Learning from House Prices: Amplification and Business Fluctuations

Ryan Chahrour* Gaetano Gaballo†

Boston College HEC Paris and CEPR

November 2, 2020

Abstract

We formalize the idea that house price changes may drive rational waves of optimism and pessimism in the economy. In our model, a house price increase caused by aggregate disturbances may be misinterpreted as a sign of higher local permanent income, leading households to demand more consumption and housing. Higher demand reinforces the initial price increase in an amplification loop that drives comovement in output, labor, residential investment, land prices, and house prices even in response to aggregate supply shocks. The qualitative implications of our otherwise frictionless model are consistent with observed business cycles and it can explain the economic impact of apparently autonomous changes in sentiment without resorting to non-fundamental shocks or nominal rigidity.

Keywords: demand shocks, house prices, imperfect information, animal spirits.

JEL Classification: D82, D83, E3.

*Department of Economics, Boston College, Chestnut Hill, MA 02467, U.S.A. Email: ryan.chahrour@bc.edu

†Department of Economic Sciences, HEC Paris, 1 rue de la Liberation, 78350 Jouy en Josas. Email: gaballo@hec.fr

"*Housing* is *the Business Cycle*."

Edward E. Leamer,
Jackson Hole Symposium, 2007

1 Introduction

House prices provide valuable information about ongoing changes in economic activity, both at the aggregate and regional levels.[1] Over the last half century in the US, real house prices and output have moved together at least half of the time (Figure 1). However, people have very different real-time information about these variables. Precise information about *local* house prices is readily available and relevant to individual choices, while the earliest measures of GDP are imprecise, released with delay, and may be less relevant to individual choices. For these reasons, people observing higher house prices may rationally become more optimistic about their own economic prospects. Through this learning channel, factors driving house price movements may also drive waves of economic optimism or pessimism.

This paper proposes a new model of housing's *informational* role in generating and amplifying demand-driven business fluctuations. The essence of the model is a price-optimism feedback channel: higher house prices beget economic optimism, which begets even higher house prices, and so on. Since any aggregate shock can activate this loop, price-quantity comovement can emerge in our model even in response to supply shocks. In this way, our learning channel offers a new source of amplification for fundamental shocks and blurs the traditional dichotomy between disturbances to supply and demand.

We embed our learning mechanism within a neoclassical model with housing. Households are located on islands and consume a traded consumption good and local housing. Traded consumption is produced using labor from all islands, while local housing is produced using land, local labor, and a traded productive factor (commodity good) whose supply is fixed. Local house prices can move either because of an increase in the future product of local labor, or because of a current aggregate disturbance to housing production.

Most fluctuations in local house prices are driven by local labor productivity, so people

[1] Leamer (2007) and Leamer (2015) make the point forcefully for aggregates, while Campbell and Cocco (2007) and Miller et al. (2011) provide evidence at the regional level.

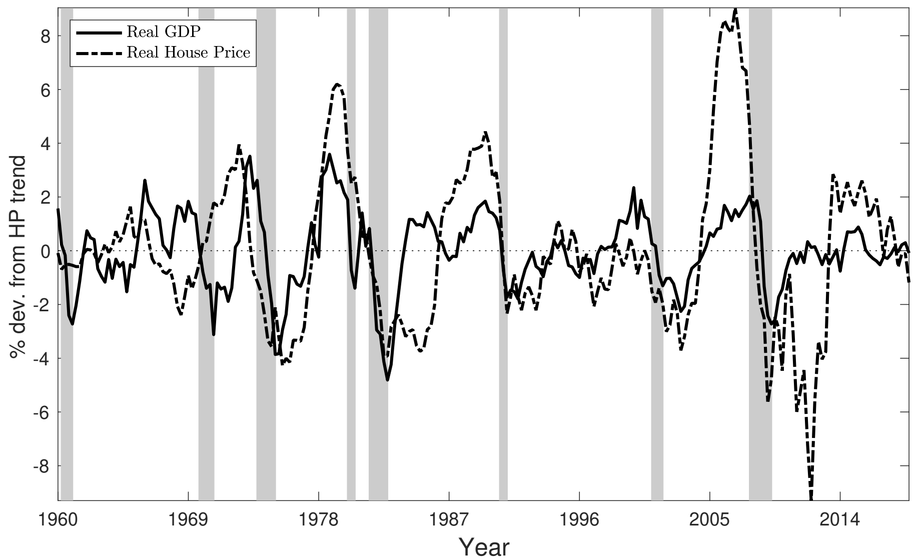

Figure 1: Real gross domestic product and the Shiller national house price index.

observing high house prices become optimistic about their own labor income prospects. However, a fall in the productivity (or availability) of the commodity factor also drives an increase in house prices across islands. In this case, the increase is misinterpreted by households as good news about future wages, increasing demand for both consumption and housing on all islands. Increasing aggregate demand further increases house prices, and consequently the price of the commodity factor, reinforcing the initial price increase. In equilibrium, what started as a small change in housing supply leads to an economy-wide increase in house prices, a boom in aggregate demand for consumption and housing, and a spike in the commodity input price.

For this increase in aggregate demand to affect real quantities, it must be associated with an intratemporal "labor wedge." We achieve this in our model by assuming that the household is split between shoppers, who make consumption and housing decisions based only on market prices, and worker-savers, who make labor supply and savings decisions based on full information. As the shopper becomes overly optimistic about wages, he spends more, thinking that the worker-saver is working less. However, since wages have not actually changed, the worker-saver is induced to work more to avoid an inefficient fall in household wealth. In this way, workers' optimal labor choices allow shoppers' elevated demand to be met with higher actual output and a boom in real quantities occurs.[2]

[2] Chahrour and Gaballo (2017) show that a similar wedge appears when some households are fully informed and others base all of their decisions — savings, labor, and consumption — just on market prices.

Our model of learning from prices has several features that make it an appealing model of the business cycle. First, our mechanism works in a flexible price model with competitive markets. This means that fluctuations in housing demand, and their real effects, are not driven by competitive or nominal frictions, or by suboptimal monetary policy. Indeed, our real economy can be interpreted as a monetary economy with a fixed nominal price level. Hence, our model aligns with recent experience in developed economies, where real fluctuations have coincided with small, largely acyclical fluctuations in inflation (e.g. Angeletos et al., 2020).

Second, the logic of our model extends to other sorts of macroeconomic fundamentals and to learning from any local price. Hence, the mechanism we propose can explain why business cycle comovements emerge irrespective of the particular type of shock hitting the economy. We show this generality by exploring an extension with an aggregate shock to consumption — rather than housing — productivity, but also refer the reader to earlier drafts of this paper that demonstrate how the mechanism works for shocks to the nominal money supply, and when learning occurs from local consumption prices.[3]

Finally, the signal structure faced by households is fully microfounded without introducing any extrinsic noise. Instead, we explicitly derive the house price signal as the outcome of competitive markets and show how fundamental shocks play the role of aggregate noise in people's inference. Thus, the model explains how people's beliefs become coordinated rather than assuming coordination, as in the literature on sunspots (e.g. Cass and Shell, 1983).

The fact that information comes from market prices, rather than from exogenously-specified signals, is crucial for our mechanism. First, it means that higher house prices can actually *spur* demand, for both consumption and housing. Indeed, learning from prices leads housing demand to be upward sloping in our model, leading to realistic comovement in house prices, housing investment, and non-house consumption.

Second, the feedback of the commodity price into local house prices allows the model to deliver strong amplification. For some calibrations, amplification can be so strong that aggregate prices and quantities exhibit sizable fluctuations in the limit of arbitrarily small aggregate shocks. To an econometrician, the fluctuations emerging at the limit of no aggregate

[3]In Chahrour and Gaballo (2017) people learn from the price of local consumption. There, we show that total factor productivity shocks can drive the business cycle and still be weakly correlated with business cycle variables, as found in the data (Angeletos et al., 2018, 2020).

shocks would appear to be driven by something akin to "animal spirits" (Shiller, 2007), "noise" (Gazzani, 2019), or "sentiment" (Angeletos and La'O, 2013; Benhabib et al., 2015).

After characterizing equilibrium in closed form, we examine the qualitative features of the economy. We show that the model implies positive comovement between output, employment, hours in the consumption and housing sectors, house prices, and land prices *for any calibration and any equilibrium* so long as aggregate shocks are small enough. Hence, the model provides a robust foundation for macroeconomic comovement across a wide range of variables.

We next enrich the model so that a portion of housing productivity is common knowledge. This allows the model to exhibit typical "supply-like" comovement in response to the common knowledge portion of the shock, while still experiencing "demand-like" fluctuations in response to the surprise component that blurs households' inference. A calibrated version of the extended model delivers qualitatively realistic (i.e. positive but imperfect) correlations among many real variables. Moreover, even though it has a unique equilibrium, the model both amplifies housing market fluctuations and generates strong fluctuations in consumption, which would disappear under full information. Indeed, amplification is strong enough that demand fluctuations dominate unconditional comovements even when the majority of productivity shocks are common knowledge.

We augment our discussion of real comovements with some non-structural evidence favoring house prices as the source of people's economic learning. For this, we use Michigan Survey of Consumer Expectations data to show that people's past house price experiences are a far better predictor of their forecasts of their own income than are people's reports about aggregate economic news that they have heard. Moreover, house price experiences modestly lead income expectations, a timing that is consistent with information flowing from house prices to income expectations. While this evidence is only suggestive, we think it indicates that our model can help guide more structural interpretations of expectations survey data.

Literature review

This paper follows an extensive literature that offers different foundations for business cycles caused by waves of economic optimism and pessimism. In this paper, changes in housing supply drive initial *consumer* optimism through their effects on house prices. The literature has considered other origins for waves of consumer optimism, including: Lorenzoni (2009) with news about future TFP; Ilut and Schneider (2014) with uncertainty shocks; and Angeletos and Lian (2020) with discount factor shocks. Several others have also modeled belief-driven fluctuations that originate on the part of firms or producers, including Angeletos and La'O (2009), Angeletos and La'O (2013) and Benhabib et al. (2015).

Unlike Lorenzoni (2009) and Ilut and Schneider (2014), the real effects of consumer optimism in our model, as in Angeletos and Lian (2020), do not rely on nominal frictions or suboptimal monetary policy. Angeletos and Lian (2020) show how a discount rate shock – a shock to the *intertemporal* margin – can be amplified when consumers' intertemporal substitution operates under imperfect information and aggregate supply is upward sloping in the real interest rate. By contrast, in our model, consumers' uncertainty maps into distortions to the *intratemporal* margin. As a result, people's correlated mistakes about private conditions can propagate, while aggregate *intertemporal* shocks affect only real interest rates, just as they would in a frictionless Real Business Cycle model (see section 5.1).

