilibrium expectations. Contrast the above condition to condition (12) obtained in the frictionless benchmark. In that benchmark, the higher-order beliefs of $\bar{p}^*$ coincide with the corresponding first-order beliefs, because all agents shared the same information. In addition, $\bar{\theta}$ is commonly known. That benchmark can therefore be nested in condition (29) above by letting $\bar{E}^2[\bar{p}^*] = \bar{E}[\bar{p}^*]$ and $\bar{E}[\bar{\theta}] = \bar{\theta}$, from which we can then solve for $\bar{E}[\bar{p}^*]$ as function of $\bar{\theta}$, namely for $\bar{E}[\bar{p}^*] = \mathcal{P}(\bar{\theta})$. In short, the combination of the REE concept with the assumption of common knowledge forces higher-order beliefs of prices to coincide with first-order beliefs of prices and ultimately pins down the later as functions of the fundamentals.

By contrast, once we depart from common knowledge, the higher-order beliefs of prices can diverge from first-order beliefs. Furthermore, because *realized* prices depend on fundamentals and first-order beliefs of prices, the latter depend on first-order beliefs of fundamentals and *second-order* beliefs of prices, which is what condition (29) states. Iterating this condition forward, and using the fact that $\mathcal{T}$ is a contraction mapping, we obtain the following characterization of the equilibrium value of $\bar{E}\left[\bar{p}^*\right]$.

Corollary 2. *In equilibrium, the average rational expectation of $\bar{p}^*$ is determined by the hierarchy of beliefs about the underlying fundamentals:*

$$\bar{E}\left[\bar{p}^*\right] = \gamma(1-\alpha)\sum_{h=1}^{\infty}\alpha^{h-1}\bar{E}^h\left[\bar{\theta}\right], \tag{30}$$

where $\alpha \in (-1,+1)$ is defined as in condition 13 and where $\gamma \equiv \frac{\partial \mathcal{P}}{\partial \theta}$ is the slope of $\bar{E}[\bar{p}^*]$ with respect to $\bar{\theta}$ in the frictionless benchmark.

Modeling the *equilibrium expectations* of future prices is therefore equivalent to modeling the *hierarchy of beliefs* of the underlying fundamentals. By assuming common knowledge, the standard practice forces higher-order beliefs to collapse to first-order beliefs. This may be convenient, but it is neither realistic nor innocuous: it imposes a very tight structure on how the rational expectations of future prices (or of any other endogenous outcome) respond to aggregate shocks. We next show that relaxing this tight structure by allowing for lack of common knowledge gives rise to a similar GE attenuation effect as the Tâtonnement modification studied in the previous section.

5.3 Lack of Common Knowledge as GE Attenuation

By combining Corollary 2 with our earlier characterization of the hierarchy of beliefs, we reach the following characterization of the equilibrium value of $\bar{E}\left[\bar{p}^*\right]$.

Lemma 9. *In equilibrium,*

$$\bar{E}\left[\bar{p}^*\right] = \bar{p}^*_{old} + \pi(\lambda)\left(\bar{p}^*_{new} - \bar{p}^*_{old}\right) \tag{31}$$

where $\bar{p}^*_{old} \equiv \mathcal{P}(\bar{\theta}_{old})$, $\bar{p}^*_{new} \equiv \mathcal{P}(\bar{\theta}_{new})$, and the function π is continuous and strictly increasing in λ, with $\pi(0) = 0$ and $\pi(1) = 1$.

By varying λ, we can thus let the post-shock expectation in the modified economy take any value in the range between the corresponding pre- and post-shock values in the frictionless benchmark. Using the fact that $\bar{q} = Q(\bar{E}[\bar{p}^*], \bar{\theta})$, we then reach the following result, which formalizes the sense in which lack of common knowledge bridges the gap between the micro and the macro elasticities of interest.

Proposition 4 (Incomplete Info)**.** *For any λ, there exists a scalar $\epsilon_{Inc}(\lambda)$ such that, for any realization $\Delta\bar{\theta}$ of the aggregate shock, the corresponding change in $\bar{q}$ is given by*

$$\Delta\bar{q} = \epsilon_{Inc}(\lambda)\Delta\bar{\theta}$$

Furthermore,

$$\epsilon_{Inc}(\lambda) = \epsilon^{micro} + \pi(\lambda)\left(\epsilon^{Macro} - \epsilon^{macro}\right), \tag{32}$$

where ϵ^{micro} and ϵ^{Macro} are the same objects as those found in Proposition 2 and π is the same function as that found in Lemma 9.

Recall that π is continuous and strictly increasing in λ, with $\pi(0) = 0$ and $\pi(1) = 1$. By varying λ between 0 and 1, we can thus span all the values between ϵ^{micro} and ϵ^{Macro} . For λ close to zero (meaning a sufficiently large departure from common knowledge), the macro elasticity of the modified model is arbitrarily close to the *micro* elasticity of the original model. But as λ increases (meaning a higher degree of common knowledge), the *macro* elasticity of the modified model gets closer and closer to the *macro* elasticity of the original model. Importantly, all these properties hold true no matter whether ϵ^{Macro} is higher or lower than ϵ^{micro}. We conclude that varying the degree of common knowledge is akin to varying the extent to which the GE effect is active, regardless of whether this effect amplifies or offsets the underlying PE effect.

Remark. In our setting, the necessary informational friction is obtained by assuming that markets are segmented and by preventing agents from observing the prices and/or quantities in markets they do not currently participate in. While these assumptions are certainly stark, it is important to keep in mind the following points. First, even if perfectly precise and perfectly public information about the outcomes of all marketplaces is readily available, lack of common knowledge about these objects can be the product of "rational inattention" along the lines of Sims (2003); Mackowiak and Wiederholt (2009), Mackowiak and Wiederholt (2009) and Myatt and Wallace (2012). And second, even if all agents were able to reach common knowledge of concurrent market outcomes, they would still not be able to reach common knowledge of the underlying fundamentals and/or of future outcomes unless they also had common knowledge of the fact that everybody behaves according to the REE concept. In our view, these points reinforce the plausibility of the assumed informational friction.

5.4 A Game-Theoretic Representation

We have already alluded to a game-theoretic interpretation of our results. We make this interpretation tight with the help of the following result, which turns out to be useful for additional purposes as well.

Proposition 5. *(i) There exists a linear function $\mathcal{BR}$ such that, for all realizations of uncertainty and all m, the equilibrium value of q_i satisfies*

$$q_m = \mathcal{BR}\left(\theta_m, E_m[\bar{\theta}], E_m[\bar{q}]\right). \tag{33}$$

(ii) If $\bar{q} = \mathcal{BR}\left(\bar{\theta}, \bar{E}[\bar{\theta}], \bar{E}[\bar{q}]\right)$ and $\bar{p}^* = P^*(\bar{q}, \bar{\theta})$, then $\bar{p}^* = \mathcal{T}(\bar{E}[\bar{p}^*], \bar{\theta})$; and conversely, if $\bar{p}^* = \mathcal{T}(\bar{E}[\bar{p}^*], \bar{\theta})$ and $\bar{q} = Q(\bar{E}[\bar{p}^*], \bar{\theta})$, then $\bar{q} = \mathcal{BR}\left(\bar{\theta}, \bar{E}[\bar{\theta}], \bar{E}[\bar{q}]\right)$.

(iii) The slope of $\mathcal{BR}$ with respect to $E_m[\bar{q}]$ equals α, the slope of $\mathcal{T}$.

Part (i) states that the quantity q_m produced in each market can be expressed as function of θ_m (which is locally known), of the local expectation of $\bar{\theta}$, and of local expectation of $\bar{q}$. Part (ii) states that looking for the fixed point of these best responses functions is equivalent to looking for the fixed point of $\mathcal{T}$, the function that pins down the rational expectations of the relevant prices. Together, these results mean that we can recast the incomplete-information solution of our dynamic Walrasian economy as the unique Bayesian-Nash Equilibrium of a fictitious static game in which the players are the markets, their actions are the local quantities, and their best response functions are given by (33). Complementing this interpretation, part (iii) ties α, the slope of the function $\mathcal{T}$ whose fixed point pins down the rational expectations of the relevant price, to the degree of strategic complementarity (if $\alpha > 0$) or substitutability (if $\alpha < 0$) in the aforementioned game.

This game is similar to the class of linear-quadratic beauty-contest games considered in, inter alia, Morris and Shin (2002); Angeletos and Pavan (2007) and Bergemann and Morris (2013). The only twist is that the best response depends on both the beliefs and the actual realization of the underlying fundamentals.[26] The restriction that $\alpha \in (-1, 1)$ guarantees that this game admits a unique rationalizable outcome. Finally, the size of the degree of strategic complementarity/substitutability in this beauty contest can be connected to the size of the GE effect in our Walrasian economy.

This connection is tightest when $\bar{p}^*$ depends on $\bar{\theta}$ only through $\bar{q}$ (i.e., when $\partial P^*/\partial\theta = 0$). Under this restriction, it is easy to show that

$$\epsilon^{Macro} = \frac{1}{1-\alpha}\epsilon^{micro}. \tag{34}$$

The macro and micro elasticities therefore share the same sign regardless of α, but their relative magnitude depends on α. When $\alpha < 0$, the fictitious game features strategic substitutability, the GE effect offsets the PE effect, and our attenuation effect translates to $|\epsilon^{Macro}| < |\epsilon^{Inco}| < |\epsilon^{micro}|$. When instead $\alpha > 0$, the fictitious game features strategic complementarity, the GE effect amplifies the PE effect, and $|\epsilon^{Macro}| > |\epsilon^{Inco}| > |\epsilon^{micro}|$. Last but not least, the following is true.

Proposition 6. *Suppose $\partial P^*/\partial\theta = 0$ and fix any $\lambda \in (0, 1)$. A higher $|\alpha|$ drives both $\frac{\epsilon^{Macro}}{\epsilon^{micro}}$ and $\frac{\epsilon^{Inco}}{\epsilon^{Macro}}$ further away from 1.[27] In this sense, the same primitives that enhance the GE effect in the frictionless benchmark also strengthen the attenuation of this effect in the considered variant.*

A higher $|\alpha|$ enhances the GE effect in the frictionless benchmark is evident from (34) and is in line with a long tradition that equates "general-equilibrium multipliers" in macroeconomic models with strategic complementarity in games. What we are adding here is the following key observation: the larger such multipliers are under the conventional common knowledge assumption, the stronger their attenuation once this assumption is relaxed.

Also note that the result applies even when the environment features strategic substitutability, which means that the GE effects offset, rather than amplify, PE effects. We conclude that, regardless of the relation between ϵ^{Macro} and ϵ^{micro} , the larger that gap between the two is under standard practice, the larger the fraction of this gap that gets "erased" once one allows for a realistic information friction, namely for lack of common knowledge of the underlying aggregate shock.

5.5 Relation to Tâtonnement

Combining Propositions 3 and 4, we have the following equivalence between two modifications of the frictionless benchmark we have studied so far.

Proposition 7. *For any Tâtonnement economy with depth $T \in (0, \infty)$, there exists an incomplete-information economy with common-knowledge degree $\lambda \in (0, 1)$ such that, for any realization of $\Delta\bar{\theta}$, the average rational expectation $\bar{E}[\bar{p}^*]$ in the latter coincides with the conjecture $\hat{p}^*$ in the former, and the two economies predict the same observable change in $\bar{q}$. The converse is also true.*

[26] To put it differently, the payoff type of player m in the relevant game is given by the pair $(\theta_m, E_m[\bar{\theta}])$.

[27] Formally, when $\alpha > 0$, $\frac{\epsilon^{Macro}}{\epsilon^{micro}}$ is higher than 1 and increases with α, whereas $\frac{\epsilon^{Inc}}{\epsilon^{Macro}}$ is lower than 1 and decreases with α; and when $\alpha < 0$, $\frac{\epsilon^{Macro}}{\epsilon^{micro}}$ is lower than 1 and decreases with $|\alpha|$, whereas $\frac{\epsilon^{Inc}}{\epsilon^{Macro}}$ is higher than one and increases with $|\alpha|$.

This explains the sense in which relaxing common knowledge offers a modern reincarnation of the ideas that motivated the older literature on Tâtonnement dynamics: what is an "off-equilibrium" adjustment in that literature has been recast as an "on-equilibrium" phenomenon; and what is a departure from rational expectations has been recast as the symptom of the imperfect coordination along the unique rational-expectations equilibrium.

This equivalence, however, pegs the following question: is there a reason to prefer the one approach over the other? The answer to this question depends largely on how much one values Lucas' critique, that is, on whether one wishes to adhere to the principle that our interpretations of the data and our policy prescriptions should not be based on the assumption that agents make systematic mistakes. On our part, we prefer the second modification over the first, not only because it is immune to Lucas' critique, but also because of the lesson provided at the end of the previous subsection: the same forces that create a gap between ϵ^{Macro} and ϵ^{micro} in standard models also help reduce this gap once we relax common knowledge. This lesson could *not* have been obtained in the Tâtonnement variant because it hinges on the interaction of rational expectations with incomplete information, and on recasting the relevant general-equilibrium effects in terms of higher-order beliefs.

6 Variants: Cobweb, Level-k, and Reflective Equilibrium

So far, we have shown that (i) dropping the assumption that the aggregate shock is common knowledge predicts a reduction of the gap between micro and macro elasticities relative to the frictionless benchmark and (ii) this prediction is shared by a *specific* relaxation of the REE concept, one that built on Tâtonnement. Clearly, this prediction need not be shared by *every* relaxation of the REE concept.

For instance, suppose that one models "bounded rationality" with ϵ-equilibrium (also know as near-Nash equilibrium). In games, this concept requires that the action of each player is "nearly rational" in the sense that it delivers a payoff that is within ϵ of the payoff delivered by the best-response action, where $\epsilon > 0$ is an exogenous scalar that can be thought of as the degree of bounded rationality. Adapting this concept to our Walrasian setting boils down to letting the "nearly rational" demand and supply functions vary around the "fully rational" ones defined in Section 2. As a result, the aggregate change, $\Delta\bar{q}$, triggered by any given shock can also vary around its frictionless counterpart. It follows the proposed relaxation of the REE concept does not share the aforementioned prediction: the gap between the micro and macro elasticities can be *larger* in the modified economy than in the frictionless benchmark.

This underscores that the approach we favor in this paper—namely, dropping common knowledge of aggregate shocks while maintaining the REE concept—imposes a *specific* structure on the departure obtained from the frictionless benchmark. This structure is precisely that the relevant GE effects *have* to be attenuated and, equivalently, that the gap between the micro and macro elasticities *have* to be decreased.

With this point in mind, we now explore whether this structure is shared by three other possible relaxations of the REE concept, which build on the following concepts from the literature:

1. Cobweb dynamics
2. Level-k thinking as in Nagel (1995) and Stahl and Wilson (1994); Stahl Dale and Wilson (1995)
3. Reflective equilibrium as in Garcıa-Schmidt and Woodford (2015)

In all these variants, we preserve the demand and supply system of the frictionless benchmark, as well as its partial-equilibrium predictions. We nevertheless modify the general-equilibrium adjustment to aggregate shocks by letting the agents act on the basis of *irrational* conjectures about how certain aggregate outcomes, namely $\bar{p}^*$ and/or $\bar{q}$, react to these shocks. In this regard, the new variants are similar to the Tâtonnement variant studied in Section 4. The difference is in the exact cognitive process that pins down the relevant price and/or quantity conjectures.

We first show that the variants that are based on Cobweb and Level-k Thinking are tightly connected with one another, but not necessarily with the Tâtonnement and the incomplete-information variants we studied before: in certain cases, the new variants predict that the relevant price or quantity conjectures can *overshoot* relative to the frictionless benchmark, which means that the GE effects get amplified instead of being attenuated. We next show that, although the variant that builds on Garcıa-Schmidt and Woodford (2015) has a similar favor as Cobweb and level-k thinking, it avoids the aforementioned "overshooting" problem and ends up delivering similar predictions as our Tâtonnement and incomplete-information variants.

On the basis of these findings, we conclude that, although GE attenuation is not a necessary implication of "bounded rationality" (in the sense of some of the considered relaxations of the REE concept), it is the necessary implication of "imperfect coordination" (in the sense of removing common knowledge of aggregate shocks and aggregate outcomes). In the next section, we then discuss why the latter approach offers a more structured—and, in our view, more appealing—approach that the alternatives *even when* these alternatives share the prediction that GE effects get attenuated.

6.1 Cobweb

A familiar alternative to the notion of Tâtonnement dynamics is that of Cobweb dynamics. Similarly to how we treated Tâtonnement in Section 4, here we recast Cobweb as an *instantaneous* cognitive process, whose outcome is a price conjecture $\hat{p}^*$ upon which the agents act.