This paper shows that rational learning from prices can help explain business cycle and housing comovements, but price-based learning has a long tradition in macroeconomics and finance, starting with Lucas (1972) and Grossman and Stiglitz (1976, 1980). Recent examples in macroeconomics include Amador and Weill (2010), Benhima and Blengini (2020), Benhima (2019), Gaballo (2016, 2018), L'Huillier (2020), Nimark (2008) and Venkateswaran (2013). Several finance papers show that price-based learning can deliver asset price amplification or multiple equilibria, including Burguet and Vives (2000), Barlevy and Veronesi (2000), Albagli et al. (2014), Manzano and Vives (2011), and Vives (2014).

Among these papers, we are the first to show extreme amplification in limit cases of noisy rational expectations equilibria. This result is connected to the sentiment equilibria of Benhabib et al. (2015), a link that we explore in Section 5.3. Other papers have documented amplification when allowing for departures from rational expectations, including Eusepi and Preston (2011) and Hassan and Mertens (2017), and Adam et al. (2011) in the housing context.

Our theory is consistent with a range of empirical evidence on housing and the business cycle. Early housing macro models, like Davis and Heathcote (2005), struggle to explain price-quantity comovement and authors have introduced housing demand shocks to match these moments (e.g. Iacoviello and Neri, 2010). Though our model is close to Davis and Heathcote (2005), the learning in our model causes prices and quantities to positively comove.[4] Our model also qualitatively accounts for the high volatility of the price of land (Davis and Heathcote, 2007) and for its strong comovement with labor markets (Liu et al., 2016).

Our paper also contributes to a long debate about the nature and size of housing wealth effects. Frictionless models typically imply that house prices should have no causal impact on consumption (e.g. Buiter, 2010) but many empirical studies suggest otherwise. For example, Muellbauer and Murphy (1990) argue that the 1980's spike in UK consumption was driven by rising house prices, while King (1990), Pagano (1990), Attanasio and Weber (1994), and Attanasio et al. (2009) argue consumption and house prices reflected people's perceptions of permanent income. In our model, these competing views coexist: high house prices drive increased consumption not because consumers expect to sell their houses at the high price, but because consumers interpret them as signaling higher permanent income.

Evidence from cross-sectional studies is also largely consistent with our theory. For example, Campbell and Cocco (2007) find that a 1% increase in an individual's home value is associated with a 1.22% increase in their non-durable consumption in the UK, while Miller et al. (2011) find a positive effect of local house prices on metropolitan-level growth in the US. The recent studies by Mian et al. (2013) and Mian and Sufi (2014) also present evidence that falling house prices are associated with consumption reductions at the ZIP code level.

Other theoretical mechanisms for a direct consumption effect of house prices have been proposed in the literature, including borrowing constraints (Iacoviello, 2005) and wealth heterogeneity with incomplete markets (Berger et al., 2017; Kaplan et al., 2017). The learning channel we formalize offers a complimentary explanation. One difference is that our channel does not depend on actual new house sales or credit contracts, which might imply a longer delay between house prices and their effects on consumption.

[4] Recently, Nguyen (2018) and Fehrle (2019) have also proposed particular types of segmentation in capital markets as solutions to these comovement challenges.

2 A housing model with learning from prices

In this section, we present a simple real business cycle model with housing. We aim as much as possible to provide analytical results and make simplifying assumptions to this end. Most of these assumptions can be relaxed; we discuss when and how as we proceed.

2.1 Preferences and technology

The economy consists of a continuum of islands, indexed by $i \in (0,1)$. Each island is inhabited by a continuum of price-taking households who consume local housing and a traded numeraire consumption good. Households provide local labor which is used in the production of both goods. On each island, a mass of competitive construction firms combine local labor and land with a traded commodity good to construct new houses, while an aggregate consumption sector combines all islands' labor to produce the traded consumption good.

Households

The representative household on island i chooses consumption, labor supply, and savings in a risk-free nominal bond to maximize the utility function:

$$U_{i0} \equiv \sum_{t=0}^{\infty} \beta^t \left\{ \log(C_{it}^{\phi} \mathcal{H}_{it}^{1-\phi}) - vN_{it} \right\}. \tag{1}$$

In the utility function above, C_{it} denotes household i's consumption of the tradable consumption good, $\mathcal{H}_{it}$ measures the total quantity of housing consumed, and N_{it} is the household's supply of labor. The household discount factor is $\beta \in (0,1)$, the share of housing in the consumption basket is $\phi \in (0,1)$, and v parameterizes the household's disutility of labor.[5]

We assume that housing consumption is composed of a sequence of housing vintages, $\Delta_{i\tau|k}$, constructed at time k and combined according to the Cobb-Douglas aggregator

$$\mathcal{H}_{it} \equiv \prod_{k=-\infty}^{t} \Delta_{it|k}^{(1-\psi)\psi^{t-k}}, \tag{2}$$

where $\psi \in (0,1)$. This formulation for housing utility adds a realistic dimension to the model, since housing vintages can have very different characteristics and are not perfect substitutes.

[5] We allow for convex disutility of labor in the Appendix.

More importantly for our purposes, however, this formulation in conjunction with log-utility implies that every housing vintage has an additive-separable impact on intertemporal utility, allowing us to analyze the dynamic model in closed form.

Each vintage of housing depreciates at a constant rate $d \in (0, 1)$, so that

$$\Delta_{i\tau+1|k} = (1-d)\Delta_{i\tau|k}$$

for $\tau \geq k$ (while, of course, $\Delta_{i\tau|k} = 0$ for $\tau < k$). The aggregate housing stock, defined as $H_{it} = \sum_{k=-\infty}^{t} \Delta_{it|k}$, then evolves according to a standard equation,

$$H_{it} = \Delta_{it|t} + (1-d)\, H_{it-1}.$$

Housing consumption can now be written $\mathcal{H}_{it} = \Delta_{it|t}^{1-\psi}(1-d)\mathcal{H}_{it-1}^{\psi}$, which we use going forward.

The choices of the household are subject to the following budget constraint,

$$\mathcal{B}_{it} \equiv \frac{B_{it}}{R_t} + C_{it} + P_{it}\Delta_{it|t} - W_{it}N_{it} - B_{it-1} - \Pi_t^c - \Pi_{it}^h \leq 0 \tag{3}$$

for $t \in \{0, 1, 2...\}$ with $B_{i-1} = 0$. Household resources come from providing local labor at wage W_{it}, from past bond holdings, from profits Π_{it}^h of locally-owned housing firms, and from profits Π_t^c of the representative consumption firm, which is evenly held across islands. The household uses its funds to purchase numeraire consumption, to acquire new housing at price P_{it}, and to save in a zero-net-supply aggregate bond with a real risk-free return R_t. We denote the price of the local housing vintages as $P_{it|k}$ and define the price of the total housing stock as $P_{it}^H = \sum_{k=-\infty}^{t} P_{it|k}\Delta_{it|k}/H_{it}$.

Notice, however, that only the price of the current vintage, $P_{it} \equiv P_{it|t}$, appears in (3). This happens because the local household is the only potential buyer and seller of past vintages, meaning that trade in houses can never generate net resources for the island. For this reason, housing wealth is not wealth in the sense of Buiter (2010). The literature has proposed several strategies to break this irrelevance; our goal is to describe a potentially complementary channel through which house prices can have a causal effect on consumption.

Housing producers

House-producing firms construct new houses using a Cobb-Douglas technology,

$$\Delta_{it} = L_{it}^{1-\alpha} X_{it}^{\alpha}, \tag{4}$$

that combines land (L_{it}) with new residential structures (X_{it}) to generate new residential units $\Delta_{it} \equiv \Delta_{it|t}$. Residential structures have share $\alpha \in (0,1)$ and are produced, in turn, via a Cobb-Douglas production function

$$X_{it} = (N_{it}^h)^\gamma (e^{-\tilde{\zeta}_t} Z_{it})^{1-\gamma} \tag{5}$$

combining local labor, N_{it}^h, with a traded commodity, Z_{it}, with share parameter $\gamma \in (0,1)$.

The housing firm maximizes profits,

$$\Pi_{it}^h \equiv P_{it}\Delta_{it} - W_{it}N_{it}^h - Q_t\left(Z_{it} - Z\right) - V_{it}L_{it}$$

subject to (4) and (5). In the above, V_{it} is the local price of land, W_{it} is the price of local labor, and Q_t is the price of the commodity good. We assume that housing firms are endowed each period with Z units of the commodity good, which trades freely across islands and depreciates fully at the end of the period. Land supply is exogenous: each period a fixed amount of residential land — normalized to one — becomes available to housing producers on the island.[6] Without loss of generality, we assume that new land is endowed to local firms.

The only aggregate shock affecting our baseline economy is a shock to productivity of the commodity good, $\tilde{\zeta}_t$.[7] This shock evolves according to a random walk, $\tilde{\zeta}_t = \tilde{\zeta}_{t-1} + \zeta_t$, with i.i.d. innovation ζ_t distributed according to $N\left(0, \sigma_\zeta^2\right)$. We focus our presentation on this shock because it has no effect on consumption under full information. Still, other shocks could play a similar role: We consider an extension with an aggregate shock to consumption productivity in Section 5 and show that $\tilde{\zeta}_t$ is isomorphic to a shock to the endowment of Z in the Appendix B.2 (see Remark 3).

Consumption sector

The numeraire consumption good is traded freely across islands and is produced by a continuum of identical competitive firms. The representative consumption producer combines labor

[6]These assumptions do not imply that land supply grows over time. Provided an appropriate transformation of the depreciation rate, this formulation is equivalent to a model in which structures are placed on a fixed stock of land and existing land becomes free as those structures depreciate. See Davis and Heathcote (2005) for details.

[7]Notice that with our sign normalization in (5), a positive $\tilde{\zeta}$ corresponds to lower productivity.

from all sectors to maximize profits,

$$\Pi_t^c \equiv Y_t - \int W_{it} N_{it}^c di$$

subject to the production function,

$$Y_t = \left(\int e^{\tilde{\mu}_{it}/\eta} {N_{it}^c}^{1-\frac{1}{\eta}} di \right)^{\frac{1}{1-\frac{1}{\eta}}}. \tag{6}$$

The quantity of local labor used is denoted by N_{it}^c, and labor types can be substituted with elasticity $\eta > 0$. Island-specific labor productivity is a random walk, and evolves according to $\tilde{\mu}_{it} = \tilde{\mu}_{it-1} + \hat{\mu}_{it}$, where $\hat{\mu}_{it}$ is i.i.d. and drawn from the normal distribution $N(0, \hat{\sigma}_\mu)$.