Let k index the "depth" of the process; for reasons that will become clear shortly, we are now forced treat k as a discrete variable, namely $k \in \mathbb{N}$. Next, let $\mathcal{T}^k$ be defined recursively by $\mathcal{T}^0(p^*,\theta) \equiv p^*$ and $\mathcal{T}^k(p^*,\theta) \equiv \mathcal{T}\left(\mathcal{T}^{k-1}(p^*,\theta),\theta\right)$ for all (p^*,θ) and all $k \geq 1$. We can then define our cognitive version of Cobweb dynamics as follows.

Definition 3. *Fix a $k \in \{0,1,2...\}$. The Cobweb(k) solution is given by a conjecture $\hat{p}^*$ and a pair (q_m,p_m) of each m such that:*

(i) For all m, the pair (q_m,p_m) is consistent with household optimality, firm optimality, and market clearing in the local market for morning goods, under the following conjectures for the afternoon prices:

$$\hat{E}_m[p^*_{m'}] = \rho^2 P^*\left(q_m,\theta_m\right) + \left(1-\rho^2\right)\hat{E}_m[\bar{p}^*] \qquad \textit{and} \qquad \hat{E}_m[\bar{p}^*] = \hat{p}^* \tag{35}$$

(ii) $\hat{p}^*$ is is given by

$$\hat{p}^* = \mathcal{T}^k(\hat{p}^*_0,\bar{\theta}_{new}), \tag{36}$$

where $\hat{p}^*_0 = \bar{p}^*_{old} \equiv \mathcal{P}(\bar{\theta}_{old})$.

Clearly, the only difference between the above definition and Definition 1 (our version of Tâtonnement) is the condition that pins down the conjecture $\hat{p}^*$, namely condition (5) above. When $k=0$, this condition gives $\hat{p}^* = \hat{p}^*_0 = \bar{p}^*_{old}$, meaning that agents behave under the conjecture that the average afternoon price will stay at its the pre-shock equilibrium level. Consider next $k=1$. Now any given agent realizes that, if other agents behave under the aforementioned conjecture, the aggregate quantity in the morning will be $\bar{q} = \bar{q}_0 \equiv Q(\hat{p}^*_0,\bar{\theta}_{new})$ and the afternoon markets will therefore clear with $\bar{p}^* = P^*\left(\bar{q}_0,\bar{\theta}_{new}\right) = \mathcal{T}(\hat{p}^*_0,\bar{\theta}_{new})$. On the basis of this argument, the initial conjecture is updated from $\hat{p} = \hat{p}^*_0$ to $\hat{p} = \hat{p}^*_1 \equiv \mathcal{T}(\hat{p}^*_0,\bar{\theta}_{new})$. By induction, the conjecture at an arbitrary round k is given by $\hat{p}^* = \hat{p}^*_k \equiv \mathcal{T}(\hat{p}^*_{k-1},\bar{\theta}_{new}) = \mathcal{T}^k(\hat{p}^*_0,\bar{\theta}_{new})$.

Similarly to what we did in the Tâtonnement variant, we can interpret k as the "depth" of the assumed cognitive process. As evident in condition (5), considering higher and higher k maps to iterating more and more times on $\mathcal{T}$. Because $\mathcal{T}$ is a contraction mapping, we know that $\hat{p}^*_k \to \bar{p}^*_{new}$ as $k \to \infty$. That is, the Cobweb process shares with the Tâtonnement process that the price conjecture converges to the post-shock equilibrium price as the "depth of reasoning" increases without bound. Unlike the Tâtonnement variant, however, this convergence need not be monotonic and the price conjecture may fall outside the range between the pre- and the post-shock equilibrium price. This is inherent to the bang-bang manner with which the price conjecture is adjusted at each round of Cobweb procedure.

Lemma 10. *There exists a sequence $\{g_k\}$, with $g_0 = 0$ and $\lim_{k\to\infty} g_k = 1$, such that the following properties hold:*

(i) For any k and any $\Delta\bar{\theta}$, the conjecture $\hat{p}_k^*$ satisfies

$$\hat{p}_k^* = \bar{p}_{old}^* + g_k \left(\bar{p}_{new}^* - \bar{p}_{old}^*\right).$$

(ii) If $\alpha > 0$, the sequence is strictly increasing and bounded between 0 and 1.

(iii) If instead $\alpha < 0$, this sequence is non-monotone, with $g_k < 1$ whenever k is even and $g_k > 1$ whenever k is odd.

The price conjecture therefore falls inside the interval between $\bar{p}_{old}^*$ and $\bar{p}_{new}^*$ (the pre-shock and the post-shock perfect-foresight prices) insofar as the economy happens to feature $\alpha > 0$. But if the economy features $\alpha < 0$, the price conjecture can fall outside this interval. By the same token, the Cobweb economy can "overshoot" in the sense of delivering a macro elasticity that exceeds (in absolute value) the one obtained in the frictionless benchmark.

Proposition 8 (Cobweb)**.** *For any k, there exists a scalar $\epsilon_{Cob}(k)$ such that, for any realization $\Delta\bar{\theta}$ of the aggregate shock, the corresponding change in the value of $\bar{q}$ that obtains in the Cobweb(k) solution is given by*

$$\Delta\bar{q} = \epsilon_{Cob}(k)\Delta\bar{\theta}$$

Furthermore,

$$\epsilon_{Cob}(k) = \epsilon^{micro} + g_k \left(\epsilon^{Macro} - \epsilon^{macro}\right), \tag{37}$$

where g_k is the same as in Lemma 10, satisfying $g_k > 1$ if $\alpha < 0$ and k is odd and $g_k \in [0, 1)$ otherwise.

We conclude that Cobweb is similar to Tâtonnement and lack of common knowledge in economies in which GE effect complements the PE effect ($\alpha > 0$), but not in economies in which the GE effect offsets the PE effect ($\alpha < 0$): in the latter class of economies, Cobweb opens the door to amplification of the GE effect.

6.2 Level-k Thinking

Level-k Thinking—also known as Limited-Depth Thinking—is a solution concept often used in the experimental literature, but also elsewhere. See Nagel (1995) and Stahl and Wilson (1994); Stahl Dale and Wilson (1995) for early contributions; Crawford, Costa-Gomes, and Iriberri (2013) for a survey; Garcıa-Schmidt and Woodford (2015) for a discussion of closely-related concepts in macroeconomics; and Farhi and Werning (2017) for an application in the context of forward guidance.

According to this concept, level-0 thinkers best-respond to the belief that other players' strategies are fixed at some "default" or "steady-state point"; level-1 thinkers best-respond to the belief that other players are level-0 thinkers; and so on. We adapt this concept to our study of GE effects by building on Proposition 5, which permits us to represent the economy as a game in quantities, and by setting the default for the average quantity $\bar{q}$ to its pre-shock equilibrium level. In particular, with the best-response function $\mathcal{BR}$ defined as in Proposition 5, we define $\mathcal{BR}^k$ recursively by $\mathcal{BR}^0(\theta_m, \theta, q) \equiv q$ and $\mathcal{BR}^k(\theta_m, \theta, q) \equiv \mathcal{BR}\left(\theta_m, \theta, \mathcal{BR}^{k-1}(\theta, \theta, q)\right)$ for all (θ_m, θ, q) and all for all $k \geq 1$, and we finally define our solution concept as follows.

Definition 4. *For any $k \in \{0,1,2...\}$, the level-k solution is given by a quantity conjecture $\hat{q}$ and a pair (q_m, p_m) of each i such that:*

(i) The quantity conjecture $\hat{q}$ is given by

$$\hat{q} = \mathcal{BR}^k(\bar{\theta}_{new}, \bar{\theta}_{new}, \bar{q}_{old}), \tag{38}$$

where $\bar{q}_{old} \equiv Q(\bar{p}^*_{old}, \bar{\theta}_{old})$.

(ii) For all i, the pair (q_i, p_i) is consistent with household optimality, firm optimality, and market clearing in the local market for morning goods, under the following conjectures for the afternoon prices:

$$\hat{E}_m[p^*_{m'}] = \rho^2 P^*(q_m, \theta_m) + \left(1-\rho^2\right)\hat{E}_m[\bar{p}^*] \qquad \textit{and} \qquad \hat{E}_m[\bar{p}^*] = P^*(\hat{q}, \bar{\theta}_{new}) \tag{39}$$

Let us explain this definition. When $k = 0$, agents expect the aggregate quantity to stay at its pre-shock equilibrium value; that is, $\hat{q} = \hat{q}_0 \equiv \bar{q}_{old}$. When $k = 1$, agents expect the other agents to act as if $k = 0$ and therefore expect the aggregate quantity to be given by the best response to $\bar{q}_{old}$; that is, $\hat{q} = \hat{q}_1 \equiv \mathcal{BR}(\bar{\theta}_{new}, \bar{\theta}_{new}, \hat{q}_0)$. Similarly, for any $k \geq 1$, agents expect the aggregate quantity to be $\hat{q} = \hat{q}_k \equiv \mathcal{BR}(\bar{\theta}_{new}, \bar{\theta}_{new}, \hat{q}_{k-1})$. This mirrors the definition of level-k thinking in games and explains part (i).[28] Part (ii) then transforms the quantity conjecture to a price conjecture and requires that demand and supply are optimal given this price conjecture: whenever an agent expects the aggregate morning quantity to be $\hat{q}$, she also expects the average afternoon price to be $P^*(\hat{q}, \bar{\theta}_{new})$, and chooses her demand or supply accordingly.

Forming conjectures about the *simultaneous* behavior of the other agents is therefore equivalent to forming conjectures about the resulting *future* prices. This permits us to go back and forth between the game-theoretic and the Walrasian representation of the economy under the level-k concept. Furthermore, because iterating on $\mathcal{BR}$ is equivalent to iterating on $\mathcal{T}$, it is evident that there is tight relation between the level-k and Cobweb solution concepts. This tight relation is formalized in the following proposition.

Proposition 9. *Suppose either that $P^*(\bar{q}, \bar{\theta})$ is invariant to $\bar{\theta}$, or that we modify the Cobweb solution concept so that the initial price conjecture is given by $\hat{p}_0 = P^*(\bar{q}_{old}, \bar{\theta}_{new})$ rather than $\hat{p}_0 = \bar{p}^*_{old} \equiv P^*(\bar{q}_{old}, \bar{\theta}_{old})$. Then, for any $k \in \{0,1,2...\}$, the level-k solution and the Cobweb(k) solution impose the same price conjectures and give rise to the same observables.*

The two concepts are *not* tautologically the same: level-k thinking is defined in the space of quantity conjecture, or of beliefs about the actions of other players; Cobweb is defined in the space of price conjectures. Furthermore, a minor discrepancy between the two emerges when $\bar{p}^*_{old} \neq P^*(\bar{q}_{old}, \bar{\theta}_{old})$, because the two iterative procedures then start from different initial guesses; this explains why the equivalence between the two has to be qualified by the first sentence in the above proposition. Notwithstanding these points, because price conjectures are tied to beliefs of the actions of others, the essence of the two connects is the same, and so are their implications with regard to the question of interest. We thus reach the following result.

Corollary 3. *Suppose $\partial P^*/\partial\theta = 0$. Similarly to Cobweb, level-k thinking attenuates the GE adjustment and reduces the gap between micro and macro elasticities in economies in which the GE effect amplifies the PE effect ($\alpha > 0$), but not in economies in which the GE effect offsets the PE effect ($\alpha < 0$).*

We illustrate this point in Figure 4. This figure contains two subfigures: one for the case of strategic complementarity ($\alpha > 0$), and another for the case of strategic substitutability ($\alpha < 0$). In each case, we draw graph of the best-response

[28] Note that, when every agent expects the aggregate quantity to be $\hat{q}$, part (ii) of the definition (which imposes individual rationality and market clearing) together with Lemma 2 (which rests on these elementary assumptions) implies that the *realized* aggregate quantity is $\bar{q} = \mathcal{BR}(\bar{\theta}_{new}, \bar{\theta}_{new}, \hat{q})$. This in turn explains why the assumed conjectures satisfy the recursion $\hat{q}_k \equiv \mathcal{BR}(\bar{\theta}_{new}, \bar{\theta}_{new}, \hat{q}_{k-1})$ for all $k \geq 1$.

function $\mathcal{BR}$ before and after the shock and use arrows to indicate the level-k thinking. The PE effect of the shock is captured by the vertical shift from point X to point Y, whereas the (frictionless) GE effect is captured by the shift from point Y to point Z. Note that the GE effect amplifies the PE effect when $\alpha > 0$ and offsets it when $\alpha < 0$. Finally, the arrows represent level-k thinking: level-0 is captured by the shift from point X to point Y (which means that level-0 coincides with PE); level-1 is captured by the shift from Y to Y'; and so on. It is then evident that level-k thinking helps capture the notion of incomplete GE adjustment when $\alpha > 0$, but opens the door to overshooting when $\alpha < 0$.[29]

6.3 Reflective Equilibrium

We now to turn attention to a different solution concept, called "reflective equilibrium", which was recently proposed by Garcıa-Schmidt and Woodford (2015). This concept has similar epistemological underpinnings as Level-k Thinking; we refer the reader to the original paper for a detailed discussion. But it also entails a different modeling choice, which which may appear to be of purely technical nature at first glance, but, as we explain here, helps avoid the "overshooting problem" discussed in the previous section.

Garcıa-Schmidt and Woodford (2015) develop and explore their concept in the context of a specific model, namely the New-Keynesian model. Moving beyond that specific context, the key idea is to let the subjective conjecture of a certain variable (e.g., inflation) to adjust *continuously* with the difference between the conjecture itself and value of that variable that gets realized if all agents act on the basis of that conjecture. For our purposes, we identify the relevant variable with the average afternoon price and adapt the concept of Garcıa-Schmidt and Woodford (2015) as follows.

Definition 5. *Fix a $T \in (0, \infty)$. The level-T reflective equilibrium is given by a conjecture $\hat{p}^*$ and a pair (q_m, p_m) of each m such that:*

(i) For all m, the pair (q_i, p_i) is consistent with household optimality, firm optimality, and market clearing in the local market for morning goods, under the following conjectures for the afternoon prices:

$$\hat{E}_m[p^*_{m'}] = \rho^2 P^*(q_m, \theta_m) + (1-\rho^2)\,\hat{E}_m[\bar{p}^*] \qquad \text{and} \qquad \hat{E}_m[\bar{p}^*] = \hat{p}^* \tag{40}$$

(ii) $\hat{p}^*$ is given by $\hat{p}^* = \hat{P}^*(T)$, where the function $\hat{P}^*$ is obtained by solving the following ODE:

$$\frac{d\hat{P}^*(t)}{dt} = \mathcal{T}\left(\hat{P}^*(t), \bar{\theta}\right) - \hat{P}^*(t) \quad \forall t \geq 0 \tag{41}$$

with initial condition $\hat{P}^*(0) = \bar{p}^*_{old} \equiv \mathcal{P}(\bar{\theta}_{old})$.

To interpret the above, note that, for every t, $\mathcal{T}\left(\hat{P}^*(t), \bar{\theta}\right)$ gives the actual average price that clears the market for afternoon goods when the quantity $\bar{q}$ is determined under the (incorrect) conjecture that this price equals $\hat{P}^*(t)$. Condition (41) therefore requires that the conjecture is adjusted upwards if the "actual" price exceeds the conjectured one, and downwards otherwise. The assumed concept is therefore similar to adaptive expectations 1956, except that the adjustments happen instantaneously, inside the minds of agents, and on the basis of *hypothetical* outcomes, as opposed to with the passage of calendar time and on the basis of the observation of *actual* past outcomes.

The assumed concept has also a similar flavor as Cobweb, in the sense that the conjecture is adjusted in the direction of the realized price. But whereas Cobweb requires the adjustment to be in discrete steps, with the conjecture in each round being *replaced* by the realized price in the previous round, the present concept lets the adjustment happen more smoothly, at infinitesimally small steps. This guarantees that the price conjecture never overshoot relative to the

[29] Corollary 3and the example in the figure assume $\partial P^*/\partial\theta = 0$. If we relax this assumption, the aforementioned overshooting can obtain even when $\alpha > 0$.

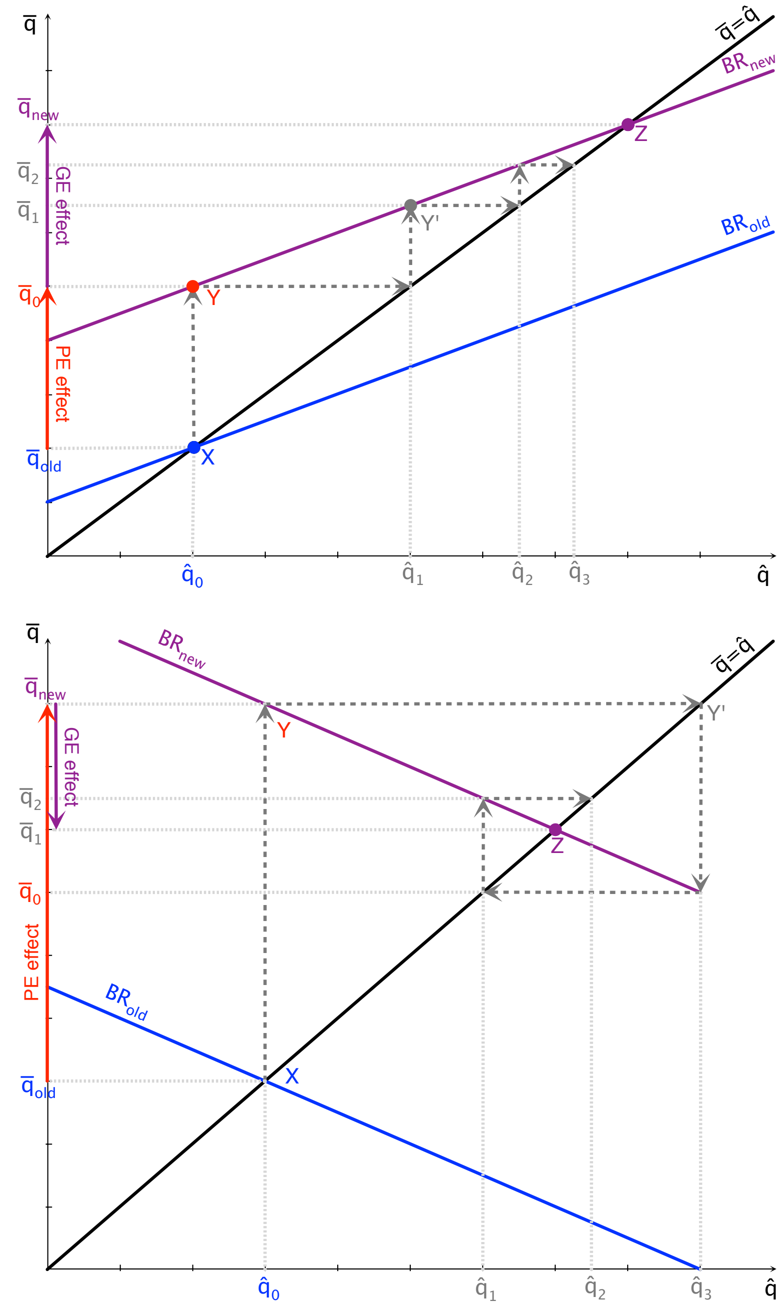

Figure 4: Level-k Thinking

frictionless benchmark, thus bypassing the aforementioned "pathology" of the Cobweb and level-k concepts: as we vary T, the price conjecture spans the entire interval between $\bar{p}^*_{old}$ and $\bar{p}^*_{new}$, in a continuous manner, and without ever overshooting outside of it, regardless of α. By the same token, the GE effect is necessarily attenuated, and the following is true.

Proposition 10 (Reflectie Equilibrium)**.** *For any $T \in (0, \infty)$, there exists a $T' \in (0, \infty)$ such that the level-T reflective equilibrium coincides with the Tâtonnement(T') solution, and vice versa.*

In the above, we have adapted the concept of Garcıa-Schmidt and Woodford (2015) in the space of price conjectures. Similar results obtain if we recast it in the space of quantity conjectures. That is, if we let the conjecture be $\hat{q} = \hat{Q}(T)$, where the function $\hat{Q}$ is obtained by solving the following ODE:

$$\frac{d\hat{Q}(t)}{dt} = \mathcal{BR}\left(\bar{\theta}_{new}, \bar{\theta}_{new}, \hat{Q}(t)\right) - \hat{Q}(t) \quad \forall t \geq 0 \tag{42}$$

with initial condition $\hat{Q}(0) = \bar{q}_{old}$. Under this perspective, the cognitive process we have introduced here can be thought of as a smooth version, not only of Cobweb, but also of level-k thinking: for any initial conjecture $\hat{q}$, the updated conjecture adjusts *smoothly* in the direction of $\mathcal{BR}(\bar{\theta}_{new}, \bar{\theta}_{new}, \hat{q})$ rather than jumping to it. As already noted, it is this kind of smoothness that avoids the overshooting problem of level-k thinking.

Remark. While we have focused exclusively on the case of a finite T, Garcıa-Schmidt and Woodford (2015) use the limit as $T \to \infty$ as selection device: they study a New-Keynesian model that admits multiple rational-expectations equilibria, only one of which happens to be selected by the reflective equilibrium, with an appropriate default point, in the limit as $T \to \infty$. In our context, the rational-expectations equilibrium is unique, so there is no selection issue. Yet, by letting $T < \infty$, we arrest the GE effect and connect this concept to the other concepts we have studied in this paper.

7 Discussion and Take-Home Lessons

The results of our paper up to this point can be summarized as follows:

Lack of Common Knowledge = Tâtonnement Dynamics = GE attenuation

Level-k Thinking = Cobweb Dynamics $\neq$ GE attenuation

In particular, the variants that are based on lack of common knowledge or Tâtonnement dynamics are able to capture the sought-after notion of GE attenuation regardless of whether the relevant GE effects amplify or offset the corresponding PE effects. By contrast, the variants that are based on Level-k Thinking and Cobweb dynamics have an inherent difficulty in capturing the sought-after notion of GE attenuation in settings in which the GE effects offset the PE effects, such as when agents compete for the same resources. That said, we also showed how the reflective-equilibrium concept of Garcıa-Schmidt and Woodford (2015) helps reconcile the *essence* of Level-k Thinking, if not its exact form, with the notion of GE attenuation. With this in mind, we would like to summarize the take-home lesson of our paper as follows.

Take-Home Lesson. *GE attenuation appears to be a robust implication of allowing either for plausible forms of bounded rationality, or for a realistic friction in the ability of the agents to reach common knowledge of the state of the economy and to coordinate their behavior.*

Does this mean that all the considered variants are inter-changeable for the kind of applied issues we are concerned with in this paper? Not necessarily.

First of all, any departure from the REE concept is, by its very nature, subject to Lucas' critique. Of course, one should not be dogmatic: there is nothing "sinful" in dropping the REE concept, nor is it a priori clear which methodological approach offers a more useful window into the real world. However, insofar as one assigns a positive value to interpretations of the data and to policy predictions that do not rest on assuming that agents make repeated, and systematic, mistakes, the balance is titled in the direction of the approach that relaxes the common-knowledge assumptions of the theory: among the considered variants, this approach is the only one that captures the sought-after notion of GE attenuation while also being immune to Lucas' critique.

Second, once this approach is adapted to multi-period settings, it necessarily predicts the following property: the GE effects of any given aggregate shock are attenuated relative to their frictionless counterparts, but get closer and closer to the latter as the time passes. The reason is that the level of common knowledge about the underlying aggregate shock increases endogenously over time as the agents observe past market outcomes. (We formalize this point in the next section.)

By contrast, the aforementioned property is *not* a necessary implication of the alternative approaches. In simple settings such as those we work with in this paper, one can reconcile Level-k Thinking with the idea that GE adjustment takes time by making the *additional* assumption that agents become "deeper thinkers" as time passes. However, this begs the question of whether this assumption is only a proxy for the aforementioned kind of learning. What is more, it is unclear how to adapt this assumption to stationary environments with recurring shocks: can agents be shallow thinkers vis-a-vis recent shocks and at the same time be deep thinkers vis-a-vis shocks that happened further in the past?

Consider next the following issue. In this paper, we defined PE as the adjustment that takes place in each marketplace holding constant the outcomes in other marketplaces. We also introduced a friction—whether in the form of removing common knowledge or in the form of relaxing the REE concept—in the conjectures that the agents in any given marketplace make about the behavior of agents in other marketplaces. We did not, however, allow a similar friction to be present *within* each marketplace. In so doing, we made sure that the relevant GE effects could be attenuated while leaving intact the underlying PE effects.

This begs the question of why one should expect the aforementioned friction to be less prevalent at the PE level than at the GE level. It is hard to tackle this question in a satisfactory manner without committing to a specific application. Yet, the following principle seems reasonable. Suppose that agents have more precise information about demand and supply conditions in the markets they themselves actively participate in, or about agents their directly interact with, than about markets and agents in "remote" markets (where "remote" could refer to geographic distance, distance in the commodity space, or distance in time). Our preferred approach then predicts that it is precisely equilibrium effects operating through more remote connections that are likely to be attenuated more.[30]

Finally, consider the relation to Gabaix (2016). That paper also considers a relaxation of the REE concept, but of a different kind that the ones we have studied here. The key assumption is that agents misperceive the law of motion of the endogenous state of the economy (equivalently, the actions of others) in a specific way: they expect it to be less responsive to aggregate shocks than what it actually does. This amounts to attenuating the relevant GE effects, in a manner similar to those we have considered here.[31] However, by assuming that the agents misperceived the law of motion in the opposite way, one is free obtain the opposite result, that is, amplification of GE effects. By contrast, our

[30] This suggests a formalization of our ideas in terms of incomplete-information networks. We leave this open for future research.

[31] Gabaix (2016)contains an additional assumption, which embeds the decision-theoretic friction proposed in Gabaix (2014): relative to a "fully rational" agent, the "sparse" agent responds less to any variation in the variables that enter her individual decision problem (prices, income, etc). This assumption amounts to attenuating the decision-theoretic, or PE, effects. Here, we are concerned with the *separate* assumption that the agents misperceive the law of motion of the endogenous state of the economy (or, equivalently, of the actions of others).

preferred approach does not offer this degree of freedom: because the variation in higher-order beliefs is necessarily bounded by the variation in first-order beliefs, the potency of the GE effects in the absence of common knowledge is *necessarily* lower than in the frictionless benchmark.

To conclude, we feel that, on the margin, the approach that removes common knowledge offers a more structured way of thinking about the issue of interest. We do not, however, wish to disparage any of the alternatives. To the contrary, we believe that both informational and behavioral frictions are realistic. We thus find it reassuring for our purposes—and doubly upsetting for the standard practice—that the two kinds of friction can complement each other in the direction of attenuating GE mechanisms.

8 GE Adjustment Takes Time

In the preceding sections, we have formulated the idea of *weak* or *incomplete* GE adjustment. We now seek to formulate the complementary idea of *slow* GE adjustment. To this goal, we introduce a variant framework, in which GE interactions happen repetitively, that is, over multiple periods after the shock has hit the economy. Like our baseline framework, the new framework is not meant to be either general or realistic. Rather, it is designed with the sole purpose of facilitating the adaptation of our earlier insights to a dynamic context.

8.1 Set Up

There is a continuum of marketplaces, indexed by $m \in [0,1]$, a double continuum of households, indexed by $i = (i_1, i_2) \in [0,1] \times [0,1]$, and multiple periods, indexed by $t \in \{0, 1, 2, \cdots, T\}$.[32] To simplify the exposition, we let $T = \infty$.[33] At any given point of time, each marketplace is populated by a measure one of households and a representative firm. Households may randomly from one marketplace to another as time passes and are the key decision-theoretic units in the model; firms are immobile and play an auxiliary role.

Households consume two goods in each period: leisure and a local final good. The latter is produced by the local firm in each marketplace, using the locally available labor and capital. Capital takes the form of multiple, imperfectly-substitutable, varieties. Each household is capable of transforming the final good into a single variety of capital, which can be used into production next period. The efficiency of this transformation depends on a local fundamental, which is specific to marketplace but stays constant over time. We denote the fundamental of marketplace m by θ_m and the corresponding aggregate by $\bar{\theta}$.

Marketplaces and matching. As in our baseline framework, the assumptions that markets are segmented but households can randomly relocate for one marketplace to another help disentangle partial- and general-equilibrium effects. To keep the analysis tractable, we now model the relocation of household as the product of random pairwise matching across the marketplaces.

At $t = 0$, there is a given allocation of households across the marketplaces. This allocation is such that each marketplace is populated by an equal measure of households. At the start of each period $t \geq 1$, each marketplace m is matched with another, randomly chosen, marketplace m'. At this point, a fraction $1 - \rho$ of the population from marketplace m relocates to marketplace m', and vice versa. Every household that relocates brings with her the capital she had accumulated in her old home. Following this relocation, the match is dissolved and each marketplace operates its own markets for the labor, the capital and the final good.

[32] Note that t and T now refer to calendar time. This should not be confused with the notation used in Section 4, where t and T referred to the number of iterations (or the depth) of Tâtonnement-like cognitive process.

[33] Note, however, that we still think of the overall time horizon as relatively short, say in the order of few quarters.

We let $M(i,t)$ denote the marketplace in which household i is located during period t. We let $I(m,t)$ denote the set of households who trade at marketplace m during period t. We finally adopt the convention that household $i=(i_1,i_2)$ is located at marketplace $m=i_1$ at $t=0$; that is, $M(i,0)=i_1$.

The firms. In each period t and each marketplace m, there is a competitive final-good firm, which employs the capital varieties and the labor of the households in $I(m,t)$. The produced quantity of the final good is given by

$$q_{m,t}=\kappa_{m,t}^{\omega}\ell_{m,t}^{1-\omega}, \tag{43}$$

where $\ell_{m,t}$ is the local supply of labor,

$$\kappa_{m,t}\equiv\left(\int_{i\in I(m,t)}k_{i,t-1}^{\frac{\sigma-1}{\sigma}}di\right)^{\frac{\sigma}{\sigma-1}},$$

is a CES composite of the local capital varieties, $k_{i,t-1}$ denotes the quantity of the capital variety owned and supplied by household i (note that this is determined in the previous period), $\sigma>0$ is the elasticity of substitution across varieties, and $\omega\in(0,1)$ is the capital share. Normalizing the price of the final good to one, we can therefore express the profits of the final-good firm in marketplace m and period t as

$$q_{m,t}-w_{m,t}\ell_{m,t}-\int_{i\in I(m,t)}p_{m,i,t}k_{i,t-1}di$$

where $w_{m,t}$ denotes the local wage and $p_{m,i,t}$ denotes the local price of the capital variety supplied by household i.

The households. Consider household i. Her preferences are given by

$$\sum_{t=0}^{T}\beta^t\left[c_{i,t}-v(n_{i,t})\right],$$

where $\beta\in(0,1)$ is her discount rate, $c_{i,t}$ denotes her consumption, $n_{i,t}$ denotes her labor supply, $v(n)=\frac{n^{1+\eta}}{1+\eta}$ measures the disutility of labor, and $\eta>0$. Her period-t budget constraint is given by

$$c_{i,t}+f\left(k_{i,t},\theta_m\right)=y_{i,t}\equiv w_{m,t}n_{i,t}+p_{m,i,t}k_{i,t-1},\quad \text{with}\quad m=M(i,t).$$

To interpret the above, recall first that $M(i,t)$ denotes the marketplace in which the household is located during period t. Next, note that $w_{m,t}n_{i,t}$ and $p_{m,i,t}k_{i,t-1}$ are, respectively, its labor and capital income. Finally, $f\left(k_{i,t+1},\theta_m\right)$ captures the cost of transforming the final good into the household's capital variety. We let $f(k,\theta)=\theta^{-\phi}\frac{k^{1+\phi}}{1+\phi}$, for some $\phi>0$. A "better local fundamental" (a higher θ) therefore means a lower cost of transforming the final good into capital.

Log-linearization, shocks, and elasticity. As in our earlier analysis, we henceforth re-interpret all the variables as log-deviations from a symmetric steady state and we work with the log-linearized version of the model. We also consider a once-and-for-all shock to $\bar{\theta}$ from some $\bar{\theta}=\bar{\theta}_{old}$ to some $\bar{\theta}=\bar{\theta}_{new}\neq\bar{\theta}_{old}$. We treat the initial fundamental, $\bar{\theta}_{old}$, as a fixed parameter and the change, $\Delta\bar{\theta}\equiv\bar{\theta}_{new}-\bar{\theta}_{old}$, as a random variable drawn from a Normal distribution centered around 0. We denote the corresponding change in the local fundamentals by $\Delta\theta_m$. We finally ask how the economy's outcomes (capital, output, etc) respond to this shock, not only on impact (at $t=0$) but also in all future periods ($t\geq 1$). The exercise conducted here is therefore similar as the one conducted before, except that now the key theoretical object is an entire impulse response function, describing the change in economic outcomes at each $t\geq 0$, as opposed to a single elasticity parameter.