Market clearing

Clearing in the local land and labor markets requires

$$L_{it} = 1 \quad \text{and} \quad N_{it} = N_{it}^c + N_{it}^h. \tag{7}$$

Per the discussion above, we omit market clearing conditions for all past housing vintages, since their trade is irrelevant at the island level. Finally, clearing in the aggregate markets for bonds, consumption, and the commodity good requires

$$Y_t = \int C_{it} di, \quad 0 = \int B_{it} di, \quad \text{and} \quad Z = \int Z_{it} di. \tag{8}$$

2.2 Timing and equilibrium

The only friction that we introduce is uncertainty in households' demand. To model this in a parsimonious way, we use the family metaphor also adopted by Angeletos and La'O (2009) and Amador and Weill (2010). The household is composed of two types of agents: a shopper, who uses household resources to buy consumption and housing, and a worker-saver, who chooses the number of hours to supply and the quantity of bonds to buy.

Both family member types act in the interest of the household, but they cannot pool their information within a time period. Hence, choices of Δ_{it} and C_{it} are conditioned on the information set of shoppers, while N_{it} and B_{it} are conditioned on the full information set of workers. Each period is composed of four stages:

1. The household splits into shoppers and worker-savers.

2. Shocks realize, namely future local productivity innovations, $\{\hat{\mu}_{i,t+1}\}_{i\in(0,1)}$, and the current aggregate shock, ζ_t. The "best available" information set, $\Omega_{it} \equiv \{\{\hat{\mu}_{i,\tau}\}_0^{t+1}, \{\zeta_\tau\}_0^t\}$, is observed by firms and worker-savers on each island, but not shoppers.

3. Production and trade take place. Shoppers and workers make their choices based on the information they have, which includes the competitive equilibrium prices in the markets in which they trade. Firms make production choices in light of realized productivity and input prices; and all markets clear.

4. Family members share information, revealing Ω_{it} to the shoppers.

Because shoppers do not immediately observe Ω_{it}, they make choices under uncertainty. However, they do observe the local price of housing in their island, P_{it}, which they use to make inference; shoppers' information set is therefore $\{P_{it}, \Omega_{it-1}\}$.[8] We derive the information about current conditions contained in P_{it} shortly.

The family metaphor is convenient but not essential. What is essential is that some agents have access to information about realized shocks: Prices cannot reveal information unless that information is already available, perhaps noisily, to some agents in the economy (Hellwig, 1980). We could have achieved the same effect by assuming that only a fraction of households on each island are informed in the spirit of Grossman and Stiglitz (1980). Nothing crucial about our results would change if did this, though the algebra is more cumbersome.[9]

The formal definition of equilibrium is the following.

Definition 1 (Equilibrium). *Given initial conditions* $\left\{\{B_{i-1}, \mathcal{H}_{i-1}, \tilde{\mu}_{i0}\}_{i\in(0,1)}, \tilde{\zeta}_{i-1}\right\}$, *a rational expectations equilibrium is a set of prices,* $\{\{P_{it}, V_{it}, W_{it}\}_{i\in(0,1)}, Q_t, R_t\}_{t=0}^{\infty}$, *and quantities,* $\{\{B_{it}, N_{it}^c, N_{it}^h, N_{it}, C_{it}, H_{it}, \Delta_{it}, X_{it}, L_{it}, Z_{it}\}_{i\in(0,1)}, Y_t\}_{t=0}^{\infty}$, *which are contingent on the realization of the stochastic processes* $\{\{\tilde{\mu}_{it}\}_{i\in(0,1)}\}_{t=0}^{\infty}$ *and* $\{\tilde{\zeta}_t\}_{t=0}^{\infty}$, *such that for each* $t \geq 0$ *and* $i \in (0,1)$*:*

(a) Shoppers and worker-savers optimize, i.e. $\{C_{it}, \Delta_{it}, N_{it}, B_{it}\}$ *are solutions to* $\max_{\{C_{it},\Delta_{it},N_{it},B_{it}\}} E[U_{it}]$ *subject to*

(i) $\mathcal{B}_{it} \leq 0$

[8]Shoppers also observe the price of old vintages for which trade does not occur in equilibrium. Nevertheless, these prices convey no new information to shoppers as these prices are a function of shoppers' local demand. We show this formally in our discussion before Proposition 3.

[9]We took this approach in our working paper, Chahrour and Gaballo (2017). Earlier drafts also showed that our mechanism could arise on the supply side of the economy, more like Lucas (1972).

(ii) C_{it}, Δ_{it} *are measurable with respect to* $\{P_{it}, \Omega_{it-1}\}$

(iii) N_{it}, B_{it} *are measurable with respect to* $\{\Omega_{it}\}$*;*

(b) Housing producers optimize, i.e. $\{N^h_{it}, Z_{it}, L_{it}, \Delta_{it}\}$ *are solutions to* $\max_{\{N^h_{it}, Z_{it}, L_{it}, \Delta_{it}\}} \Pi^h_{it}$ *subject to (4) and (5);*

(c) Consumption producers optimize, i.e. $\{N^c_{it}\}_{i\in(0,1)}$ *are solutions to* $\max_{\{N^c_{it}\}_{i\in(0,1)}} \Pi^c_t$ *subject to (6);*

(d) Markets clear, i.e. equations (7) - (8) hold.

The measurability constraints above imply that the consequences of a particular choice must be evaluated by averaging across states which remain uncertain under the relevant measure. For example, let Λ_{it} be the Lagrange multiplier associated with constraint (i), which the worker-saver correctly evaluates based on Ω_{it}. Shopper optimality requires equating the marginal utility of consumption with the expectation of this multiplier conditional to the shopper's information set, i.e. $\phi C_{it}^{-1} = E[\Lambda_{it}|P_{it}, \Omega_{it-1}]$. This condition equates the average benefits and costs of a marginal change in C_{it}, weighted by the probability of states indistinguishable to the shopper.

2.3 Linearized model

We now derive conditions describing an approximation to equilibrium in the economy, in which we assume that deviations from the deterministic steady state of the economy are sufficiently small. Going forward, lower-case variables refer to log-deviations from this steady-state and we refer to the shoppers' information set as p_{it}.

Shoppers demand consumption and housing goods according to the following:

$$c_{it} = -E[\lambda_{it}|p_{it}] \tag{9}$$

$$\delta_{it} = -E[\lambda_{it}|p_{it}] - p_{it}, \tag{10}$$

where λ_{it} is the marginal value of household i's resources — known by the worker but not the shopper — and $E[\cdot|p_{it}]$ denotes the shopper's expectation conditional on the local house price p_{it} and observations of the past. Equations (9) and (10) show that, ceteris paribus, a higher perceived value of resources lowers shoppers' demand for both consumption and housing.

Optimality of worker-shopper choices requires:

$$w_{it} = -\lambda_{it} \tag{11}$$

$$\lambda_{it} = E[\lambda_{it+1}|\Omega_{it}] + r_t. \tag{12}$$

The worker provides any quantity of labor demanded, so long as the offered wage equals the household Lagrangian, and purchases bonds until the interest rate reflects the difference between the current and the expected future marginal value of budget resources, which the worker-saver forecasts based on Ω_{it}, the full current information set.

We pause here to observe that the condition for *intratemporal* optimality (marginal rate of substitution equals marginal product of labor) will hold only on average in our economy because consumption and labor choices are conditioned on different information (see Remark 1 in Appendix A.2). This means that our model can generate a time-varying labor wedge, an observation we explore in Section 4.3.

Housing producer optimality conditions are standard:

$$z_{it} + q_t = p_{it} + \delta_{it}, \tag{13}$$

$$n^h_{it} + w_{it} = p_{it} + \delta_{it} \tag{14}$$

$$v_{it} = p_{it} + \delta_{it} \tag{15}$$

with production technology given by

$$\delta_{it} = \alpha\gamma n^h_{it} + \alpha\left(1-\gamma\right)\left(z_{it} - \tilde{\zeta}_t\right), \tag{16}$$

after imposing the fact that $l_{it} = 0$.

Consumption producer optimality requires:

$$n^c_{it} = \tilde{\mu}_{it} - \eta\left(w_{it} - w_t\right) + n^c_t \tag{17}$$

$$y_t = n^c_t \tag{18}$$

$$w_t = 0, \tag{19}$$

where w_t denotes the average log-wage in the economy. Condition (17) captures firms' demand for island-specific labor. Firms demand more of a type of labor whenever its productivity is high or its wage is low compared to the average, or if they demand more labor overall. Notice that the wage for the aggregate labor bundle is constant, since there are no shocks to aggregate consumption productivity; we relax this assumption in the Appendix.

All relations above obtain as exact log transformations. Only the island resource constraint needs to be log-linearized as follows:[10]

$$\beta b_{it} + C(c_{it} - c_t) = C(w_{it} - w_t) + C(n^c_{it} - n^c_t) - Qz_{it} + b_{it-1}. \tag{20}$$

In equation (20), C and Q represent the deterministic steady state values used in the linearization. We reiterate here that neither the local housing stock (h_{it}) nor new house production (δ_{it} and n^h_{it}) appear in (20): since housing is non-tradable, housing adjustments can never be used to raise island-level consumption. Market clearing conditions $0 = \int z_{it}di$, $0 = \int b_{it}di$, and $n = \int n^c_{it}di + \int n^h_{it}di$ complete the description of equilibrium in the linearized economy.

3 Learning from prices

This section presents the main theoretical results regarding the inference problem of shoppers. We derive the value of household resources as a function of exogenous shocks, characterize the shoppers' price signal, and then show the implications for inference.

3.1 Marginal value of budget resources

The only friction in the economy is shoppers' uncertainty regarding the marginal value of household budget resources. Without this friction, the model is a standard real business cycle economy. Lemma 1 expresses the value of resources, λ_{it}, as determined by the choices of worker-savers. It depends on the income prospects of the household and end-of-period wealth.

Lemma 1. *In equilibrium,*

$$\lambda_{it} = E[\lambda_{it+\tau}|\Omega_{it}] = -\omega_\mu \tilde{\mu}_{it+1} - \omega_b b_{it} \qquad \textit{and} \qquad r_t = 0 \tag{21}$$

for any $\tau \geq 0$ and any $i \in (0, 1)$. In addition, $\omega_\mu > 0$ and $\omega_b > 0$, with $\lim_{\beta\to 1} \omega_b = 0$.

Proof. Proved in Appendix B □

[10] We linearize bond holdings in levels because B_{it} can take negative values.

Intuitively, the intertemporal arbitrage carried out by worker-savers allows them to equalize the marginal value of budget resources across time. One important implication of Lemma 1 is that the real interest rate does not react to housing productivity shocks. This is again a consequence of the fact that housing wealth cannot be sold across islands.