8.2 Preliminaries

We start the analysis by imposing only Assumptions 1 and 2 along with the following assumption.[34] With these assumptions in place, we henceforth let $\hat{E}_{m,t}$ denote the common subjective expectation of all the agents in marketplace m during period t.

Fix a period t and a marketplace m. Because of the linearity of preferences in consumption, we can express the optimal (log-linearized) choices of any household $i \in I(m,t)$ as follows:

$$n_{i,t} = \tfrac{1}{\eta} w_{m,t} \qquad \text{and} \qquad k_{i,t} = \theta_m + \tfrac{1}{\phi} \hat{E}_{m,t} \left[p_{M(i,t+1),i,t+1} \right]. \tag{44}$$

The first part of this condition gives the optimal supply of labor; the second condition is the Euler condition and characterizes optimal capital accumulation. As households in the same marketplace share the same belief, we have, for all $i \in I(m,t)$, $n_{i,t} = \ell_{m,t}$ and

$$k_{i,t} = k_{m,t} \equiv \int_{i \in I(m,t)} k_{i,t} di. \tag{45}$$

Turning to the local firm, on the other hand, we have

$$w_{m,t} = q_{m,t} - \ell_{m,t}, \qquad p_{m,t} = q_{m,t} - \kappa_{m,t} \qquad \text{and} \qquad p_{m,i,t} - p_{m,t} = \tfrac{1}{\sigma} \left(\kappa_{m,t} - k_{i,t-1} \right), \tag{46}$$

where $p_{m,t} = \int_{i \in I(m,t)} p_{m,i,t} di$ is the (log-linearized) ideal price index for the capital composite in marketplace m. The first part of the above condition gives the firm's demand for labor; the second part gives its demand for the capital composite; the third part gives its demand for each particular capital variety.

By imposing market clearing in the labor market, and using the household's supply of labor and the firm's demand for labor, we get

$$\ell_{m,t} = \tfrac{1}{1+\eta} q_{m,t} \qquad \text{and} \qquad w_{m,t} = \tfrac{\eta}{1+\eta} q_{m,t}. \tag{47}$$

By the production function, on the other hand, we have $q_{m,t} = \omega \kappa_{m,t} + (1-\omega)\ell_{m,t}$. It follows that

$$q_{m,t} = \psi \kappa_{m,t}, \tag{48}$$

where $\psi \equiv \frac{(1+\eta)\omega}{\eta+\omega} \in (0,1)$. By the firm's demand for the capital composite, we then get

$$p_{m,t} = -\chi \kappa_{m,t}, \tag{49}$$

where $\chi \equiv \frac{\eta(1-\omega)}{\eta+\omega} \in (0,1)$. Combing the above results, we infer that the local capital stock $\kappa_{m,t}$ is a sufficient statistic for local quantities and local prices. By the same token, the aggregate investment in period $t-1$ (which is the capital stock in period t) is a sufficient statistic for aggregate outcomes in period t.

Let $\bar{q}_t \equiv \int q_{m,t} dm$ and $\bar{\ell}_t \equiv \int \ell_{m,t} dm$ denote the aggregate levels of, respectively, output and employment; let $\bar{w}_t \equiv \int w_{m,t} dm$ and $\bar{p}_t \equiv \int p_{m,t} dm$ denote, respectively, the average wage and the average price of capital; and finally let $\bar{k}_{t-1} \equiv \int k_{m,t-1} dm$ denote the aggregate investment in period $t-1$ (with the convention that $\bar{k}_{-1} = 0$). We summarize the analysis so far as follows.

Lemma 11. *There exist a known linear vector function F such that, for every m and every t,*

$$(q_{m,t}, \ell_{m,t}, w_{m,t}, p_{m,t}) = F(\kappa_{m,t}) \qquad \textit{and} \qquad (\bar{q}_t, \bar{\ell}_t, \bar{w}_t, \bar{p}_t) = F(\bar{k}_{t-1}).$$

[34] Note that Assumptions 1 and 2 are satisfied, not only in the frictionless benchmark, but also maintained in the considered modifications.

Note that this characterizations has relied only on Assumptions 1 and 2.

To make further progress, we now add the following assumption:

Assumption 9. *Agents have first-order knowledge of the structure of the economy (i.e., of the preference and technology parameter and of the processes of matching and sampling) and of Assumption 1 (i.e., that markets clear and that other agents optimize given subjective beliefs).*

Fix a period t and a marketplace m, consider any household $i \in I(m,t)$, and let $m' = M(i, t+1)$ denote the location of that household in period $t+1$. Thanks to the above assumption, the household can reason that, regardless of what m' turns out to be, the price for her own capital will satisfy

$$p_{m',i,t+1} = p_{m',t+1} + \tfrac{1}{\sigma}\left(\kappa_{m',t+1} - k_{i,t}\right) = \left(1 - \tfrac{1}{\chi\sigma}\right)p_{m',t+1} - \tfrac{1}{\sigma}k_{i,t}, \tag{50}$$

Next, because a household expects to stay in her current marketplace with probability ρ and to be randomly reallocated with the remaining probability, her expectation of $p_{m',t+1}$ must satisfy

$$\hat{E}_{m,t}\left[p_{m',t+1}\right] = \rho\hat{E}_{m,t}[p_{m,t+1}] + (1-\rho)\hat{E}_{m,t}[\bar{p}_{t+1}], \tag{51}$$

Furthermore, because the household knows that $p_{m,t+1}$ will satisfy $p_{m,t+1} = -\chi\kappa_{m,t+1}$, and because the household also knows that $\kappa_{m,t+1}$ will be given by mixture of the investment made by the current households in marketplace m and of the investment made in a random other marketplace (which will become m's match in period $t+1$), the following is true:

$$\hat{E}_{m,t}[p_{m,t+1}] = -\chi\hat{E}_{m,t}\left[\rho k_{m,t} + (1-\rho)\,\bar{k}_t\right].$$

Finally, because the household can reason that all other households in her marketplace choose the same investment as herself, and because $\bar{p}_{t+1} = -\chi\bar{k}_t$, we can re-write that above as follows:

$$\hat{E}_{m,t}[p_{m,t+1}] = -\chi\rho k_{i,t} + (1-\rho)\hat{E}_{m,t}[\bar{p}_{t+1}]. \tag{52}$$

Combining (51) and (52), we infer that

$$\hat{E}_{m,t}\left[p_{m',t+1}\right] = -\rho^2\chi k_{i,t} + (1-\rho^2)\hat{E}_{m,t}[\bar{p}_{t+1}]. \tag{53}$$

Plugging the above into household's FOC, namely condition (44), solving the resulting equation for $k_{i,t}$ as function of θ_m and $\hat{E}_{m,t}[\bar{p}_{t+1}]$. and using the fact that $k_{i,t} = k_{m,t}$, we reach at the following result.

Lemma 12. *There exists a linear function K such that, for all m and t,*

$$k_{m,t} = K\left(\ \hat{E}_{m,t}[\bar{p}_{t+1}]\ ,\ \theta_m\ \right).$$

To sum up, so far we have made the following elementary assumptions: that agents maximize; that markets clear; that agents have first-order knowledge of these facts and of the structure of economy; and that subjective beliefs are homogeneous within each marketplace. Importantly, we have not yet assumed either that the agents across different marketplaces have common knowledge of these facts, or that they have common knowledge of the aggregate fundamentals, or that their subjective beliefs satisfy the REE concept. On the basis the aforementioned, and much weaker, assumptions, we have obtained Lemmas 11 and 12. The first lemma gives output, employment, and prices in each period and each marketplace as functions of the local capital stock. The second lemma tells us how capital evolves as function of the subjective beliefs of a key general-equilibrium object, the average price of capital. To complete the characterization of the model's outcomes, what remains to do is to specify how these subjective beliefs are formed and how the adjust to the underlying aggregate shock. We do so in the sequel, under different assumptions about the solution concept and/or the level of common knowledge in the economy.

8.3 Frictionless Benchmark

Similarly Section 3, we define the frictionless benchmark by imposing the REE concept together with commonly-shared information (and therefore common knowledge of the fundamentals).

Definition 6. *The frictionless outcomes are those obtained when agents share the same information and form rational expectations.*

In this benchmark, we have that, for all i and all t,

$$\hat{E}_{m,t}[\bar{p}_{t+1}] = E_t[\bar{p}_{t+1}],$$

where E_t denotes the rational expectation conditional on the information that is commonly available in period t. Using the above in Lemma 12 and aggregating, we get

$$\bar{k}_t = K\left(E_t[\bar{p}_{t+1}], \bar{\theta}\right).$$

Because $\bar{p}_{t+1} = -\chi \bar{k}_t$, we can rewrite the above as

$$\bar{p}_{t+1} = \mathcal{T}\left(E_t[\bar{p}_{t+1}], \bar{\theta}\right), \tag{54}$$

where $\mathcal{T}(\cdot) \equiv -\chi K(\cdot)$. Since both $\bar{\theta}$ and $E_t[\bar{p}_{t+1}]$ are commonly known in period t, so is $\bar{p}_{t+1}$. We can thus restate the above as

$$\bar{p}_{t+1} = \mathcal{T}\left(\bar{p}_{t+1}, \bar{\theta}\right).$$

The slope of $\mathcal{T}$ is given by

$$\alpha \equiv \frac{\partial \mathcal{T}}{\partial p} = \frac{(1-\sigma\chi)\left(1-\rho^2\right)}{\sigma\phi + \sigma\chi\rho^2 + (1-\rho^2)}, \tag{55}$$

where $\chi = \frac{\eta(1-\omega)}{\eta+\omega}$, as defined above. The mapping $\mathcal{T}$ and the scalar α have similar interpretations as the corresponding objects in our baseline model, although they of course admit different functional forms.

It is easy to verify that α is necessarily less than $+1$; it can be either positive or negative;[35] and is higher than -1 as long as ϕ is high enough and/or σ is low enough.[36] We henceforth assume that the latter property holds so as to guarantee that $\mathcal{T}$ defines a contraction mapping.

For every $t \geq 1$, we then have that

$$\bar{p}_t = \mathcal{P}(\bar{\theta}),$$

where $\mathcal{P}$ is the fixed point of $\mathcal{T}$. Using this into Lemma 12, we infer that

$$k_{m,t} = K(\mathcal{P}(\bar{\theta}), \theta_m), \tag{56}$$

[35] Given that $\rho \in (0,1)$, we have that $\alpha \in (0,1)$ if and only if $\sigma\chi < 1$, $\alpha = 0$ if and only if $\sigma\chi = 1$, and $\alpha < 0$ if and only if $\sigma\chi > 1$.

[36] When $\phi + \chi(2\rho^2 - 1) \geq 0$, $\alpha > -1$ regardless of σ; and when $\phi + \chi(2\rho^2 - 1) < 0$, $\alpha > -1$ if and only if $\sigma < \tilde{\sigma}$, for some $\tilde{\sigma} > 1$.

for all m and all t. From Lemma 11, we can then express the local outcomes $(q_{m,t}, \ell_{m,t}, w_{m,t}, p_{m,t})$ as linear functions of the pair $(\theta_m, \bar{\theta})$ and the corresponding aggregates as linear function of $\bar{\theta}$.

We are now ready to characterize the response of the economy to the aggregate shock. Without serious loss of generality, we focus on investment as the observables of interest; the responses of all other variables can be inferred from Lemma 11 and feature a similar disconnect between micro and macro effects.

Proposition 11. *There exist scalars ϵ^{Macro} and ϵ^{micro}, with $\epsilon^{Macro} \neq \epsilon^{micro}$, such that, for all t,*

$$\bar{k}_t = \bar{k}_{new} \equiv \bar{k}_{old} + \epsilon^{Macro} \Delta\bar{\theta},$$

and, for all t and all m,

$$k_{m,t} = \bar{k}_{new} + \epsilon^{micro} \left(\Delta\theta_m - \Delta\bar{\theta}\right).$$

Part (i) means that the aggregate outcomes feature a very simple dynamic response to the aggregate shock. On impact, aggregate investment $\bar{k}_t$ jumps form $\bar{k}_{old} = K(\mathcal{P}(\bar{\theta}_{old}), \bar{\theta}_{old})$ to $\bar{k}_{new} = K(\mathcal{P}(\bar{\theta}_{new}), \bar{\theta}_{new})$ and stays constant thereafter. And because aggregate employment and output are determined by previous-period investment, $\bar{n}_t$ and $\bar{q}_t$ exhibit exactly the same IRF as $\bar{k}_t$, only lagged by one period. In short, the impulse response functions (IRFs) of all these aggregate quantities to the aggregate shock are totally flat, at a level equal to scalar ϵ^{Macro} for investment and at a level proportional to ϵ^{Macro} for all other quantities. The scalar ϵ^{Macro} therefore has a similar meaning as the macro elasticity in our baseline, except that it now describes the entire dynamic response of the economy.

Part (ii) turns attention to the cross section and identifies the scalar ϵ^{micro} as the dynamic analogue of the micro elasticity in our baseline framework: this scalar summarizes the IRF of local investment to local shocks or, equivalently, to cross-sectional variation in the exposure to the aggregate shocks.

It is straightforward to check that $\epsilon^{Macro} \neq \epsilon^{micro}$ if and only if $\alpha \neq 0$. The macro elasticities of interest differ from their micro counterparts because of the mobility of capital across marketplaces. Depending on parameters, this GE effect can cause ϵ^{Macro} to be either higher or lower than ϵ^{micro}. Intuitively, the are two opposite forces at work. On the one hand, investment choices tend to be strategic substitutes because an increase in the aggregate capital stock raises wages and depresses the aggregate return to capital. On the other hand, investment choices tend to be strategic complements because an increase in the aggregate capital stock raises the demand for each individual variety (insofar as the different varieties are imperfect substitutes in production). When the first effect dominates, $\alpha < 0$ and $\epsilon^{Macro} < \epsilon^{micro}$; otherwise, $\alpha < 0$ and $\epsilon^{Macro} > \epsilon^{micro}$.

Remark. In our setting, the frictionless benchmark lacks any interesting dynamic patterns in the responses either of aggregate outcomes to aggregate shocks or of local outcomes to local shocks: the IRFs are flat and boring. In applications, interesting dynamic patterns can emerge from various forms of adjustment costs embedded in preferences (e.g., desire to smooth consumption) or technology (e.g., adjustment costs to labor or capital). By abstracting from such effects, we sharpen the comparison between the frictionless benchmark and the modifications studied in the sequel: in these modifications, the aggregate IRFs are non-flat because and only because of the kind of GE attenuation we are interested in.

8.4 Tâtonnement Dynamics

In the aforementioned benchmark, $\bar{k}_t$ jumps for from $\bar{k}_{old}$ to $\bar{k}_{new}$ as soon as the aggregate shock has hit the economy. Behind this instantaneous adjustment in quantities, there is an instantaneous adjustment in the expected and the actual price of capital: $\mathbb{E}_t[\bar{p}_{t+1}]$ and $\bar{p}_{t+1}$ alike jump from $\bar{p}_{old} \equiv \mathcal{P}(\bar{\theta}_{old})$ to $\bar{p}_{new} \equiv \mathcal{P}(\bar{\theta}_{new})$ as soon as the shock hits. We now slow down this adjustment by relaxing the solution concept.

Similarly to what we did in Section 4, we let the aforementioned beliefs be the product of a Tâtonnement process rather than being pinned down by rational expectations (or perfect foresight). The only difference is that now the Tâtonnement process take places in "real time": additional rounds of updating require more time.

Let $\mathcal{N}\left(\bar{p},\bar{\theta}\right) \equiv -\frac{1}{\chi}\bar{p} - K\left(\bar{p},\bar{\theta}\right)$ measure the excess aggregate demand for capital that obtains in any given period if both the average price of capital and the previous-period expectation of it are given by p.[37] Next, let the function $\hat{P}(\cdot)$ be the solution to the following ODE:

$$\frac{d\hat{P}(\tau)}{d\tau} = \mathcal{N}\left(\hat{P}(\tau),\bar{\theta}_{new}\right) \quad \forall\tau \geq 0 \tag{57}$$

with initial condition $\hat{P}(0) = \bar{p}_{old}$. That is, $\hat{P}(\tau)$ is defined in exactly the same fashion as in Definition 1 and identifies the price conjecture generated from the Tâtonnement-like cognitive process when its depth, or the number of rounds, is τ. Finally, specify the period-t conjecture in our dynamic economy as follows:

$$\hat{E}_{i,t}\left[\bar{p}_{t+1}\right] = \hat{p}_{t+1} \equiv \hat{P}(f(t)) \tag{58}$$

where $f : \mathbb{N} \to \mathbb{R}^{+}$ is a function that maps the calendar time to Tâtonnement rounds. This function is assumed to be strictly increasing—so that more rounds require more time or, conversely, time helps the agents become "deeper thinkers"—but is otherwise a "free parameter" that controls the speed of adjustment in the relevant conjectures.

This construction yields a simple translation of the static analysis in Section 4to the dynamic framework under consideration. Indeed, it is straightforward to check that the following variant of Lemma 5 applies in the present context.

Lemma 13. *There exists a strictly increasing sequence $\{w_t\}_{t=0}^{\infty}$, with $w_t \in (0,1)$ for all t, such that the period-t conjecture about the period-$(t+1)$ price of capital satisfies*

$$\hat{p}_{t+1} = \bar{p}_{old} + w_t(\bar{p}_{new} - \bar{p}_{old}),$$

By Lemma 12, the local and the aggregate level of investment is then give by, respectively,

$$k_{m,t} = K\left(\hat{p}_{t+1},\theta_m\right) \qquad \text{and} \qquad \bar{k}_t = K\left(\hat{p}_{t+1},\bar{\theta}\right), \tag{59}$$

where K is the same function as before.

It then also follows that we can express the dynamic response of $\bar{K}_t$ to the shock as follows:

Proposition 12. *In the Tâtonnement variant described above, the dynamic response of aggregate investment to the aggregate shock is given by*

$$\bar{k}_t = \bar{k}_{old} + \left\{\epsilon^{micro} + w_t(\epsilon^{Macro} - \epsilon^{micro})\right\}\Delta\bar{\theta}, \tag{60}$$

where $\{w_t\}_{t=0}^{\infty}$ is the same sequence as in Lemma 13

We illustrate this result in Figure 5 for two economies in which $f(0) \approx 0$ and $f(\infty) \approx 1$ (meaning that the relevant conjecture does not move at all on impact but converges to the REE counterpart as time passes). The economy in the left panel features $\epsilon^{Macro} > \epsilon^{micro} > 0$, meaning that the GE effect reinforces the PE effect; the economy in the right

[37] To understand the formula for $\mathcal{N}$, take any $t \geq 1$. From condition (49), the average demand for capital is given by $\bar{\kappa}_t = -\frac{1}{\chi}\bar{p}_t$. Next, suppose that all agents expect, in period $t-1$, that the period-t average price of capital will be $\bar{p}_t$.

Then, from Lemma 12, we have that the average investment in period $t-1$, and therefore also the average supply of capital in period t, is given by $\bar{k}_{t-1} = K(\bar{p}_t,\bar{\theta})$. It follows that the excess demand for capital in period t equals $-\frac{1}{\chi}\bar{p}_t - K(\bar{p}_t,\bar{\theta})$, which explains the formula for $\mathcal{N}$ used above. Future reference, let us also note that the excess demand is necessarily downward sloping: $\frac{\partial\mathcal{N}}{\partial p} = -\frac{1}{\chi}(1-\alpha) < 0$.

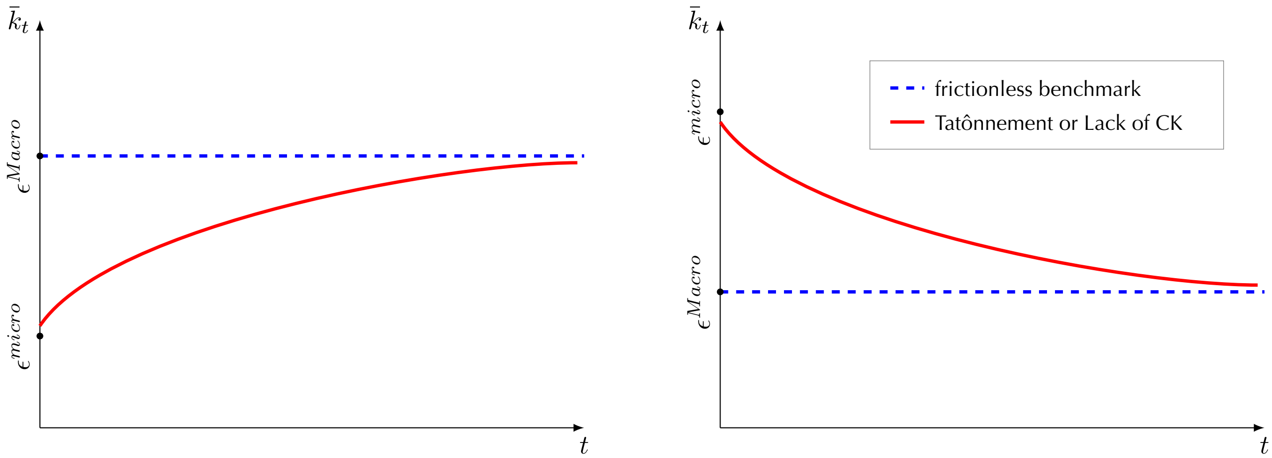

Figure 5: GE adjustment takes time. (Left panel: GE amplifies PE. Right panel: GE offsets PE.)

panel features the opposite property. We normalize the size of the shock to $\Delta\bar{\theta} = 1$ and draw the dynamic response of $\bar{k}_t$ under two scenarios: the Tâtonnement variant under consideration (solid line), and the frictionless benchmark studied earlier (dashed line).

In that benchmark, as soon as the shock hits the economy, $\bar{k}_t$ jumps up by a quantity equal to ϵ^{Macro} and stays at this higher level for ever after. In the Tâtonnement variant, instead, the initial jump in $\bar{k}_t$ is approximately equal to ϵ^{micro}. Depending on whether we consider the economy on the left panel or the one on the right, this initial response can represent either an under-reaction or an over-reaction relative to the frictionless benchmark. In both cases, however, $\bar{k}_t$ converges monotonically from ϵ^{micro} to ϵ^{Macro} as time passes: the under- or over-reaction is corrected with the passage of time. This dynamic pattern reflects that property that, due to the sluggish adjustment in the relevant price conjectures, the GE effect is inactive in the short run, but builds force as time passes.

This completes the analysis of our Tâtonnement variant. We find this variant to be useful for two reasons. First, it helps capture the notion that the adjustment for one equilibrium point to another may take time in a very direct, albeit old-fashioned, manner. Second, it has similar methodological underpinnings as the other variants that drop the REE concept, such Level-k Thinking or Reflective Equilibrium. However, as anticipated in Section 7, all these variants raise difficult conceptual and practical questions, which instead can be bypassed with our preferred approach. In the sequel, we show how this approach can generate the same dynamic response as the one seen in Figure 5 without violating rational expectations.

8.5 Dynamics without Common Knowledge

We now bring back the REE concept but remove common knowledge of the aggregate shock. Similarly to Section 5, we summarize the information that is available to marketplace m at $t = 0$ in a local signal s_m given by

$$s_m = \Delta\bar{\theta} + v_m,$$

where $\Delta\bar{\theta}$ is the underlying change in the fundamentals, itself modeled as a random variable drawn from $N(0, \sigma_\theta^2)$, and v_i is an idiosyncratic noise term, drawn from $N(0, \sigma_v^2)$, i.i.d. across marketplaces, and independent of $\Delta\bar{\theta}$.[38] Unlike our earlier analysis, however, we must now take into account the fact that the information structure changes endogenously over time, as agents relocate and trade.

[38] The signal s_m includes the information contained in the local fundamental θ_m.

In general, the task of characterizing the dynamics of beliefs in settings with strategic or GE interactions, dispersed private information, and endogenous learning can be rather taunting. For example, Huo and Takayama (2015) prove that an exact finite-state-space solution is impossible for a large class of such settings. Here, we have keep the aforementioned task manageable thanks to two key assumptions: that there is a single, once-and-for-all aggregate shock; and that the endogenous learning reduces, in effect, to a direct exchange of the private information of any two marketplaces within any realized match.

Let us explain. Fix a period $t \geq 1$, a particular marketplace m, and a particular realization of the aforementioned pairwise matching, and let m' denote the match of marketplace m. At the moment its local markets for capital and labor open, marketplace m is populated by two types of agents: those that we previously located in m (the "locals") and those that were previously located in m' (the "foreigners"). By observing the local prices $p_{m,t}$ and/or $w_{m,t}$, both type of agents can infer the local capital stock, $\kappa_{m,t}$ Recall that the latter is given by a weighted average of the investment made by the two type of agents, i.e., $\kappa_{m,t} = \rho k_{m,t-1} + (1-\rho^2)k_{m',t-1}$. Furthermore, each type knows its own investment, i.e., the locals know $k_{m,t-1}$ and the foreigners know $k_{m',t-1}$. It follows that the observation of the local prices perfectly reveals $k_{m,t-1}$ to the foreigners and $k_{m',t-1}$ to the locals. Next, note that $k_{m,t-1} = K(\theta_m, \mathbb{E}_{m,t-1}[\bar{p}_t])$ and $k_{m',t-1} = K(\theta_{m'}, \mathbb{E}_{m',t-1}[\bar{p}_t])$, where K is a commonly known function. Furthermore, the foreigners directly observe θ_m as soon as they arrive in marketplace m. It follows that, by observing the local prices and learning $k_{m,t-1}$, the foreigners also learn $\mathbb{E}_{m,t-1}[\bar{p}_t]$. To simplify the analysis, we assume the foreigners also tell the locals what $\theta_{m'}$ was. It follows that, by learning $k_{m',t-1}$, the locals learn $\mathbb{E}_{m',t-1}[\bar{p}_t]$. Furthermore, because there is no aggregate shock other than the one in the $\bar{\theta}$, the equilibrium $\bar{p}_t$ is a known, albeit time-varying, function of $\bar{\theta}$. It follows that the local learn $\mathbb{E}_{m',t-1}[\bar{\theta}]$ and the foreigners learn $\mathbb{E}_{m,t-1}[\bar{\theta}]$. That is, it is *as if* the two types of agents exchange their (pre-trading) posterior beliefs about $\bar{\theta}$.

We are thus able to show the following.

Lemma 14. *There exists a deterministic sequence $\{\lambda_t\}_{t=0}^{+\infty}$, with $0 < \lambda_t < \lambda_{t+1} < 1$ for all t, such that*

$$\bar{E}_t^h[\bar{\theta}] = \bar{\theta}_{old} + \lambda_t^h \Delta\bar{\theta} \quad \forall t \geq 0 \text{ and } h \geq 1,$$

where $\bar{E}_t^h[\cdot]$ denote that h-th order average belief operator at period t.[39] *Moreover, holding all other parameters constant, $\lambda_t \to \infty$ as $t \to \infty$.*

The characterization of sequence $\{\lambda_t\}_{t=0}^{T}$ can be found in the Appendix. For the present purposes, it suffices to note that the lack of common knowledge diminishes over time and eventually vanishes. That is, for any given $h \geq 1$, the average h-th order belief of the aggregate fundamental moves monotonically from a value closer to $\bar{\theta}_{old}$ to a value closer to $\bar{\theta}_{new}$ as time passes. But the higher h is, the more anchored the h-th order belief is to $\bar{\theta}_{old}$ and the more time it takes for it to cover any given distance between $\bar{\theta}_{old}$ and $\bar{\theta}_{new}$.

We now explore what these properties imply for the observables of the economy. From Lemma 12, local investment is now given by

$$k_{m,t} = K\left(E_{m,t}\left[\bar{p}_{t+1}\right], \theta_m\right), \tag{61}$$

where K is the same function as before and $E_{m,t}[\cdot]$ is the rational expectation conditional on the information that is available to marketplace m in period t. Together with condition (49), we get the following fixed-point relation between the realized price $\bar{p}_{t+1}$ and the average expectation of it in period t

[39] This defined recursively by $\bar{E}_t^1[\cdot] \equiv \int E_{m,t}[\cdot]dm$ and $\bar{E}_t^h[\cdot] \equiv \int E_{m,t}[\bar{E}_t^{h-1}[\cdot]]di$ for all $h \geq 2$, where $E_{m,t}[\cdot]$ denotes the rational expectation conditional on the information that is available to marketplace m in period t.

$$\bar{p}_{t+1} = \mathcal{T}\left(\bar{E}_t\left[\bar{p}_{t+1}\right], \bar{\theta}\right),$$

where $\mathcal{T}$ is the same mapping as before. Similarly to Corollary 2 in our baseline framework, we can thus express the equilibrium expectations of the price of capital at any given period t as a linear combination of the contemporaneous hierarchy of beliefs of the underlying fundamentals:

$$\bar{E}_t\left[\bar{p}_{t+1}\right] = \gamma(1-\alpha)\sum_{h=1}^{\infty}\alpha^{h-1}\bar{E}_t^h\left[\bar{\theta}\right], \tag{62}$$

where $\gamma \equiv \frac{\partial \mathcal{P}}{\partial \theta}$ is the elasticity of the price with respect to $\bar{\theta}$ in the frictionless benchmark and where $\alpha \equiv \frac{\partial \mathcal{T}}{\partial p}$ is the slope of the mapping $\mathcal{T}$. Combining this result with Lemma 14, we reach at the following characterization of the dynamics of the equilibrium expectations.

Lemma 15. *For all $t \geq 0$, the equilibrium expectations of the price of capital are given by*

$$\bar{E}_t\left[\bar{p}_{t+1}\right] = \bar{p}_{old} + \pi\left(\lambda_t\right)\left(\bar{p}_{new} - \bar{p}_{old}\right), \tag{63}$$

where $\pi\left(\lambda\right)$ is continuous and strictly increasing in λ, with $\pi(0) = 0$ and $\pi(1) = 1$.

The following is then a direct implication.

Proposition 13. *In the Tâtonnement variant described above, the dynamic response of aggregate investment to the aggregate shock is given by*

$$\bar{k}_t = \bar{k}_{old} + \left\{\epsilon^{micro} + \pi\left(\lambda_t\right)\left(\epsilon^{Macro} - \epsilon^{micro}\right)\right\}\Delta\bar{\theta}, \tag{64}$$

where π is the same function as the one appearing in Lemma 15.

This provides our second, and preferable, formalization of the notion that "GE adjustment takes times": for t low enough, the change in $\bar{k}_t$ relative to the change in $\bar{\theta}$ is close to ϵ^{micro}; but as time passes, this change gets closer and closer to ϵ^{Macro}.

Clearly, this pattern is qualitative similar to the one seen in Figure 5. In fact, if we were free to choose any increasing sequence $\{\lambda_t\}$ in $(0, 1)$, we could rationalize any increasing pattern in the response of aggregate investment. In fact, it is straightforward to see that the following analogue of Proposition 7 applies: for any function f in the Tâtonnement variant, there exists an increasing sequence $\{\lambda_t\}_{t=0}^{+\infty}$ such that the dynamic adjustment in the price conjectures and the observables of the incomplete-information variant coincide with those in the Tâtonnement variant, and vice versa. The only caveat is that, whereas in the Tâtonnement variant one is free to choose the function f and the resulting IRFs at will, here the sequence $\{\lambda_t\}_{t=0}^{+\infty}$ and the associated IRFs are dictated by the trading structure along with Bayesian updating.

To sum up, the message of Figure 5 survives in the following sense: what was an ad hoc, *off-equilibrium,* adjustment process has now been recast as the *equilibrium* dynamics of a modified economy, in which common knowledge of the underlying aggregate shock is attained only gradually. Two additional lessons, however, have also emerged. The first is that the speed at which the GE effect kicks in, and therefore also the speed at which the gap between micro and macro elasticities grows, depends on the quality of the learning that takes place through markets or other means. The second is that, precisely because of this learning, the aforementioned gap *has* to increase with time. By contrast, in the Tâtonnement variant, we could have kept the gap constant by assuming that the "depth of reasoning" does not increase with time.

8.6 Variants

It is straightforward to extend the analysis of this section to the variants studied in Section 6. In particular, we can adapt Definitions 3 and 4 (i.e., our Cobweb and Level-k concepts) to the current setting by introducing a monotonic transformation between the calendar time, t, and the depth, k, of the price and quantity conjectures defined in, respectively, conditions (36) and (38). Following similar steps as in Subsection 6, we can then show that the IRF of $\bar{k}_t$ in the modified economy converges monotonically from ϵ^{micro} to ϵ^{Macro} in the case in which $\alpha > 0$, but not otherwise: when $\alpha < 0$, the Cobweb and Level-k variants predict that the IRF of $\bar{k}_t$ in the modified economy oscillates around its frictionless counterpart, thus failing to capture the sought-after notion. This is because of the overshooting of the price and/or quantity conjectures that we described in Section 6—a "pathology" that, as already explained, can be avoided with the variant that builds on Garcıa-Schmidt and Woodford (2015), provided, of course, that we let that variant's depth of reasoning to increase with the time lag since the shock hits the economy.