By contrast, local labor and bonds can be traded across islands in exchange for consumption. Therefore, islands with more productive labor or higher savings have better consumption prospects and a lower marginal value of resources. Thus, Lagrangian multipliers depends on future labor productivity, $\tilde{\mu}_{it+1}$, and on bond holdings at the end of the period, b_{it}.

As β approaches one, λ_{it} becomes independent of bond holdings. This happens because, as β tends towards unity, bond wealth generates no interest earnings and is rolled over indefinitely. To simplify exposition, we present derivations in the case of $\beta \to 1$ from below so that λ_{it} is approximately exogenous. However, all the results in our propositions are stated for all $\beta \in (0, 1)$.

We conclude this section with a remark on the distinction between local and aggregate productivity in the consumption sector. Our model resembles a standard real business cycle model, in that an *aggregate* shock to future productivity in the consumption sector would drive the future value of resources and the real interest rate in opposite directions, leaving λ_{it} and current consumption unchanged. This is why papers looking for business cycle effects of productivity news require either real adjustment frictions (e.g. Jaimovich and Rebelo, 2009) or nominal frictions along with suboptimal monetary policy (e.g. Lorenzoni, 2009). In our environment, however, *local* news has an effect on λ_{it}. The information friction we describe below transforms the effects of local news into fluctuations in aggregate demand.

3.2 Local housing price

We now derive the signal that shoppers use to make their inferences about $\hat{\mu}_{it+1}$. To economize notation, we solve for equilibrium assuming that at time t, $\tilde{\mu}_{it} = \tilde{\zeta}_{t-1} = 0$, so that $\tilde{\mu}_{it+1} = \hat{\mu}_{it+1}$ and $\tilde{\zeta}_t = \zeta_t$. Since past shocks are common knowledge, nothing in the description of equilibrium changes when we relax this.

Rearranging first order conditions from the housing sector, we recover the standard Cobb-

Douglass result that the price is a linear combination of input costs weighted by their elasticity:

$$p_{it} \;=\; (1-\alpha)v_{it} + \alpha\gamma w_{it} + \alpha\left(1-\gamma\right)\left(\zeta_t + q_t\right). \tag{22}$$

We wish to rewrite (22) in terms of the exogenous variables and expectations thereof. We substitute (21) into the local wage in (11) and, recalling that $\beta \to 1$ implies $\omega_b = 0$, conclude

$$w_{it} = \omega_\mu \hat{\mu}_{it+1} \equiv \mu_i. \tag{23}$$

Equation (23) says that workers who expect higher future local productivity demand higher wages today, while (22) shows that higher wages increase house prices. Going forward, we use the definition of $\mu_i \sim N(0, \sigma_\mu^2)$ above and drop time subscripts for contemporaneous relations.

Importantly, the price of local land only reflects shoppers' local housing demand, since equations (10), (15) and (21) can be combined to get $v_i = E[\mu_i|p_i]$. Hence, although shoppers' do not observe v_i, they can predict it exactly. By contrast, the price of the traded commodity good varies with the aggregate appetite for housing across islands, since market clearing for the commodity good and (13) together imply

$$q = \int v_i di = \int E[\mu_i|p_i]di. \tag{24}$$

Using (23) and (24), shoppers' observation of the house price p_{it} is informationally equivalent to observing the signal:

$$s_i \;=\; \gamma\mu_i + (1-\gamma)\left(\zeta + \int E[\mu_i|p_i]di\right). \tag{25}$$

The crucial feature of the signal in (25) is that it conflates house prices changes caused by local conditions with those caused by aggregate shocks. Moreover, since the correlated portion of the price signal contains an endogenous component, a common change in expectations feeds back into local prices, thereby further shifting the inference of all consumers.

3.3 Equilibrium

We now solve the shopper's inference problem. The main challenge is the self-referential nature of the signal, as its precision depends on the *equilibrium* volatility of the commodity price.

Following the related literature, we focus on linear equilibria. We therefore conjecture that the optimal individual expectation is linear in s_i and takes the form

$$E[\mu_i|p_i] = as_i = a\left(\gamma\mu_i + (1-\gamma)\left(\int E[\mu_i|p_i]di + \zeta\right)\right). \tag{26}$$

In (26), a measures the weight the shopper places on the price signal in forming his forecast. Since the signal is *ex ante* identical for all shoppers, each uses a similar strategy. Integrating across the population yields

$$\int E[\mu_i|p_i]di = a\,(1-\gamma)\left(\int E[\mu_i|p_i]di + \zeta\right). \tag{27}$$

Equation (27) is useful for summarizing how changes in aggregate expectations are amplified by the endogenous signal structure: as the weight a grows from to zero towards $(1-\gamma)^{-1}$, initial changes in expectations experience increasingly strong amplification. The case where $a = (1-\gamma)^{-1}$ is particularly extreme, as any initial perturbation (i.e. by a non-zero productivity shock ζ) must lead to infinitely large fluctuations in $\int E[\mu_i|p_i]di$.

When a does not equal $(1-\gamma)^{-1}$, equation (27) can be solved for the average expectation,

$$\int E[\mu_i|p_i]di = \frac{a\,(1-\gamma)}{1-a\,(1-\gamma)}\zeta, \tag{28}$$

which is a nonlinear function of the weight a. The fact that the average expectation is normally distributed confirms the conjectured form of the optimal individual forecast.

Integrating consumption demand in (9) shows that aggregate consumption equals the average forecast, i.e $c = \int E[\mu_i|p_i]di$. Hence, as long as households put nonzero weight on their signal s_i, aggregate consumption moves with housing productivity, and its variance is

$$\sigma_c^2(a) = \left(\frac{a\,(1-\gamma)}{1-a\,(1-\gamma)}\right)^2 \sigma^2, \tag{29}$$

where $\sigma_c^2 \equiv var(\int E[\mu_i|p_i]di)/\sigma_\mu^2$ and $\sigma^2 \equiv \sigma_\zeta^2/\sigma_\mu^2$ are the variances of the average expectation and the aggregate shock after each is normalized by the variance of the idiosyncratic fundamental. Substituting (28) into the price signal described in equation (25), we get an expression for the local signal exclusively in terms of exogenous shocks:

$$s_i(a) = \gamma\mu_i + \frac{1-\gamma}{1-a\,(1-\gamma)}\zeta, \tag{30}$$

whose precision with regard to μ_i is given by

$$\tau(a) \equiv \left(\frac{\gamma\,(1-a\,(1-\gamma))}{(1-\gamma)\sigma}\right)^2.$$

We next compute the shopper's optimal inference, taking the average weight of other households as given. We seek an a^* such that the covariance between the signal and forecast error is zero, i.e. $\mathrm{E}[s_i(a)(\mu_i - a^* s_i(a))] = 0$, which implies that information is used optimally. The individual best-response weight is thus given by

$$a^*(a) = \frac{1}{\gamma}\left(\frac{\tau(a)}{1+\tau(a)}\right). \tag{31}$$

The function $a^*(a)$ captures the individual's best reply to the profile of others' actions. An equilibrium of the model is characterized by a fixed point, $\hat{a} = a^*(\hat{a})$, and there are as many equilibria as intersections between $a^*(a)$ and the 45° line. In the two top panels of Figure 2 we plot the best-response weight $a^*(a)$ for two different values of σ. The case $\gamma > 1/2$ appears in panel (a) and the case $\gamma < 1/2$ in panel (b). We now provide existence conditions for these equilibria and provide intuition for the different cases.

Unique equilibrium

Our first proposition concerns the case in which local house prices respond relatively strongly to local conditions, i.e. the labor share in construction is greater than one half. In this case, the model always has a unique equilibrium.

Proposition 1. *For $\gamma \geq 1/2$ and any $\beta \in (0,1)$, there exists a unique REE equilibrium, which is characterized by $a_u \in (0, \gamma^{-1})$. Moreover, $\lim_{\sigma\to\infty} a_u = 0$ and $\lim_{\sigma\to 0} a_u = \gamma^{-1}$ with $\partial a_u/\partial\sigma < 0$.*

Proof. Given in Appendix C. □

The negative slope of the best response in the range $a \in [0, (1-\gamma)^{-1}]$ is crucial for understanding the forces behind the equilibrium. Negative slope entails *substitutability* in people's use of information: a higher average response to the signal lowers the individual's optimal weight. This happens because a higher a amplifies the effect of aggregate noise, making s_i less informative about private conditions. This result contrasts with the complementarity featured by other models, like Amador and Weill (2010), and explains why our model can deliver a unique equilibrium for any variance of the the aggregate shock.