With these qualifications in mind, one can view all our non-REE variants as close substitutes to the one studied in Subsection 8.5: the common message is that "GE adjustment takes time". Nevertheless, we favor, on the margin, the variant that is based on lack of common knowledge for the reasons already discussed in Section 7: this variant appears to offer, not only a natural explanation of why the potency of the GE effect has to increase with time, but also a method for accommodating the desired notions in stationary settings with recurring shocks.

Remark. Earlier works such as Woodford (2003) and Angeletos and La'O (2010) have already studied settings featuring recurring shocks, lack of common knowledge, and strategic complementarity; they have then proceeded to show how the inertia of higher-order beliefs can generate sluggishness in the response of certain aggregate outcomes to the underlying aggregate shocks. Clearly, our analysis in this section builds heavily on these earlier works. However, these earlier works have not provided the PE-vs-GE perspective of the present paper, nor have they tried to build the kind of bridges we have built to various other literatures. Furthermore, by focusing on settings with strategic complementarity, these earlier works have not highlighted that, in settings with strategic substitutability, the inertia in higher-order beliefs has the *opposite* observable implication in the following regard: in such settings, lack of common knowledge causes aggregate outcomes to respond *more* than in the frictionless setting. As we have explained in this paper, the robust prediction is that lack of common knowledge attenuates GE effects. In some applications, this attenuation will manifest as under-reaction to shocks; in others, the same mechanism will manifest as over-reaction; but in all cases, it will reduce the gap between micro and macro elasticities in the short run.

9 Conclusions

General-equilibrium effects that operate at the economy-wide level are central to understanding the response of macroeconomic outcomes to aggregate shocks, as well as to policy changes. Such effects limit the usefulness of partial-equilibrium intuitions. They also introduce a gap between the macroeconomic effects of interest and the kind of micro or local elasticities that are estimated in a growing empirical literature.

In this paper, we sought to reduce this gap and to reinforce the usefulness of partial-equilibrium intuitions by operationalizing the notion that general-equilibrium effects may be less potent, or may take more time to settle in, than what is presumed in standard modeling practice. We attained this goal by building on existing insights from various strands of the literature, but also by blending them in a new—and hopefully refreshing—manner.

More specifically, we considered an elementary Walrasian economy, in which trading was sequential and decentralized. We fixed its microeconomic foundations in terms of the specification of preferences, technologies, and market structures. In so doing, we took the relevant demand and supply structures are given. We characterized the response of this economy to an aggregate shock under a benchmark specification that, in line with standard practice, imposed rational expectations along with common knowledge of the aggregate shock. We then departed from that benchmark in two distinct ways. In one, we dropped the rational-expectations solution concept in favor of certain cognitive processes, which mimicked Tâtonnement, Cobweb dynamics, level-k thinking, or the reflective-equilibrium concept proposed by Garcıa-Schmidt and Woodford (2015). In the other, we maintained the rational-expectations solution concept but removed common knowledge of the aggregate shocks. We explored the similarities and the differences of the various alternatives. We concluded that, although both approaches help accommodate the sought-after notion, the latter approach—that is, the one that maintains rational expectations but removes common knowledge—appears to do so in a more natural and more structured manner.

Our framework was deliberately simple and abstract. The intended goals were to simplify the analysis and to deliver the key insights in a transparent and flexible manner. The obvious downside is that the assumed level of abstraction prevented any concrete application. We thus view the contribution of the present paper as a "proof of concept" and explore a few applications in companion work.

In Angeletos and Lian (2016a), we show how relaxing the common-knowledge assumption of an otherwise standard New-Keynesian model can attenuate the general-equilibrium effects of monetary policy that operate both on the demand side (the consumers) and the supply side (the firms). We then proceed to show how this attenuation helps resolve the forward guidance puzzles and the paradox of flexibility. In Angeletos and Lian (2016c), we show how a similar methodological approach can help the RBC framework accommodate the popular notion—if not the apparent fact—that a drop in consumer spending, such as the one triggered by deleveraging, can trigger a recession. In Angeletos and Lian (2017), we study the implications of our insights for the validity of Ricardian Equivalence and for the effects of fiscal policy.

Compared to the framework used in the present paper, the aforementioned applications feature richer microfoundations, more realistic forward-looking behavior, and a multitude of GE effects. In so doing, they help deliver concrete policy lessons that are not possible in the present paper. They also help build tighter, and context-specific, mappings between the theory and the data. They nevertheless all share the central theme of this paper, namely the attenuation of GE mechanisms and the reduction of the associated gap between microeconomic and macroeconomic effects. We hope that the combination of these applications with the broader perspective developed in this paper will inspire further research in the direction of revisiting the quantitative importance of general-equilibrium mechanisms and their policy implications, under appealing relaxations of either the REE solution concept or the standard common-knowledge assumptions.

Appendix: Proofs

Proof of Lemma 1. Consider a firm it hat trades in markets m and m' in, respectively, the morning and the afternoon. Because the technology is convex, the following conditions are necessary and sufficient for optimality:

$$\hat{E}_m\left[\frac{\partial K}{\partial q}(q_i, q_i^*; \theta_m)\right] = p_m, \tag{65}$$

$$\frac{\partial K}{\partial q^*}(q_i, q_i^*; \theta_m) = p_{m'}^*, \tag{66}$$

where $\hat{E}_m[\cdot]$ denotes subjective—potentially irrational—belief of agents who trade in market m in the morning (By Assumption 2, agents in the same market share the same belief in the morning),

Now, consider the optimal behavior of a household i that trades in markets m and m' in, respectively, the morning and the afternoon. It is pinned down by the solution to the following first-order conditions together with the budget constraint (3):[40]

$$\hat{E}_m\left[\frac{\partial U}{\partial c}(c_i, c_i^*; \theta_m)\right] = p_m, \tag{67}$$

$$\frac{\partial U}{\partial c^*}(q_i, q_i^*; \theta_m) = p_{m'}^*, \tag{68}$$

Similarly as the main text, with abuse of notation, we re-interpret all the variables as log-deviations from a symmetric steady state in which all marketplaces have the same fundamentals, and we work with the log-linearized demand and supply system. In effect, this means that we henceforth work with the log-linear approximation of the original model.

Solving and (log-linearizing) conditions (66) and (68), we can find linear functions D^* and S^* such that

$$c_i^* = D^*\left(c_i, p_{m'}^*, \theta_m\right) \qquad \text{and} \qquad q_i^* = S^*\left(q_i, p_{m'}^*, \theta_m\right).$$

By individual rationality, we can then substitute the previous condition into conditions (65) and (67). We can find linear functions D and S such that

$$c_m = c_i = D\left(p_m, \hat{E}_m[p_{m'}^*], \theta_m\right) \qquad \text{and} \qquad q_m = q_i = S\left(p_m, \hat{E}_m[p_{m'}^*], \theta_m\right),$$

where we use the fact that consumers (firms) of any given marketplace m are identical in the morning.

Now let us consider the demand and supply in any afternoon market m. Note that the demand in afternoon market m has two components: one reflecting the agents who were in this market from the morning; and another reflecting the agents who were relocated from other markets. The former have mass ρ and their demand is given by $D^*\left(c_m, p_m^*, \theta_m\right)$; the latter have mass $1-\rho$ and their average demand is given $\int D^*\left(c_{m'}, p_m^*, \theta_{m'}\right) dm' = D^*\left(\bar{c}, p_m^*, \bar{\theta}\right)$. As a result,

$$c_m^* = \rho D^*\left(c_m, p_m^*, \theta_m\right) + (1-\rho) D^*\left(\bar{c}, p_m^*, \bar{\theta}\right).$$

The same logic applies on the supply side.

Proof of Lemma 2. From Assumption 4, there exists a unique linear function $\tilde{P}(p^*, \theta)$ such that, for all (p^*, θ),

$$N\left(\tilde{P}(p^*, \theta), p^*, \theta\right) = 0.$$

[40] These conditions are both necessary and sufficient thanks to the concavity of U.

As a result, from condition (6), we know $p_m = \tilde{P}\left(\hat{E}_m[p^*_{m'}], \theta_m\right)$ clears the market in the morning. The market clearing quantity is then given by

$$c_m = q_m = D\left(\tilde{P}\left(\hat{E}_m[p^*_{m'}], \theta_m\right), \hat{E}_m[p^*_{m'}], \theta_m\right) \equiv \tilde{Q}\left(\hat{E}_m[p^*_{m'}], \theta_m\right).$$

Proof of Lemma 3. From conditions (**??**) and (8), we know the average market clearing price, $\bar{p}^*$, in the afternoon is given by $\bar{p}^* = P^*(\bar{q}, \bar{\theta})$. The quantity of afternoon goods is then given by

$$\begin{aligned}\bar{q}^* &= \int \left(\rho D^*\left(q_m, p^*_m, \theta_m\right) + (1-\rho) D^*\left(\bar{q}, p^*_m, \bar{q}\right)\right) dm \\ &= D^*\left(\bar{q}, \bar{p}^*, \bar{\theta}\right) = D^*\left(\bar{q}, P^*(\bar{q}, \bar{\theta}), \bar{\theta}\right) \equiv Q^*(\bar{q}, \bar{\theta}).\end{aligned}$$

Now we consider market clearing in a particular, marketplace m in the afternoon. If $p^*_m = \rho P^*\left(q_m, \theta_m\right) + (1-\rho)\,\bar{p}^*$, because N^* is linear, from condition (7), we have

$$\begin{aligned}n^*_m &= N^*\left(\rho q_m + (1-\rho)\,\bar{q}, \rho P^*\left(q_m, \theta_m\right) + (1-\rho)\,\bar{p}^*, \rho\theta_m + (1-\rho)\,\bar{\theta}\right). \\ &= \rho N^*\left(q_m, P^*\left(q_m, \theta_m\right), \theta_m\right) + (1-\rho)\, N^*\left(\bar{q}, \bar{p}^*, \bar{\theta}\right) \\ &= 0.\end{aligned}$$

From Assumption 5, we know such p^*_m is the only price that will clear afternoon goods market in marketplace m. Finally,

$$\begin{aligned}q^*_m &= \rho D^*\left(q^*_m, p^*_m, \theta_m\right) + (1-\rho)\, D^*\left(\bar{q}, p^*_m, \bar{\theta}\right) \\ &= \rho D^*\left(q_m, P^*\left(q_m, \theta_m\right), \theta_m\right) + (1-\rho)\, D^*\left(\bar{q}, P^*(\bar{q}, \bar{\theta}), \bar{\theta}\right) \\ &= \rho Q^*(q_m, \theta_m) + (1-\rho)\, Q^*\left(\bar{q}, \bar{p}^*\right),\end{aligned}$$

where Q^* is defined above in the proof of this Lemma.

Proof of Lemma 4. Using condition (10), Lemma 2 and the linearity of $\tilde{Q}$, we have

$$q_m = \rho^2 \tilde{Q}\left(P^*\left(q_m, \theta_m\right), \theta_m\right) + \left(1-\rho^2\right)\tilde{Q}\left(\mathbb{E}\left[\bar{p}^*\right], \theta_m\right).$$

Suppose $\rho^2\tilde{\alpha} = \rho^2 \frac{\partial P^*}{\partial q}\frac{\partial \tilde{Q}}{\partial p^*} \neq 1$. The above fixed point problem with respect to q_m has a unique solution:

$$q_m = Q(\mathbb{E}\left[\bar{p}^*\right], \theta_m)\ \forall m,$$

where $Q\left(p^*, \theta\right)$ is a linear function such that, for all $\left(p^*, \theta\right)$,

$$Q\left(p^*, \theta\right) = \rho^2 \tilde{Q}\left(P^*\left(Q\left(p^*, \theta\right), \theta\right), \theta\right) + \left(1-\rho^2\right)\tilde{Q}\left(p^*, \theta\right). \tag{69}$$

Similarly, using condition (10), Lemma 2 and the linearity of $\tilde{P}$, we have

$$
\begin{aligned}
p_m &= \tilde{P}\left(\rho^2 P^*\left(q_m, \theta_m\right) + \left(1-\rho^2\right)\mathbb{E}\left[\bar{p}^*\right], \theta_m\right) \\
&= \tilde{P}\left(\rho^2 P^*\left(Q(\mathbb{E}\left[\bar{p}^*\right], \theta_m), \theta_m\right) + \left(1-\rho^2\right)\mathbb{E}\left[\bar{p}^*\right], \theta_m\right) \\
&\equiv P\left(\mathbb{E}\left[\bar{p}^*\right], \theta_m\right).
\end{aligned} \tag{70}
$$

By aggregation, we then have

$$
\bar{q} = Q(\mathbb{E}\left[\bar{p}^*\right], \bar{\theta}) \qquad \text{and} \qquad \bar{p} = P(\mathbb{E}\left[\bar{p}^*\right], \bar{\theta}).
$$

Proof of Proposition 1. Let me first prove condition (13) in the main text. First, as $\mathcal{T}\left(p^*, \theta\right) \equiv P^*\left(Q\left(p^*, \theta\right), \theta\right)$ for all (p^*, θ), we have

$$
\frac{\partial \mathcal{T}}{\partial p^*} = \frac{\partial P^*}{\partial q}\frac{\partial Q}{\partial p^*}. \tag{71}
$$

Then, from the definition of Q, condition (69), we have

$$
\frac{\partial Q}{\partial p^*} = \frac{1-\rho^2}{1-\tilde{\alpha}\rho^2}\frac{\partial \tilde{Q}}{\partial p^*}. \tag{72}
$$

Together with condition (71), we have

$$
\frac{\partial \mathcal{T}}{\partial p^*} = \frac{1-\rho^2}{1-\tilde{\alpha}\rho^2}\frac{\partial P^*}{\partial q}\frac{\partial \tilde{Q}}{\partial p^*} = \frac{1-\rho^2}{1-\tilde{\alpha}\rho^2}\tilde{\alpha}.
$$

This finishes the proof of condition (13) in the main text. Together with Assumption 6 and the fact$\rho \in [0, 1)$, we have

$$
-1 < \alpha = \frac{1-\rho^2}{1-\tilde{\alpha}\rho^2}\tilde{\alpha} < 1.
$$

As a result, $\mathcal{T}$ is a contraction mapping. The unique solution of condition (12) can then be represented as

$$
\mathbb{E}[\bar{p}^*] = \mathcal{P}(\bar{\theta}),
$$

where $\mathcal{P}(\theta)$ is a linear function such that $\mathcal{P}(\theta) = \mathcal{T}\left(\mathcal{P}(\theta), \theta\right)$ for all θ. This finishes the proof of part (i) of the Proposition. Part (ii) can be then derived from Lemma 4.

Proof of Proposition 2. Use condition (15) and consider the change relative to its pre-shock value, we have

$$
\Delta q_m = Q(\mathcal{P}\left(\Delta\bar{\theta}\right), \Delta\theta_m); \tag{73}
$$

$$
\Delta\bar{q} = Q(\mathcal{P}\left(\Delta\bar{\theta}\right), \Delta\bar{\theta}) \equiv \epsilon^{Macro}\Delta\bar{\theta}, \tag{74}
$$

where

$$
\epsilon^{Macro} = \frac{\partial Q}{\partial \theta} + \frac{\partial Q}{\partial p^*}\frac{\partial \mathcal{P}}{\partial \theta}. \tag{75}
$$

Subtract condition (74) from condition (73), we have

$$
\Delta q_m = \Delta\bar{q} + \epsilon^{micro}\left(\Delta\bar{\theta}_m - \Delta\bar{\theta}\right),
$$

where

$$\epsilon^{micro} = \frac{\partial Q}{\partial \theta}. \tag{76}$$

Proof of Corollary 1. The Proposition follows from conditions (75) and (76), and the fact that $\int \delta_m dm = 1$.

Proof of Lemma 5. Let us prove first that $\mathcal{N}(p^*, \cdot)$ is decreasing in p^*. First, by definition of $\mathcal{N}$, we have $\frac{\partial \mathcal{N}}{\partial p^*} = \frac{\partial N^*}{\partial q}\frac{\partial Q}{\partial p^*} + \frac{\partial N^*}{\partial p^*}$. Then, by condition (72), we have

$$\frac{\partial \mathcal{N}}{\partial p^*} = \frac{1-\rho^2}{1-\tilde{\alpha}\rho^2}\frac{\partial N^*}{\partial q}\frac{\partial \tilde{Q}}{\partial p^*} + \frac{\partial N^*}{\partial p^*}. \tag{77}$$

Moreover, by taking partial derivatives with respect to q in the definition of P^*, condition (8), we have

$$\frac{\partial N^*}{\partial q} + \frac{\partial N^*}{\partial p^*}\frac{\partial P^*}{\partial q} = 0.$$

Together with condition (77) and Assumptions 5 and 6, we have $\frac{\partial \mathcal{N}}{\partial p^*} = \frac{\partial N^*}{\partial p^*}\left(1 - \frac{1-\rho^2}{1-\tilde{\alpha}\rho^2}\tilde{\alpha}\right) = \frac{\partial N^*}{\partial p^*}(1-\alpha) < 0$. This proves that $\mathcal{N}(p^*, \cdot)$ is decreasing in p^*.

Now we turn to the proof of Lemma 5. Note that $\mathcal{N}\left(\hat{P}^*(t), \bar{\theta}_{new}\right) = \mathcal{N}\left(\hat{P}^*(t), \bar{\theta}_{new}\right) - \mathcal{N}\left(\bar{p}^*_{new}, \bar{\theta}_{new}\right) = \frac{\partial \mathcal{N}}{\partial p^*}\left(\hat{P}^*(t) - \bar{p}^*_{new}\right)$. From condition (18), we have

$$\frac{d\left(\hat{P}^*(t) - \bar{p}^*_{old}\right)}{dt} = \frac{\partial \mathcal{N}}{\partial p^*}\left(\left(\hat{P}^*(t) - \bar{p}^*_{old}\right) - \left(\bar{p}^*_{new} - \bar{p}^*_{old}\right)\right) \quad \forall t \geq 0.$$

Together with $\hat{P}^*(0) = \bar{p}^*_{old}$, we have

$$\hat{P}^*(t) - \bar{p}^*_{old} = \left(\bar{p}^*_{new} - \bar{p}^*_{old}\right)\left(1 - e^{\frac{\partial \mathcal{N}}{\partial p^*}t}\right).$$

As a result,

$$w(T) = 1 - e^{\frac{\partial \mathcal{N}}{\partial p^*}T}. \tag{78}$$

Therefore, $w(T)$ is continuous strictly increasing in T, $w(0) = 0$ and $\lim_{T\to\infty} w(T) = 1$.

Proof of Lemma 6. Based on part (i) in Definition 1 and Lemma 2, we have

$$q_m = \rho^2 \tilde{Q}\left(P^*\left(q_m, \theta_m\right), \theta_m\right) + \left(1-\rho^2\right)\tilde{Q}\left(\hat{p}^*, \theta_m\right).$$

From Assumption (6), we know$\rho^2 \frac{\partial P^*}{\partial q}\frac{\partial \tilde{Q}}{\partial p^*} \neq 1$. As a result, the above fixed point problem has the unique solution:

$$q_m = Q\left(\hat{p}^*, \theta_m\right),$$

where Q is the same as the one in Lemma 4, defined in condition (69). Similarly, based on part (i) in Definition 1 and Lemma 2, we have

$$\begin{aligned} p_i &= \tilde{P}\left(\rho^2 P^*\left(q_m, \theta_m\right) + \left(1-\rho^2\right)\hat{p}^*, \theta_m\right) \\ &= \tilde{P}\left(\rho^2 P^*\left(Q(\hat{p}^*, \theta_m), \theta_m\right) + \left(1-\rho^2\right)\hat{p}^*, \theta_m\right) \\ &= P\left(\hat{p}^*, \theta_m\right), \end{aligned}$$

where P is the same as the one in Lemma 4, defined in condition (70).

Proof of Proposition 3. The Proposition follows directly from condition (21), Lemma 5, and the fact that $\epsilon^{micro} = \frac{\partial Q}{\partial \theta}$ and $\epsilon^{Macro} - \epsilon^{micro} = \frac{\partial Q}{\partial p^*}\frac{\partial \mathcal{P}}{\partial \theta}$.

Proof of Lemma 7. Lemmas 2 and 3 derived in Section 2 do not use the common knowledge assumption in Section 3, so they are still valid here. As a result, we have

$$\begin{aligned} q_m &= \rho \tilde{Q}\left(E_m\left[p_m^*\right], \theta_m\right) + (1-\rho)\tilde{Q}\left(E_m\left[\bar{p}^*\right], \theta_m\right) \\ &= \rho^2 \tilde{Q}\left(P^*\left(q_m, \theta_m\right), \theta_m\right) + \left(1-\rho^2\right)\tilde{Q}\left(E_m\left[\bar{p}^*\right], \theta_m\right). \end{aligned}$$

From Assumption (6), we know$\rho^2 \frac{\partial P^*}{\partial q}\frac{\partial \tilde{Q}}{\partial p^*} \neq 1$. As a result, the above fixed point problem has the unique solution:

$$q_m = Q\left(E_m\left[\bar{p}^*\right], \theta_m\right),$$

where Q is the same as the one in Lemma 4, defined in condition (69). Similarly, we have

$$\begin{aligned} p_m &= \tilde{P}\left(\rho^2 P^*\left(q_m, \theta_m\right) + \left(1-\rho^2\right)E_m\left[\bar{p}^*\right], \theta_m\right) \\ &= \tilde{P}\left(\rho^2 P^*\left(Q(E_m\left[\bar{p}^*\right], \theta_m), \theta_m\right) + \left(1-\rho^2\right)E_m\left[\bar{p}^*\right], \theta_m\right) \\ &= P\left(E_m\left[\bar{p}^*\right], \theta_m\right), \end{aligned}$$

where P is the same as the one in Lemma 4, defined in condition (70).

Proof of Lemma 8. The Lemma follows directly from taking average expectations of both sides of condition (28).

Proof of Corollary 2. Iterate condition (28), we have

$$\bar{p}^* = \frac{\partial \mathcal{T}}{\partial \theta}\sum_{h=1}^{\infty}\left(\frac{\partial \mathcal{T}}{\partial p^*}\right)^{h-1}\bar{E}^{h-1}\left[\bar{\theta}\right], \tag{79}$$

where for notational simplicity we let $\bar{E}^0\left[\bar{\theta}\right] = \bar{\theta}$.

By condition (13), we have $\alpha = \frac{\partial \mathcal{T}}{\partial p^*}$. Using condition (14) and taking derivatives with respect to $\bar{\theta}$ in condition (12), we have

$$\gamma = \frac{\partial \mathcal{P}}{\partial \theta} = \frac{\frac{\partial \mathcal{T}}{\partial \theta}}{1-\alpha}. \tag{80}$$

Substitute into condition 79, we have

$$\bar{p}^* = \gamma(1-\alpha)\sum_{h=1}^{\infty}\alpha^{h-1}\bar{E}^{h-1}\left[\bar{\theta}\right].$$

Taking average expectation of $\bar{p}^*$ based on the previous expression, Corollary 2 is then proved.

Proof of Lemma 9. Substitute condition (24) into condition (2), we have

$$\bar{E}\left[\bar{p}^*\right] = \gamma\bar{\theta}_{old} + \frac{\gamma(1-\alpha)\lambda}{1-\alpha\lambda}\Delta\bar{\theta} = \bar{p}^*_{old} + \pi(\lambda)\left(\bar{p}^*_{new} - \bar{p}^*_{old}\right),$$

where

$$\pi(\lambda) = \frac{(1-\alpha)\,\lambda}{1-\alpha\lambda}. \tag{81}$$

For $-1 < \alpha < 1$, π is continuous and strictly increasing in λ, with $\pi(0) = 0$ and $\pi(1) = 1$.

Proof of Proposition 4. The result directly follows from condition (26) and Lemma 9.

Proof of Proposition 5. Substituting condition (27) into condition (25), we have

$$q_m = \mathcal{BR}\left(\theta_m, E_m[\bar{\theta}], E_m[\bar{q}]\right),$$

where

$$\mathcal{BR}\left(\theta_m, \theta, q\right) = Q\left(P^*(q, \theta), \theta_m\right) \ \forall \left(\theta_m, \theta, q\right).$$

This proves part (i) of the Proposition. Taking partial derivatives with respect to q in the above condition, together with condition (71), we have

$$\frac{\partial \mathcal{BR}}{\partial q} = \frac{\partial P^*}{\partial q}\frac{\partial Q}{\partial p^*} = \frac{\partial \mathcal{T}}{\partial p^*} = \alpha.$$

This proves part (iii) of the Proposition.

Now let us try to prove part (ii) of the Proposition. If $\bar{q} = \mathcal{BR}\left(\bar{\theta}, \bar{E}[\bar{\theta}], \bar{E}[\bar{q}]\right) = Q\left(P^*(\bar{E}[\bar{q}], \bar{E}[\bar{\theta}]), \bar{\theta}\right)$ and $\bar{p}^* = P^*(\bar{q}, \bar{\theta})$, we have

$$\bar{E}[\bar{p}^*] = P^q(\bar{E}\left[\bar{q}\right], \bar{E}\left[\bar{\theta}\right]),$$

$$\bar{p}^* = P^*(\bar{q}, \bar{\theta}) = P^*(Q\left(\bar{E}[\bar{p}^*], \bar{\theta}\right), \bar{\theta}) = \mathcal{T}(\bar{E}[\bar{p}^*], \bar{\theta}).$$

Conversely, if $\bar{p}^* = \mathcal{T}(\bar{E}[\bar{p}^*], \bar{\theta})$ and $\bar{q} = Q(\bar{E}[\bar{p}^*], \bar{\theta})$, we have

$$P^*\left(\bar{q}, \bar{\theta}\right) = P^*\left(Q(\bar{E}[\bar{p}^*], \bar{\theta}), \bar{\theta}\right) = \mathcal{T}(\bar{E}[\bar{p}^*], \bar{\theta}) = \bar{p}^*.$$

As a result,

$$P^*\left(\bar{E}[\bar{q}], \bar{E}[\bar{\theta}]\right) = \bar{E}[\bar{p}^*].$$

Finally, from $\bar{q} = Q(\bar{E}[\bar{p}^*], \bar{\theta})$, we have

$$\bar{q} = Q\left(P^*(\bar{E}[\bar{q}], \bar{E}[\bar{\theta}]), \bar{\theta}\right) = \mathcal{BR}\left(\bar{\theta}, \bar{E}[\bar{\theta}], \bar{E}[\bar{q}]\right).$$

This finishes the proof.

Proof of Proposition 6. From conditions (32) and (34), we have

$$\frac{\epsilon^{Macro}}{\epsilon^{micro}} = \frac{1}{1-\alpha};$$

$$\frac{\epsilon^{Inc}}{\epsilon^{Macro}} = \frac{\epsilon^{micro} + \pi(\lambda)\frac{\alpha}{1-\alpha}\epsilon^{micro}}{\frac{1}{1-\alpha}\epsilon^{micro}} = \frac{1 + \pi(\lambda)\frac{\alpha}{1-\alpha}}{\frac{1}{1-\alpha}} = (1-\alpha) + \frac{(1-\alpha)\,\alpha\lambda}{1-\alpha\lambda} = \frac{1-\alpha}{1-\alpha\lambda}.$$

The Proposition then follows.

Proof of Proposition 7. To prove part (i), note that for any $T \in (0, \infty)$, there exists a unique $\lambda \in (0,1)$ such that $w(T) = \pi(\lambda)$. This is because (i) w is continuous and strictly increasing in T, with $w(0) = 0$ and $\lim_{T\to\infty} w(T) = 1$ and (ii) π is continuous and strictly increasing in λ, with $\pi(0) = 0$ and $\pi(1) = 1$. The theorem then directly follows conditions (19), (22), (31) and (32).

To prove part (ii), similarly, note that for any $\lambda \in (0,1)$, there exists a unique $T \in (0,\infty)$ such that $\pi(\lambda) = w(T)$. The theorem then directly follows conditions (19), (22), (31) and (32).

Proof of Lemma 10. Because $\mathcal{T}^0(\bar{p}^*_{old}, \bar{\theta}_{new}) = \bar{p}^*_{old}$, we have $g_0 = 0$. Now proceed by induction. Suppose for $k \geq 0$, we have $\hat{p}^*_k = \bar{p}^*_{old} + g_k \left(\bar{p}^*_{new} - \bar{p}^*_{old}\right)$. We have

$$\begin{aligned}
\hat{p}^*_{k+1} &= \mathcal{T}^{k+1}(\bar{p}^*_{old}, \bar{\theta}_{new}) \\
&= \mathcal{T}(\hat{p}^*_k, \bar{\theta}_{new}) \\
&= \bar{p}^*_{old} + \frac{\partial \mathcal{T}}{\partial p^*} g_k \left(\bar{p}^*_{new} - \bar{p}^*_{old}\right) + \frac{\partial \mathcal{T}}{\partial \theta}\left(\bar{\theta}_{new} - \bar{\theta}_{old}\right) \\
&= \bar{p}^*_{old} + \left(\alpha g_k + (1-\alpha)\right)\left(\bar{p}^*_{new} - \bar{p}^*_{old}\right),
\end{aligned}$$

where the last equation uses the fact that

$$\bar{p}^*_{new} - \bar{p}^*_{old} = \frac{\partial \mathcal{P}}{\partial \theta}\left(\bar{\theta}_{new} - \bar{\theta}_{old}\right) = \frac{\frac{\partial \mathcal{T}}{\partial \theta}}{1-\alpha}\left(\bar{\theta}_{new} - \bar{\theta}_{old}\right), \tag{82}$$

according to condition (80). As a result, we have

$$\hat{p}^*_{k+1} = \bar{p}^*_{old} + g_{k+1}\left(\bar{p}^*_{new} - \bar{p}^*_{old}\right),$$

where

$$g_{k+1} = \alpha g_k + (1-\alpha). \tag{83}$$

This proves part (i) in the Lemma 10.

From condition (83) and the fact that $g_0 = 0$, we have, for $k \geq 0$,

$$g_k = (1-\alpha)\left(\frac{1-\alpha^k}{1-\alpha}\right) = 1 - \alpha^k. \tag{84}$$

From this formula, it is easy to see that $\lim_{k\to\infty} g_k = 1$. Moreover, if $\alpha > 0$, the sequence is strictly increasing and bounded between 0 and 1. Finally, if instead $\alpha < 0$, this sequence is non-monotone, with $g_k < 1$ whenever k is even and $g_k > 1$ whenever k is odd.

Proof of Proposition 8. Similarly as the derivation of Lemma 6, we have

$$q_m = Q\left(\hat{p}^*, \theta_m\right) \qquad \text{and} \qquad \bar{q} = P\left(\hat{p}^*, \bar{\theta}\right).$$

Then, this Proposition follows directly from Lemma **10.**

Proof of Proposition 9. Similarly as the derivation of Lemma 6, based on part (ii) of the Definition 4, we have, for all $k \geq 0$,

$$q_m = Q\left(\hat{E}_m[\bar{p}^*], \theta_m\right) \qquad \text{and} \qquad \bar{q} = P\left(\hat{E}_m[\bar{p}^*], \bar{\theta}\right),$$

where $\hat{E}_m[\bar{p}^*] = P^*(\mathcal{BR}^k(\bar{\theta}_{new}, \bar{\theta}_{new}, \bar{q}_{old}), \bar{\theta}_{new}) = P^*(\hat{q}_k, \bar{\theta}_{new})$. Note that, for $k \geq 1$, we have

$$\begin{aligned}
P^*(\hat{q}_k, \bar{\theta}_{new}) &= P^*(\mathcal{BR}\left(\bar{\theta}_{new}, \bar{\theta}_{new}, \hat{q}_{k-1}\right), \bar{\theta}_{new}) \\
&= P^*(Q\left(P^*(\hat{q}_{k-1}, \bar{\theta}_{new}), \bar{\theta}_{new}\right), \bar{\theta}_{new}) \\
&= \mathcal{T}(P^*(\hat{q}_{k-1}, \bar{\theta}_{new}), \bar{\theta}_{new}) \\
&= \cdots \\
&= \mathcal{T}^k(P^*(\hat{q}_0, \bar{\theta}_{new}), \bar{\theta}_{new}) \\
&= \mathcal{T}^k(P^*(\bar{q}_{old}, \bar{\theta}_{new}), \bar{\theta}_{new}). \qquad (85)
\end{aligned}$$

Now we consider two cases.

(i) If $P^*(\bar{q}, \bar{\theta})$ is invariant to $\bar{\theta}$, we have $P^*(\bar{q}_{old}, \bar{\theta}_{new}) = \bar{p}^*_{old}$. As a result, $\hat{E}_m[\bar{p}^*] = \mathcal{T}^k(\bar{p}^*_{old}, \bar{\theta}_{new})$. This is exactly the conjecture in condition (36) in the definition of Cobweb(k) solution. The level-k solution and the Cobweb(k) solution impose the same price conjectures and give rise to the same observables.