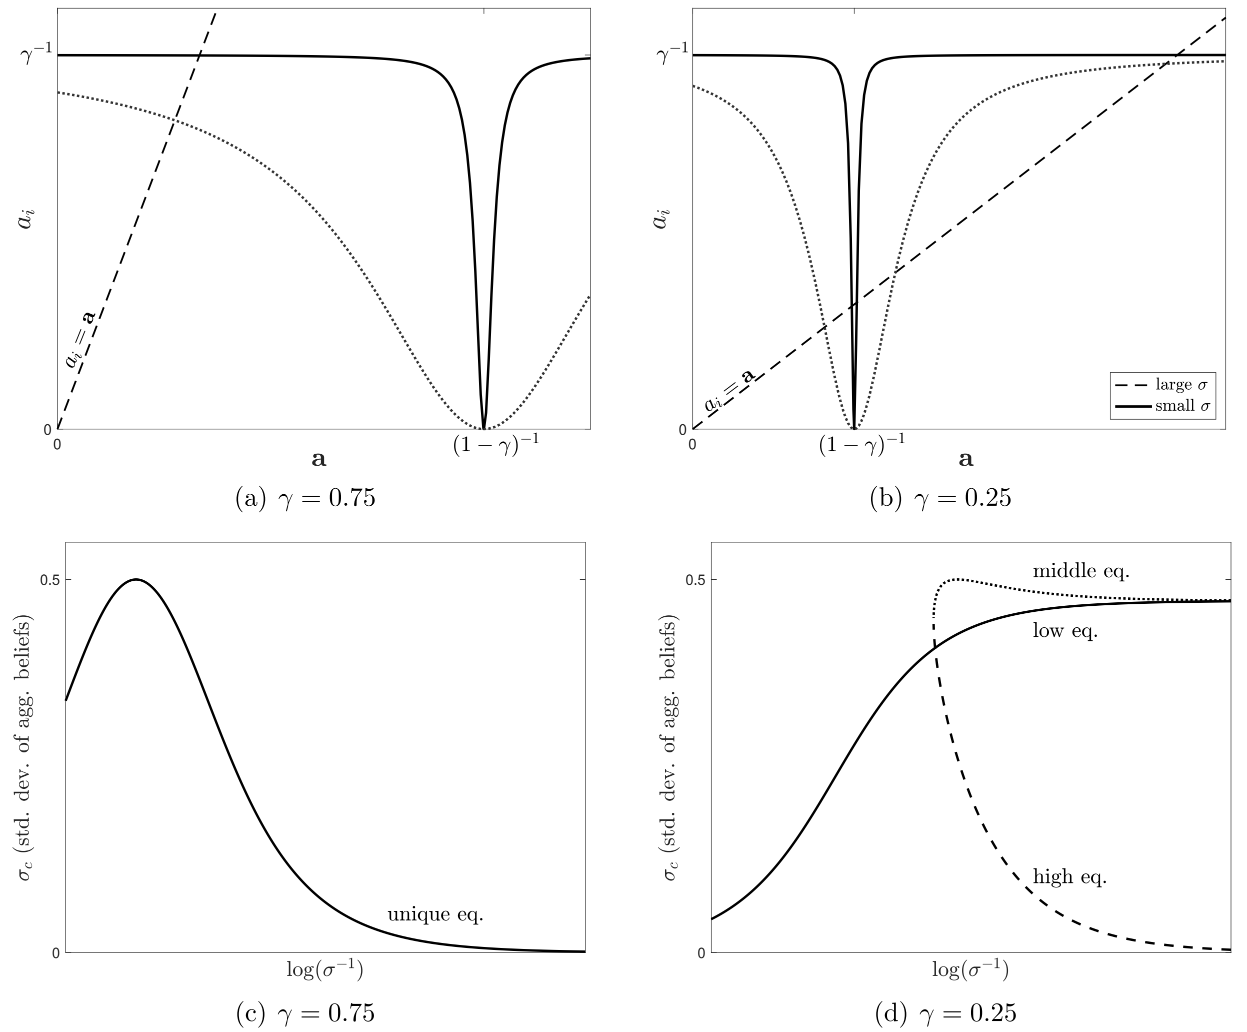

(a) $\gamma = 0.75$ (b) $\gamma = 0.25$

(c) $\gamma = 0.75$ (d) $\gamma = 0.25$

Figure 2: Top panels illustrate the best weight function $a^*(a)$ in a case with unique equilibrium (a) and with multiplicity (b) for two different values of σ. Bottom panels show the evolution of aggregate consumption volatility in a case with a unique equilibrium (c) and with multiplicity (d) as the relative standard deviation of the aggregate shock, σ, ranges from roughly ten (left extreme) to approximately zero (right extreme).

Panel (c) of Figure 2 plots the variance of aggregate beliefs as a function of σ^{-1}. The figure shows that the relationship is non-monotonic, as the equilibrium weight on ζ grows as σ shrinks. Nevertheless, the latter effect eventually dominates so that, in the limit $\sigma \to 0$, average beliefs exhibit no fluctuations. In this limit $a = \gamma^{-1}$ and the local price signal is perfectly informative about μ_i.

Multiple equilibria

When local house prices respond strongly to aggregate conditions, i.e. the local labor share in construction is less than one half, the feedback loop between demand and the commodity price can be so strong that multiple equilibria exist. Proposition 2 summarizes this result.

Proposition 2. *For $\gamma < 1/2$ there always exists a "low" REE equilibrium characterized by $a_- \in (0, (1-\gamma)^{-1})$; in addition, there exists a threshold $\bar{\sigma}^2(\beta)$ with $\partial\bar{\sigma}^2(\beta)/\partial\beta \geq 0$ such that, for any $\sigma^2 \in (0, \bar{\sigma}^2(\beta))$, a "middle" and a "high" REE equilibrium also exist characterized by $a_\circ$ and $= a_+$, respectively, both lying in the range $((1-\gamma)^{-1}, \gamma^{-1})$. In the limit $\sigma^2 \to 0$:*

i. the "high" equilibrium converges to a point with no aggregate volatility:

$$\lim_{\sigma^2 \to 0} a_+ = \min\left(\frac{1}{\gamma}, \frac{1}{1-\gamma}\right) \qquad \lim_{\sigma^2 \to 0} \sigma_c^2(a_+) = 0.$$

ii. the "low" and "middle" equilibria converge to the same point and exhibit non-trivial aggregate volatility:

$$\lim_{\sigma^2 \to 0} a_{\circ,-} = \frac{1}{1-\gamma} \qquad \lim_{\sigma^2 \to 0} \sigma_c^2(a_{\circ,-}) = \frac{\gamma(1-2\gamma)}{(1-\gamma)^2}. \tag{32}$$

Proof. Given in Appendix C. □

The best weight function in this case is plotted in panel (b) of Figure 2. It shows that the function yields three intersections with the 45° line provided the variance of productivity shocks σ is sufficiently low. We demonstrate in the proof that a lower β is isomorphic to considering a larger σ at any a, so $\beta \to 1$ turns out to be the case most favorable to multiplicity.

Importantly, the qualitative features of the unique equilibrium also hold for the "low" equilibrium described above: (i) there is substitutability between individual and average weights; and (ii) amplification increases as the variance of aggregate shocks falls. However, since now $\gamma^{-1} > (1-\gamma)^{-1}$, the "low" equilibrium must be distinct from the full-information equilibrium, implying the model has multiple equilibria in the limit.

Substitutability in information use is key in generating incomplete information in the limit of zero noise. In models like Amador and Weill (2010), which feature complementarity in the use of public information, decreasing exogenous aggregate noise always improves the precision of the signal so that any limit equilibrium is fully revealing. By contrast, here and in Gaballo (2018), strengthening feedbacks offset the direct effect of reducing exogenous noise, generating substitutability in information use and potentially leading to noisy equilibria even in the limit.

Panel (d) of Figure 2 illustrates the volatility of these equilibria. Consumption volatility in the "high" equilibrium case converges to zero as σ^{-1} goes to infinity. By contrast, consumption volatility in the "middle" and "low" equilibria converges to a positive, finite number. Surprisingly, the low and middle limit equilibria have the same stochastic properties as the extrinsic sentiment equilibrium described by Benhabib et al. (2015). In our case, however, fluctuations are driven by infinitesimally-small fundamental shocks, whose realizations coordinate sizable fluctuations in agents' expectations. We elaborate on this connection in Section 5.3.

3.4 The economic forces behind equilibrium

We now summarize the economic forces behind the equilibria described in Propositions 1 and 2. The goal is to clarify the propagation of belief fluctuations to real variables, and illuminate the role of the traded good in generating amplification and multiple equilibria.

To start, consider an island experiencing a positive idiosyncratic shock to tomorrow's productivity ($\mu_{t+1} > 0$). The current worker-saver observes this, and demands a higher wage to supply her labor: she is more productive tomorrow and, since the shopper is already out shopping, today's consumption is not directly affected by her labor supply. High local wages feed into high local house prices, however, which *is* observed by the shopper. So from the perspective of the shopper, a high local house price could signal high future local productivity, and thereby encourage higher consumption expenditures today.

Now, suppose the economy has experienced an aggregate negative shock to construction productivity. That also leads to a higher house price, now on all islands. Since housing is separable and non-tradable, the shopper does not wish to adjust their consumption in response to this shock. But, since the shopper observing higher house prices cannot be sure of the source, he attributes at least some of the change to improved local productivity. Because of this mistake, the shopper increases his demand for both housing and consumption.

In order for quantities to rise, the shopper's optimism must drive an actual increase in labor supply, rather than a change in relative prices. Barro and King (1984) show such real effects cannot happen if all workers and shoppers have the same information. In our model, however, the worker-saver understands that the shopper is making a mistake by spending too much and that, other things equal, savings will fall. This potential fall in wealth increases the marginal utility of budget resources and induces the worker to supply more labor. The worker's choices thereby support the increase in consumption demand with higher actual output. Thus, even though household savings and labor choices are taken under full information, information heterogeneity within the household drives a wedge in the standard intratemporal optimality condition between consumption and leisure. We expand on the model's implications for this labor wedge in Section 4.3.

The mechanism described thus far does not require particular assumptions about housing production, but the addition of the tradable input provides for strong amplification. For, whenever optimism drives up demand for housing across islands, demand for the traded input also rises along with the input price, q. The higher input price pushes up all housing prices, thereby driving an even larger change in expectations, and so on. The strength of this aggregate feedback depends on the elasticity of housing production to inputs, as captured by γ, and the reaction of shoppers to the house price, the weight a. In particular, a unit increase in expectations triggers an increase in the housing price signal of $1-\gamma$, and hence an increase in average expectations of $a(1-\gamma)$.

Our propositions show that, when production relies more on traded than on local factors, amplification be can so strong that multiple equilibria arise. To see the intuition for this multiplicity, it is helpful to focus on the situation in which the fluctuations of housing productivity are vanishingly small. Consider first the case in which we conjecture that q does not move. Then the price signal is perfectly informative about local conditions and and no correlated fluctuations in house or commodity pries may emerge. Hence, an equilibrium in which all shoppers are perfectly informed — and q never moves — must exist.

Consider now the case in which we conjecture that q fluctuates. In this case, shoppers' price signals will be "polluted" by changes in q that act as aggregate noise in inference. This noise correlates demand across islands, drivings further fluctuations in q. When individual expectations react exactly one-to-one to changes to in q, self-fulfilling fluctuations become possible and an equilibrium with imperfect information — and volatile q — also exists.

We have demonstrated that a one-to-one reaction can only occur with a large enough share for the traded input, i.e. $\gamma < 1/2$. To see why this cutoff is important, observe first that the largest weight rational shoppers could ever place on their price signal is $1/\gamma$. This implies that the feedback from average expectations back into average expectations is bounded above by $(1-\gamma)/\gamma$. When $\gamma > 1/2$, this bound implies that aggregate feedback is always strictly less than one and amplification can never translate arbitrarily small productivity shocks into non-trivial fluctuations.

By contrast, when the commodity share is large ($\gamma < 1/2$) even vanishing productivity shocks engender substantial fluctuations. To see why, note that with small fluctuations, agents would place a weight on their signal close to the upper bound of $1/\gamma$. However, this makes the aggregate feedback larger than one. Without a countervailing force, such a situation would imply that any correlated fluctuation in expectations would be indefinitely magnified via the q feedback. Yet, optimal signal extraction provides such a countervailing force: as the effects of aggregate disturbances transmitted by q grow larger, shoppers' optimal reaction to the price signal must shrink. In this way, signal extraction picks the unique variance for aggregate outcomes that is consistent with one-for-one feedback, i.e. $a(1-\gamma) = 1$, and belief fluctuations are sustained.

4 Business cycle fluctuations

In this section, we show that many features of the business cycle can be explained by our model. Our analysis also demonstrates that the learning-from-prices mechanism can qualitatively change the comovement properties of fundamental shocks, implying that many strategies for disentangling shocks may give misleading results if learning from prices is important.