(ii) If we modify the Cobweb solution concept so that the initial price conjecture is given by $\hat{p}_0 = P^*(\bar{q}_{old}, \bar{\theta}_{new})$. Then the RHS of the the conjecture in condition (36) in the definition of Cobweb(k) solution becomes exactly $\mathcal{T}^k(P^*(\bar{q}_{old}, \bar{\theta}_{new}), \bar{\theta}_{new}$ $P^*(\hat{q}_k, \bar{\theta}_{new})$. As a result, the level-k solution and the Cobweb(k) solution impose the same price conjectures and give rise to the same observables.

Proof of Corollary 3. The result follows from Proposition 8 and case (i) in the proof of Proposition 9.

Proof of Proposition 10. From conditions (5) and (82), we have, for all $t \geq 0$,

$$\begin{aligned}
\frac{d\left(\hat{P}^*(t) - \bar{p}^*_{old}\right)}{dt} &= \left(\frac{\partial \mathcal{T}}{\partial p^*} - 1\right)\left(\hat{P}^*(t) - \bar{p}^*_{old}\right) + \frac{\partial \mathcal{T}}{\partial \theta}\left(\bar{\theta}_{new} - \bar{\theta}_{old}\right) \\
&= -(1-\alpha)\left(\left(\hat{P}^*(t) - \bar{p}^*_{old}\right) - \left(\bar{p}^*_{new} - \bar{p}^*_{old}\right)\right).
\end{aligned}$$

Together with $\hat{P}^*(0) = \bar{p}^*_{old}$, we have

$$\hat{P}^*(t) - \bar{p}^*_{old} = \left(\bar{p}^*_{new} - \bar{p}^*_{old}\right)\left(1 - e^{-(1-\alpha)t}\right),$$

As a result, we have, in the reflective equilibrium economy,

$$\hat{p}^* = \bar{p}^*_{old} + w_{ref}(T)\left(\bar{p}^*_{new} - \bar{p}^*_{old}\right), \qquad (86)$$

in which $w_{ref}(T) = \left(1 - e^{-(1-\alpha)T}\right)$. Together with Lemma 5 and condition (78), we have, for any $T \geq 0$, there exists a T' such that the level-T reflective equilibrium conjecture $\hat{p}^*$ coincides with the Tâtonnement(T') conjecture $\hat{p}^*$, and vice versa. Given part (i) of both Definition (5) and Definition (1), the level-T reflective equilibrium also shares the same economic outcomes with the Tâtonnement(T') solution.

Proof of Lemma 11. The result follows directly from conditions (47), (48), (49) and the fact that $\int \kappa_{m,t} dm = \int k_{m,t-1} dm$.

Proof of Lemma 12. Plugging conditions (50) and (53) into household's FOC, condition (44), we have

$$k_{i,t} = \theta_m + \frac{1}{\phi}\left(\left(1 - \frac{1}{\chi\sigma}\right)\left(-\rho^2\chi k_{i,t} + (1-\rho^2)\hat{E}_{m,t}[\bar{p}_{m+1}]\right) - \frac{1}{\sigma}k_{i,t}\right),$$

where i is any household, $m = M(i,t)$ and $m' = M(i,t+1)$. Collecting terms and using the fact that $k_{i,t} = k_{m,t}$, we have[41]

$$k_{m,t} = \frac{\sigma\phi}{\sigma\phi + (\sigma\chi - 1)\rho^2 + 1}\theta_m + \frac{\left(\sigma - \frac{1}{\chi}\right)(1-\rho^2)}{\sigma\phi + (\sigma\chi - 1)\rho^2 + 1}\hat{E}_{m,t}[\bar{p}_{t+1}].$$

Proof of Proposition 11. Let $\epsilon^{micro} = \frac{\partial K}{\partial \theta}$ and $\epsilon^{Macro} = \frac{\partial K}{\partial \theta} + \frac{\partial K}{\partial P}\frac{\partial \mathcal{P}}{\partial \theta}$. The result then follows from condition (56) directly.

Proof of Lemma 13. From condition (57), we have

$$\hat{P}(t) - \bar{p}_{old} = (\bar{p}_{new} - \bar{p}_{old})\left(1 - e^{\frac{\partial \mathcal{N}}{\partial p}t}\right).$$

As a result,

$$w_t = 1 - e^{\frac{\partial \mathcal{N}}{\partial p}f(t)}$$

satisfies $\hat{p}_{t+1} = \bar{p}_{old} + w_t(\bar{p}_{new} - \bar{p}_{old})$ for all t. Because $\frac{\partial \mathcal{N}}{\partial p} < 0$, $f(t) > 0$ and $f(t)$ is strictly increasing, we have $w_t \in (0,1)$ for all t and w_t is strictly increasing in t.

Proof of Proposition 12. The result follows directly from condition (59), Lemma 12 and the definition of ϵ^{micro} and ϵ^{Macro} in the proof of Lemma 11.

Proof of Lemma 14. As discussed in the main text, starting $t \geq 1$, it is *as if* agents each marketplace exchanges its information with another randomly matched marketplace. As a result, at period $t = 1$, agents in any marketplace m will have two signals about the aggregate shock $\Delta\bar{\theta}$,

$$s_m = \Delta\bar{\theta} + v_m \quad \text{and} \quad s_{m'} = \Delta\bar{\theta} + v_{m'},$$

where m' is the marketplace with which m is matched at period t, v_m and $v_{m'}$ is an idiosyncratic noise term, drawn i.i.d. from $N(0, \sigma_v^2)$. Similarly, at period $t = 2$, agents in any marketplace m will have 4 signals about $\Delta\bar{\theta}$. Two of them are the signals they already receive at $t = 1$. The other two are new, from the new marketplace with which m is matched at period $t = 2$. By induction, it is as if, at period $t \geq 0$, each marketplace m have a total of 2^t signals distributed as

$$s_i = \Delta\bar{\theta} + v_i,$$

where v_i is an idiosyncratic noise term, drawn from $N(0, \sigma_v^2)$. As we have a continuum of markets but discrete time, the probability of a marketplace receives a "repetitive" signals through the matching place is always 0, we can view the 2^t signals as i.i.d. As a result, we have

$$\bar{E}_t[\Delta\bar{\theta}] = \lambda_t\Delta\bar{\theta} \text{ and } \bar{E}_t[\bar{\theta}] = \bar{\theta}_{old} + \lambda_t\Delta\bar{\theta} \;\; \forall t \geq 0, \tag{87}$$

[41] Note that, the denominator, $\sigma\phi + (\sigma\chi - 1)\rho^2 + 1 = \sigma\phi + \sigma\chi\rho^2 + 1 - \rho^2 > 0$.

where $\lambda_t = \frac{2^{2t}\sigma_v^{-2}}{2^{2t}\sigma_v^{-2}+\sigma_\theta^{-2}}$.

Iterating condition (87) by taking average expectations of both sides, we have

$$\bar{E}_t^h[\Delta\bar{\theta}] = \lambda_t^h \Delta\bar{\theta} \text{ and } \bar{E}_t^h[\bar{\theta}] = \bar{\theta}_{old} + \lambda_t^h \Delta\bar{\theta} \quad \forall t \geq 0 \text{ and } h \geq 1.$$

Proof of Lemma 15. Using condition (62) together with Lemma 14, we have

$$\bar{E}_t\left[\bar{p}_{t+1}\right] = \gamma\bar{\theta}_{old} + \frac{\gamma\left(1-\alpha\right)\lambda}{1-\alpha\lambda}\Delta\bar{\theta} = \bar{p}_{old} + \pi\left(\lambda_t\right)\left(\bar{p}_{new} - \bar{p}_{old}\right),$$

where

$$\pi(\lambda) = \frac{\left(1-\alpha\right)\lambda}{1-\alpha\lambda}. \tag{88}$$

For $-1 < \alpha < 1$, π is continuous and strictly increasing in λ, with $\pi(0) = 0$ and $\pi(1) = 1$.

Proof of Proposition 13. The result follows directly from aggregating condition (61) and Lemma 15.

References

Abreu, Dilip and Markus K Brunnermeier. (2003). "Bubbles and Crashes." *Econometrica* 71 (1):173–204.

Acemoglu, Daron, David Autor, David Dorn, Gordon H Hanson, and Brendan Price. (2016). "Import competition and the great US employment sag of the 2000s." *Journal of Labor Economics* 34 (S1):S141–S198.

Akerlof, George A and Janet L Yellen. (1985)a. "Can Small Deviations from Rationality Make Significant Differences to Economic Equilibria?" *The American Economic Review* 75 (4):708–720.

———. (1985)b. "A Near-rational Model of the Business Cycle, with Wage and Price inertia." *The Quarterly Journal of Economics* 100 (Supplement):823–838.

Allen, Franklin, Stephen Morris, and Hyun Song Shin. (2006). "Beauty Contests and Iterated Expectations in Asset Markets." *Review of financial Studies* 19 (3):719–752.

Angeletos, George-Marios and Jennifer La'O. (2010). "Noisy Business Cycles." In *NBER Macroeconomics Annual 2009, Volume 24*. University of Chicago Press, 319–378.

———. (2013). "Sentiments." *Econometrica* 81 (2):739–779.

Angeletos, George-Marios and Chen Lian. (2016)a. "Forward Guidance without Common Knowledge." *NBER Working Paper No. 22785* .

———. (2016)b. "Incomplete Information in Macroeconomics: Accommodating Frictions in Coordination." *Handbook of Macroeconomics* 2:forthcoming.

———. (2016)c. "A (Real) Theory of Keynesian Multipliers." *Work in progress* .

———. (2017). "Fiscal Policy without Common Knowledge." *Work in progress* .

Angeletos, George-Marios and Alessandro Pavan. (2007). "Efficient Use of Information and Social Value of Information." *Econometrica* 75 (4):1103–1142.

Autor, David, David Dorn, and Gordon H Hanson. (2013). "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *The American Economic Review* 103 (6):2121–2168.

Bacchetta, Philippe and Eric van Wincoop. (2006). "Can Information Heterogeneity Explain the Exchange Rate Determination Puzzle?" *The American Economic Review* 96 (3):552.

Beraja, Martin, Erik Hurst, and Juan Ospina. (2016). "The Aggregate Implications of Regional Business Cycles." *NBER Working Paper No. 21956* .

Bergemann, Dirk and Stephen Morris. (2013). "Robust Predictions in Games with Incomplete Information." *Econometrica* 81 (4):1251–1308.

Caballero, Ricardo J and Eduardo MRA Engel. (1999). "Explaining Investment Dynamics in US Manufacturing: a Generalized (S, s) Approach." *Econometrica* 67 (4):783–826.

Cagan, Phillip. (1956). "The Monetary Dynamics of Hyperinflation." *Studies in the Quantity Theory of Money* :25–117.

Caplin, Andrew S and Daniel F Spulber. (1987). "Menu Costs and the Neutrality of Money." *The Quarterly Journal of Economics* 102 (4):703–725.

Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber. (2011). "Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins." *The American Economic Review* 101 (3):471–475.

———. (2013). "Does Indivisible Labor Explain the Difference between Micro and Macro Elasticities? A Meta-analysis of Extensive Margin Elasticities." *NBER Macroeconomics Annual* 27 (1):1–56.

Crawford, Vincent P, Miguel A Costa-Gomes, and Nagore Iriberri. (2013). "Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications." *Journal of Economic Literature* 51 (1):5–62.

Duffie, Darrell, Nicolae Garleanu, and Lasse Heje Pedersen. (2005). "Over-the-Counter Markets." *Econometrica* 73 (6):1815–1847.

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

Evans, George W and Garey Ramey. (1992). "Expectation Calculation and Macroeconomic Dynamics." *American Economic Review* 82 (1):207–24.

———. (1995). "Expectation Calculation, Hyperinflation and Currency Collapse." *The New Macroeconomics: Imperfect Markets and Policy Effectiveness* :307–336.

Farhi, Emmanuel and Iván Werning. (2017). "Monetary Policy, Bounded Rationality, and Incomplete Markets." *NBER Working Paper No. 23281* .

Gabaix, Xavier. (2014). "A Sparsity-Based Model of Bounded Rationality." *The Quarterly Journal of Economics* 129 (4):1661–1710.

———. (2016). "A Behavioral New Keynesian Model." *NBER Working Paper No. 22954* .

Garcıa-Schmidt, Mariana and Michael Woodford. (2015). "Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis." *NBER Working Paper No. 21614* .

Golosov, Mikhail and Robert E Lucas Jr. (2007). "Menu Costs and Phillips Curves." *Journal of Political Economy* 115 (2):171–199.

Grandmont, Jean-Michel. (1977). "Temporary General Equilibrium Theory." *Econometrica* 45 (3):535–72.

Guesnerie, Roger. (1992). "An Exploration of the Eductive Justifications of the Rational-expectations Hypothesis." *The American Economic Review* :1254–1278.

———. (2008). "Macroeconomic and Monetary Policies from the Eductive Viewpoint." *Central Banking, Analysis, and Economic Policies Book Series* 13:171–202.

Keane, Michael, Richard Rogerson et al. (2012). "Micro and Macro Labor Supply Elasticities: A Reassessment of Conventional Wisdom." *Journal of Economic Literature* 50 (2):464–76.

Lagos, Ricardo and Randall Wright. (2005). "A Unified Framework for Monetary Theory and Policy Analysis." *Journal of Political Economy* 113 (3):463–484.

Mackowiak, Bartosz and Mirko Wiederholt. (2009). "Optimal Sticky Prices under Rational Inattention." *American Economic Review* 99 (3):769–803. URL `http://www.aeaweb.org/articles.php?doi=10.1257/aer.99.3.769`.

Mankiw, N Gregory and Ricardo Reis. (2002). "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *Quarterly Journal of Economics* :1295–1328.

Mas-Collell, Andreu, Michael Whinston, and Jerry R Green. (1995). *Microeconomic Theory*. Oxford University Press.

Mian, Atif, Kamalesh Rao, and Amir Sufi. (2013). "Household Balance Sheets, Consumption, and the Economic Slump." *The Quarterly Journal of Economics* :qjt020.

Mian, Atif and Amir Sufi. (2012). "The Effects of Fiscal Stimulus: Evidence from the 2009 Cash for Clunkers Program." *The Quarterly Journal of Economics* :qjs024.

———. (2014). "What Explains the 2007–2009 Drop in Employment?" *Econometrica* 82 (6):2197–2223.

Morris, Stephen and Hyun Song Shin. (1998). "Unique Equilibrium in a Model of Self-fulfilling Currency Attacks." *American Economic Review* :587–597.

———. (2001). "Rethinking Multiple Equilibria in Macroeconomic Modeling." *NBER Macroeconomics Annual* 15:139–161.

———. (2002). "Social Value of Public Information." *The American Economic Review* 92 (5):1521–1534.

———. (2006). "Inertia of Forward-looking Expectations." *The American Economic Review* :152–157.

Myatt, David P and Chris Wallace. (2012). "Endogenous Information Acquisition in Coordination Games." *The Review of Economic Studies* 79 (1):340–374.

Nagel, Rosemarie. (1995). "Unraveling in Guessing Games: An Experimental Study." *The American Economic Review* 85 (5):1313–1326.

Nakamura, Emi and Jon Steinsson. (2014). "Fiscal Stimulus in a Monetary Union: Evidence from US regions." *The American Economic Review* 104 (3):753–792.

Nimark, Kristoffer. (2008). "Dynamic Pricing and Imperfect Common Knowledge." *Journal of Monetary Economics* 55 (2):365–382.

———. (2017). "Dynamic Higher Order Expectations." *Cornell Univeristy mimeo* .

Sims, Christopher A. (2003). "Implications of Rational Inattention." *Journal of Monetary Economics* 50 (3):665–690.

Stahl, Dale O. (1996). "Boundedly Rational Rule Learning in a Guessing Game." *Games and Economic Behavior* 16 (2):303–330.

Stahl, Dale O and Paul W Wilson. (1994). "Experimental Evidence on Players' Models of Other Players." *Journal of Economic Behavior & Organization* 25 (3).

Stahl Dale, O and Paul W Wilson. (1995). "On Players' Models of Other Players: Theory and Experimental Evidence." *Games and Economic Behavior* 10 (1):218–254.

Venkateswaran, Venky. (2014). "Heterogeneous Information and Labor Market Fluctuations." *NYU mimeo* .

Woodford, Michael. (2003). "Imperfect Common Knowledge and the Effects of Monetary Policy." *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps* .

———. (2013). "Macroeconomic Analysis Without the Rational Expectations Hypothesis." *Annual Review of Economics* 5:303–346.