4.1 Public news

Before proceeding to our analysis, we introduce an anticipated (common-knowledge) component of aggregate housing productivity. The decomposition of productivity into a common-knowledge and surprise component serves two purposes. First, it allows us to isolate the effects of the learning channel in our model, as the common-knowledge component of productivity transmits as a standard supply-side shock. Second, by combining the responses of the economy to forecasted and surprise productivity shocks, the model can generate a rich and realistic correlation structure among business cycle variables.

We assume that housing productivity is composed of two independent components

$$\zeta = \zeta^n + \zeta^s;$$

with $\zeta^n \sim (N, \sigma^2_{\zeta^n})$, $\zeta^s \sim (N, \sigma^2_{\zeta^s})$ and $\sigma^2_{\zeta^n} + \sigma^2_{\zeta^s} = \sigma^2_{\zeta}$. The first term ($\zeta^n$) is "news"; it corresponds to the common-knowledge component of productivity, and is known to all agents before they make consumption choices. The second term (ζ^s) is the "surprise"; it is unknown to shoppers and they seek to forecast it using their observation of prices.[11] For future reference, let $\sigma^2_n \equiv \sigma^2_{\zeta^n}/\sigma^2_{\mu}$, and $\sigma^2_s \equiv \sigma^2_{\zeta^s}/\sigma^2_{\mu}$ be the normalized variances of the news and surprise components of productivity respectively.

Only modest modifications are necessary to characterize equilibrium in this case. Shoppers refine the information contained in the price signal by "partialing-out" the known portion of productivity. We can thus rewrite households' expectations as

$$E[\mu_i|p_i] = a(s_i - (1-\gamma)\zeta^n), \tag{33}$$

where $s_i - (1-\gamma)\zeta^n$ captures the information available to the shopper, after he has controlled for the effect of ζ^n. The equilibrium values $\hat{a}$ and the conditions for their existence are the same as in the baseline economy once σ^2_s takes the place of σ^2.

For a given total variance of productivity, $\sigma^2 = \sigma^2_s + \sigma^2_n$, we can now span the space between two polar cases, from the case in which productivity occurs as a pure "surprise" to the case in which the productivity shock is common-knowledge "news". Thus, overall comovements in the economy will represent a mix of demand and supply shocks. We note here that, because the

[11] Chahrour and Jurado (2018) show that this information structure is equivalent to assuming that agents observe a noisy aggregate signal, $s = \zeta + \vartheta$.

Table 1: Business Cycle Comovements

	GDP	**Cons**	**Hours**	**ResInv**	**House Pr**	**ResInv Pr**	**Cons TFP**
GDP	1.00	0.93	0.88	0.64	0.51	0.53	-0.17
Cons		1.00	0.80	0.65	0.47	0.47	-0.02
Hours			1.00	0.50	0.54	0.66	-0.35
ResInv				1.00	0.62	0.37	-0.11
House Pr					1.00	0.81	-0.37
ResInv Pr						1.00	-0.43
Cons TFP							1.00

Note: Data are real per-capita output, real per-capita consumption, per-capita hours in the non-farm business sector, real per-capita residential investment, Case-Schiller real house price index, real price of residential investment, and relative TFP in the construction sector from the World KLEMS databse (http://www.worldklems.net/data.htm). All data are annual log-levels, HP-detrended using smoothing parameter $\lambda = 10$. Date range: 1960 to 2018, except for construction TFP which ends in 2010. Details on data construction can be found in Appendix D.

two components of productivity transmit very differently in the economy, moments generated by projecting variables onto total productivity ζ could give very misleading inference on productivity's effects. Econometric identification of the distinct components of productivity represents a substantial empirical challenge, for which Chahrour and Jurado (2019) provide some guidance in related contexts.

4.2 Demand-driven fluctuations

Table 1 summarizes unconditional correlations between business cycle variables in US data. Although these are simple raw statistics, the table summarizes several facts that have been documented by more sophisticated empirical analysis. In particular, the table demonstrates that business cycles are dominated by demand-like fluctuations with real quantities, house prices, and residential investment all substantially comoving. Meanwhile, construction productivity is at most weakly negatively related to any of these variables.

In the model, the emergence of demand fluctuations can be seen intuitively by analyzing the aggregate demand and aggregate supply schedules. Using equations (9), (10), (14) and (16), we can express aggregate demand and supply in the housing market as

$$\delta = c - p, \tag{34}$$

$$\delta = \frac{\alpha\gamma}{1-\alpha\gamma}p - \frac{\alpha(1-\gamma)}{1-\alpha\gamma}\zeta. \tag{35}$$

Moreover, because of the learning channel, we know that aggregate consumption shifts upwards in response to a correlated increase in price signals across island,

$$c = \int E[\mu_i|p_i]di = a(s - (1-\gamma)\zeta^n).$$

Note this expression implies c *does not* move with the news component of housing productivity, as ζ^n is being removed from the price signal.

To derive the implications of shopper inference for housing demand, use $p = (1-\alpha)v + \alpha s$ and $v = c$ to express $s = (p + (1-\alpha)a(1-\gamma)\zeta^n)/((1-\alpha)a + \alpha)$. Substituting the expression for c into (34) we get

$$\delta = \frac{\alpha(a-1)}{(1-\alpha)a+\alpha}p + \frac{\alpha a(1-\gamma)}{(1-\alpha)a+\alpha}\zeta^n. \tag{36}$$

When aggregate conditions do not feed into shoppers' beliefs ($a = 0$), equation (36) entails a standard downward-sloping aggregate demand relation in the housing market, while consumption and working hours are invariant to housing sector productivity. By contrast, when learning from prices is sufficiently important—i.e. whenever a is larger than one—equation (36) shows that δ and p must comove in response to surprise shocks.

We can now solve for equilibrium consumption, residential investment, and the price of new housing as functions of shocks and the equilibrium inference coefficient:

$$c = \frac{a(1-\gamma)}{1-a(1-\gamma)}\zeta^s \tag{37}$$

$$p = \alpha(1-\gamma)\zeta + (1-\alpha\gamma)c \tag{38}$$

$$\delta = -\alpha(1-\gamma)\zeta + \alpha\gamma c. \tag{39}$$

Expressions (37) - (39) above are useful for disentangling the direct effects of productivity from the learning channel. Equation (37) shows that a correlated mistake due to a surprise in aggregate productivity moves consumption. Equations (38) and (39) show how this change in beliefs transmits into the housing market, moving prices for new housing and residential investment in the same direction. Under full information ($a = 0$) these spillovers across markets would disappear. Meanwhile, the appearance of housing productivity ζ in (38) and (39) is independent of a, and captures the standard neoclassical channel through which productivity changes drive prices and quantities in opposite directions.

With a few more lines of algebra, we have that

$$c = \int \big(\lambda_i - E[\lambda_i|p_i]\big)di = \int n_i^h di = \int n_i^c di. \tag{40}$$

Equation (40) implies that an increase in consumption corresponds to an increase in working hours in both sectors. In times of optimism, shoppers' spending increases but wages do not, so production increases.

Since empirical house price measures include both new and existing homes, we also derive the connection between the price of new housing, p, and the price of the total housing stock, p^H. In the Appendix, we show that the price of each vintage moves with shoppers' expected Lagrangian, $p_{i|k} = -E[\lambda_i|p_i]$. This happens because the supply of past vintages is fixed and prices must completely absorb any change in expectations. We therefore find that $p^H = \kappa p + (1-\kappa)E[-\lambda_i|p_i]$ where $\kappa \in (0,1)$ is the steady state fraction of new houses in the total housing stock.

Collecting these results, it is straightforward to demonstrate the following:

Proposition 3. *For σ_s^2 sufficiently small, surprise aggregate productivity shocks drive positive comovement of consumption, employment (in both sectors), residential investment, prices for new and existing housing, commodity prices, and the price of land.*

Proof. Given in appendix C. □

In sum, our model exhibits comovements of aggregate business cycle variables in response to sufficiently small productivity shocks, *in any equilibrium* and for *any configuration of parameters*. To an outside observer, the economy would appear to be buffeted by recurrent shocks to aggregate demand.

Proposition 3 requires aggregate shocks to be "sufficiently small". Intuitively, this is needed because price signals must be informative enough that shoppers put substantial weight on them. Yet, Proposition 1 shows that for $\gamma \geq 1/2$ aggregate fluctuations still disappear in the limit $\sigma \to 0$. Taken together, these results raise the question: can the unique equilibrium model deliver comovement and realistically large business cycle fluctuations at the same time? The answer is yes. As we show in the following section, even if the surprise component accounts for a small fraction of realized productivity, demand driven fluctuations may dominate unconditional comovements.

4.3 Business cycles under unique equilibrium

In this section, we discuss the model's business cycle properties when it has a unique equilibrium. We organize the discussion around three pictures illustrating its implications for business cycle comovements, amplification, and correlations with productivity. Our goal is to show that our model can qualitatively account for the empirical patterns reported in Table 1.

While we do not undertake a full quantitative evaluation of the model, we wish to demonstrate the mechanism can be very powerful for reasonable parameterizations. To this end, we calibrate a set of parameters to standard values and/or long run targets in the data. We set the model period to one year. We set $\beta = 0.96$ consistent with an annual real interest rate of roughly 4%. We set $\phi = 0.66$, to be consistent with 2013-2014 CPI relative importance weight placed on shelter. Estimates of η, the elasticity of local labor demand, range in the literature from below one (Lichter et al., 2015) to above twenty (Christiano et al., 2005). We use $\eta = 2$ as a baseline, and note that the aggregate effects of changing η can be offset one-for-one by changing the volatility of local productivity.

For the housing sector, we follow Davis and Heathcote (2005) in fixing $\alpha = 0.89$ to match the evidence that land accounts about 11% of new home prices.[12] We pick the residential investment labor share parameter $\gamma = .526$ by computing the ratio of labor input costs to materials and energy costs in the construction sector, using Bureau of Labor Statistics data from 1997-2014. Finally, we select the volatility of local productivity shocks relative to aggregate shocks $\text{std}(\hat{\mu}_i)/\text{std}(\zeta) = 10$, implying $\sigma = 0.228$.

Comovement in business cycle variables

Figure 3 plots the unconditional correlations and volatilities of several variables in the economy. On the horizontal axis of each panel we vary the ratio between the forecastable and non-forecastable components of productivity, going from pure "surprise" on the left to pure "news" on the right, while holding the total variance σ constant.

Panel (a) of the figure plots the correlation of consumption and house prices with residential investment. Towards the left of the panel, when productivity is mostly unanticipated,

[12] For existing homes, Davis and Heathcote (2007) find that land prices accounts for a larger portion of home prices.

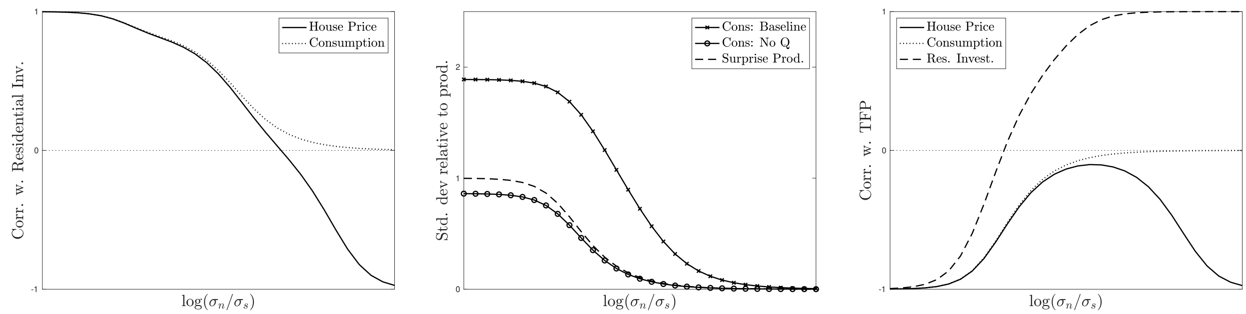

(a) Comovements (b) Amplification (c) Relation with productivity

Figure 3: Panels illustrate unconditional correlation and volatility of business cycle variables as a function of the ratio between volatility of the forecastable and non forecastable component for the baseline case of $\gamma = 0.526$.

our learning channel dominates: residential investment, house prices and consumption all perfectly comove. Given the results derived above, this also implies comovement in hours in both sectors, the average price of land, and the price of commodities.

By contrast, when productivity is largely common knowledge, prices and quantities in the housing market exhibit the negative correlation associated with supply-driven fluctuations, while consumption does not move. Therefore, the more housing productivity is anticipated, the more the economy behaves like a standard real business cycle model. In between these two extremes, the model generates positive but imperfect correlations, consistent with the data reported in Table 1.

Amplification

What is the role of the endogeneity of the signal in generating amplification? Panel (b) of Figure 3 plots the standard deviation of consumption relative to that of aggregate productivity, as a function of the share of productivity that is forecastable. The panel contrasts two cases (i) the baseline model and (ii) the counter-factual case in which the price signal, $\tilde{s}_i = \gamma\mu_i + (1-\gamma)\zeta^s$, excludes its dependence on q. This comparison is useful to evaluate the role of q in amplifying the impact of surprise shocks. To highlight this aspect we also draw the standard deviation of the surprise component of productivity, which by construction falls from one to zero going from left to right.

The comparison is striking. With a completely exogenous price signal, the volatility of consumption, while positive, would be strictly less than the volatility of the surprise component of productivity. This is not the case for our baseline calibration, when the signal is endogenous. The surprise component is amplified substantially, such that consumption remains more volatile than aggregate productivity even when more than 90% of productivity fluctuations are anticipated (near the middle of the horizontal axis)!

The source of amplification can also be seen in our analytical results via equation (37). That equation shows there is a range of parameters where aggregate consumption responds more than one-to-one to productivity shocks.[13] This result depends on the endogenous precision of the signal and, in particular, on having the commodity price q enter in local house prices. One can easily verify that, with a constant q, the reaction of expectations to productivity shocks cannot exceed unity, provided $\gamma > 1/2$.[14]

Relationship with construction TFP

In our model, the noise in people's inference comes from a fundamental shock: housing productivity. One major advantage of our approach to microfounding information is that it provides testable implications about how beliefs fluctuations should relate to measurable economic fundamentals. In this section, we explore this potential by showing that the data are generally consistent with the model's implications for one direct (i.e. model-independent) measure of a fundamental shock: construction TFP. Other shocks may play an important role in the cycle and, as we show in Section 5.1, can induce the same comovements via the learning channel. However, here we emphasize how learning from prices qualitatively changes the transmission of supply shocks and offers one possible interpretation of TFP's contractionary effects.[15]

To this end, the last column of Table 1 reports business cycle correlations with relative productivity in the construction sector — the data analogue to ζ — using the USA KLEMS productivity data of Jorgensen et al. (2012). Overall, the column shows that this measure of housing-sector productivity is negatively, but weakly, correlated with business cycle variables.

[13]This occurs when $\hat{a} \in (1/2(1-\gamma), 1/\gamma)$ with $\gamma \in (1/2, 2/3)$ then $\partial c/\partial \zeta^s > 1$.

[14]To see, suppose that q is fixed, so that the price signal corresponds to $s_i(0)$ in (30) having a precision $\tau(0)$. Then $E[\mu_i|s_i(0)] = \gamma^{-1}\tau(0)(1+\tau(0))^{-1}s_i(0)$, so that $\partial E[\mu_i|s_i(0)]/\partial\zeta = (1-\gamma)\gamma^{-1}\tau(0)(1+\tau(0))^{-1} < 1$.

[15]Galí (1999) and Basu et al. (2006) find that aggregate productivity is contractionary for hours, while Basu et al. (2014) find evidence that investment-specific productivity has contractionary effects across many variables. Angeletos and La'O (2009) propose a different dispersed information mechanism by which employment can fall in response to positive productivity shocks.

Most notably, residential investment is somewhat negatively correlated with this measure of productivity, a result that would be difficult to reproduce in a full information environment.

Panel (c) of Figure 3 illustrates the correlations of residential investment, the price of housing, and consumption as a function of the ratio between the volatilities of the news and surprise components of productivity. These correlations depend on the fraction of anticipated productivity and, as in the data, are generally not perfect. Correlations with total productivity are imperfect because the two components of productivity – surprise and news – are transmitted very differently in the economy. In particular, so long as a sufficient portion of productivity is unanticipated, all of these variables are negatively correlated with productivity. When instead productivity is mostly common knowledge, consumption and hours do not move while residential investment and house prices move in opposite directions.

Implications for the labor wedge

How does our model address the (Barro and King, 1984) challenge and generate realistic business-cycle comovement without relying on contemporaneous changes to productivity? The answer is that the model generates a counter-cyclical distortion of the *intratemporal* margin or "labor wedge," so that both hours and consumption can rise at the same time. We draw out this implication below.

Frictionless real business cycle models usually include, as a condition of intratemporal optimality, that the marginal product of labor should equal the household marginal rate of substitution. The labor wedge measures deviations from this condition:

$$\tau_t \equiv \log\left(\frac{MPN_t}{MRS_t}\right), \tag{41}$$

where MPN_t is the marginal product of labor and MRS_t is the marginal rate of substitution between leisure and consumption. Several authors have argued that empirical analogues to this quantity are counter-cyclical, i.e. that τ_t is high during recessions. Following Karabarbounis (2014) (and ignoring labor taxes) this wedge can be decomposed into two terms,

$$\tau_t^F \equiv \log(MPN_t/W_t) \qquad \text{and} \qquad \tau_t^H \equiv \log(W_t/MRS_t), \tag{42}$$

so that $\tau_t = \tau_t^F + \tau_t^H$.

The first term is the "firm-side" wedge and describes the failure of marginal product to

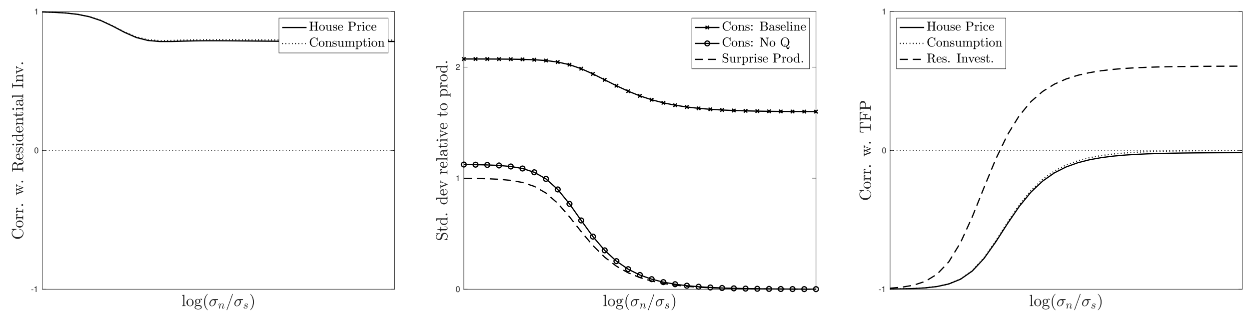

(a) Comovements (b) Amplification (c) Relation with productivity

Figure 4: Panels illustrate correlation and the unconditional volatility of business cycle variables as a function of the ratio between volatility of the forecastable and non forecastable component for the case of $\gamma = 0.45$.

equal the wage. The second is the "household-side" labor wedge, and corresponds with failure of the marginal rate of substitution to equal the wage. Since labor markets in our model are competitive with full information, we immediately know that $\tau_t^F = 0$.

Using our functional forms and information assumptions, we have that

$$\begin{aligned}\tau_t^H &= w_t - c_t \\ &= \int E[\lambda_{it}|p_{it}]di - \lambda_t \\ &= \int E[\mu_i|p_{it}]di.\end{aligned}$$

In words, booms in our economy correspond to moments when people are optimistic about their local conditions — $\int E[\mu_i|p_{it}]di$ is positive —, where c_t grows faster than w_t, and where $\tau_t^H < 0$. This pattern for τ^H is exactly the qualitative pattern described by Karabarbounis (2014).

4.4 Multiple equilibria: supply shocks or animal spirits?

In this section, we explore the properties of one equilibrium when $\gamma < 1/2$ as an illustration of the amplification power of our mechanism. We focus on the "low" equilibrium, characterized by a_- in Proposition 2, since this equilibriums turns out to be learnable in the sense of the adaptive learning literature (see Section 5.4.)

In Figure 4, we present correlations and amplification plots for the case of the "low" equilibrium, changing only $\gamma = 0.45$ with respect to our baseline calibration. Panel (b) shows that, in contrast to our original calibration, consumption remains roughly twice as volatile as realized productivity even as the variance of its surprise component goes to zero. This happens because even infinitesimal surprises drive large fluctuations in beliefs. Note also that the endogeneity of the price signal is crucial to this result: if inference were based on the counter-factual signal $\tilde{s}_i$ that excludes q, the model could deliver large fluctuations in consumption, but these would disappear as σ_s shrinks.

The housing demand and supply relations in (35) and (36) provide an alternative perspective on this powerful amplification. As a approaches $1/(1-\gamma)$, the slope of the curves coincide, implying the two curves overlie one another. In this case, the model exhibits extreme amplification of vanishingly-small shocks, as any point along the coincident upward-sloping curves represents a market clearing allocation and equilibrium volatilities are pinned down by the conditions for optimal inference.

Since belief fluctuations do not disappear with σ_s in this parameterization, the model has very different implications for aggregate comovements. First, panel (a) shows that house prices and consumption remain positively correlated with residential investment even when nearly all of realized productivity is anticipated. Second, as shown in panel (c), when more of productivity is anticipated, the correlation of consumption and house prices with productivity becomes very small. This happens because, though consumption and house prices move substantially with surprise productivity, surprise shocks themselves play a small role in total productivity. Hence, it is with more public information that consumption and house prices appear most disconnected from fundamentals!

In the limit of a small surprise component, house prices and residential investment are moved by infinitesimal productivity surprises. An econometrician looking at the data generated by our model would be unable to measure such small revisions in productivity and would probably conclude that the housing market is moved by animal spirits in the vein of Burnside et al. (2016); Shiller (2007) or sentiments as in Angeletos and La'O (2013) and Benhabib et al. (2015). Our model shows how demand-driven waves can be the result of extreme amplification of small fundamental shocks sustained by the feedback loop of learning from prices. With respect to earlier models of sentiments, the different is sharp: the degree of optimism or pessimism in the economy in our model is fully determined by (potentially

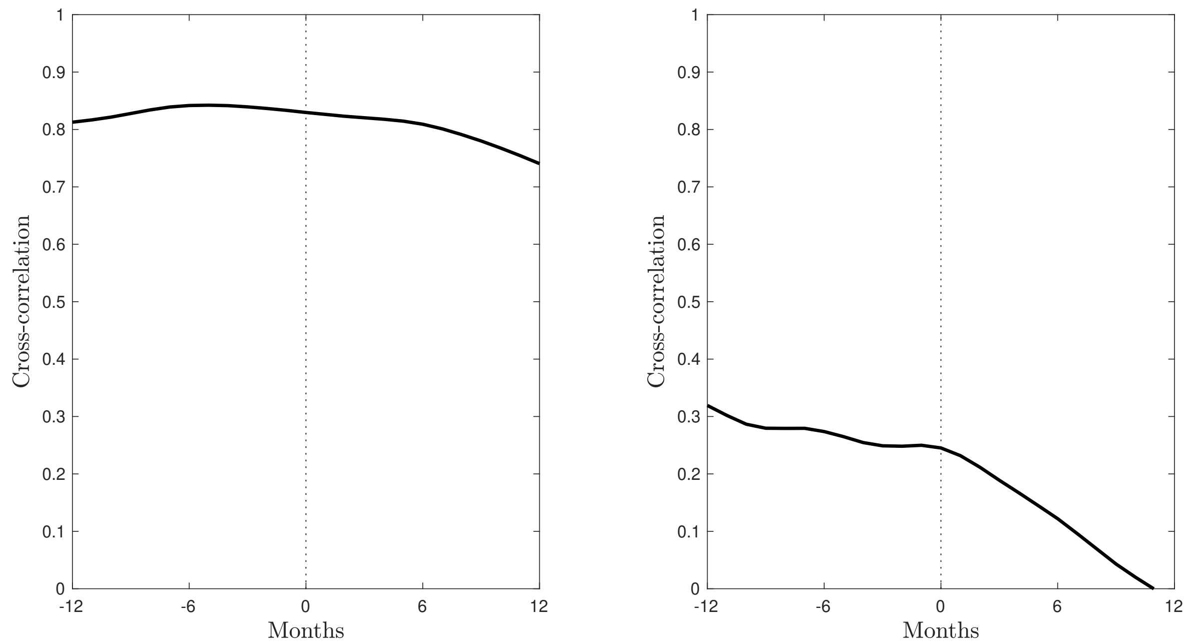

(a) Income expectations at time t vs house price experiences at time $t+h$.

(b) Income expectations at time t vs economic news heard $t+h$.

Figure 5: Auto-correlations of survey measure of own income expectations with respect to own house price experience (panel a) and with respect to news heard about the economy (panel b).

small) fundamental changes rather being totally erratic or “animal”.

4.5 Evidence from survey data

The essential feature of our model is that people’s expectations about their future prospects depend on their own market experiences, particularly housing. We provide here one piece of evidence from survey data that suggests this mechanism may be important in practice.

To this end, we use evidence from the Michigan Survey of Consumer Expectations. Survey participants are asked each month about (i) their perceptions of local house price growth over the last year (ii) whether they have heard good or bad news about overall economic conditions and (iii) what they expect regarding their own real income growth over the coming year. The survey then produces index numbers from the answers to these questions, essentially subtracting those who experienced/heard/expect about negative outcomes from those who have experienced/heard/expect positive ones.

Panel (a) of Figure 5 plots the autocorrelation structure of people’s current expectations about future income, with respect to their recent experiences in the housing market. Negative numbers on the horizontal axis reflect past responses to the housing experience question, while positive numbers reflect future responses. Panel (a) shows that the two series are *extremely* strongly correlated, with past housing experiences leading income expectations by roughly half a year (as measured by the peak correlation.) This result suggests a strong connection between peoples' past experiences in the housing market and their expectations about their own income, exactly as our model predicts.

By contrast, Panel (b) of the figure plots the correlation structure of people's current expectations of their own income with respect to what they report having heard about aggregate economic developments. The correlation in this picture is *much* smaller than in Panel (a), suggesting that what people have heard about the aggregate economy (if they've heard anything) plays a much smaller role in forming people's expectations about their own prospects.

While these results are far from dispositive on the merits of our mechanism, we think they provide some initial evidence that learning from prices is plausible in the context of housing.

5 Extensions

This section presents several extensions that demonstrate the mechanism is robust to various modeling details. In Section 5.1, we explore the impact of contemporaneous and future aggregate shocks to consumption production. In Section 5.2, we allow households to observe additional private information about local conditions and show that our results do not rely on excluding exogenous sources of information. In Section 5.3, we explore whether extrinsic noise may drive fluctuations jointly with aggregate productivity and conclude that this is never the case. Finally, Section 5.4 studies the issue of stability under adaptive learning for the various equilibria of the baseline model.

5.1 Aggregate shocks in consumption production

For this extension, we modify the production function of the consumption sector to allow for aggregate shocks to labor productivity,

$$Y_t = \tilde{\zeta}_t^c \left(\int e^{\tilde{\mu}_{it}/\eta} {N_{it}^c}^{1-\frac{1}{\eta}} di \right)^{\frac{1}{1-\frac{1}{\eta}}} . \tag{43}$$

Consumption productivity is a random walk with i.i.d. disturbance $\zeta_t^c \sim N(0, \sigma_\zeta)$. To simplify our exposition, we focus on time t and assume that workers in island i, but not shoppers, know $\{\zeta_{t+1}^c, \zeta_t^c, \mu_{it+1}\}$ and abstract from the presence of other aggregate shocks. A few lines of algebra shows that

$$\lambda_{it} = -\omega_\mu \mu_{it+1} - \omega_b b_{it} - \zeta_t^c \tag{44}$$

We note immediately that a contemporaneous productivity shock in consumption is equivalent to an increase in consumption spending (measured in consumption units). Given the properties of log utility, an increase in consumption spending induces an increase in housing spending as well. In other words, a productivity shock to consumption production is equivalent to an exogenous demand shock in the housing sector.

Including the future realization of aggregate productivity helps to clarify that the model cannot generate demand shocks in the form of news about aggregate productivity as in Lorenzoni (2009). To see this, notice that

$$r_t = \lambda_t - \lambda_{t+1} = -\zeta_t^c + \zeta_{t+1}^c. \tag{45}$$

Thus, the real interest rate adjusts to equalize the return on savings in the two periods and anticipation of higher productivity in the future has no effect on the *intertemporal* margin, i.e. on consumption choices today (see also Remark 2 in Appendix B.2). This is a feature that our model shares with frictionless real economies, as Angeletos (2018) clarifies. By contrast, news about future productivity creates a demand shock in Lorenzoni (2009) because of the presence of nominal rigidities and monetary policy that is suboptimal. A corollary to this result is that no current variable in the economy, other than the real interest rate, moves with anticipated aggregate consumption shocks, so shoppers will not be able to learn about them in advance.

Contemporaneous consumption productivity shocks, by contrast, decrease the marginal value of households resources, pushing up the real wages demanded by workers. In the Appendix we show that the price signal in this case is:

$$s_{it} = \gamma(\mu_{it+1} + \zeta_t^c) + (1-\gamma)\int E[\mu_{it+1} + \zeta_t^c | s_{it}]di, \tag{46}$$

where again we present the case $\lim_{\beta\to 1}\omega_\mu = 0$ with $\tilde{\mu}_{it+1}$ normalize by ω_μ.

One again, correlated fundamentals generate confusion between the idiosyncratic and common components of the signal. As before, the individual expectation of a household is formed according to the linear rule $E[\mu_{it+1} + \zeta_t^c | s_{it}] = as_i$. Hence, the signal embeds the average expectation, which causes the precision of the signal to depend on the average weight a. Following our earlier analysis, the realization of the price signal can be rewritten as

$$s_i = \gamma\mu_{it+1} + \frac{\gamma}{1-a(1-\gamma)}\zeta_t^c, \tag{47}$$

where a represents the average weight placed on the signal by other shoppers. The average expectation is given by

$$\int E[\mu_{it+1} + \zeta_t^c | s_{it}]di = \frac{\gamma a}{1-a(1-\gamma)}\zeta_t^c, \tag{48}$$

which is slightly different from (29). The shopper's best response function is now given by

$$a^*(a) = \frac{1}{\gamma}\left(\frac{(1-a(1-\gamma))^2 + (1-a(1-\gamma))\,\sigma^2}{(1-a(1-\gamma))^2+\sigma^2}\right). \tag{49}$$

While the best-response function in equation (49) is slightly different than in (31), the characterization of the limit equilibria is identical.

Proposition 4. *In the limit $\sigma_\mu^2 \to 0$, the equilibria of the economy converge to the same points as the baseline economy. For $\gamma > 1/2$: there exists a unique equilibrium $\hat{a}$ such that $\lim_{\sigma^2_{\zeta^c}\to 0} a^\mu = \gamma^{-1}$ with $\lim_{\sigma^2_{\zeta^c}\to 0}\sigma_c^2 = 0$. For $\gamma < 1/2$ instead three equilibria exist such that*

$$\lim_{\sigma^2_{\zeta^c}\to 0}\hat{a} \in \{a_-, a_\circ, a_+\} \qquad \textit{with} \lim_{\sigma^2_{\zeta^c}\to 0}\sigma_c^2(\hat{a}) \in \{\sigma_c^2(a_-), \sigma_c^2(a_\circ), \sigma_c^2(a_+)\}.$$

Proof. Follows from the fact that the best response in (49) converges to the best response in (31). □

The proposition has a straightforward intuition. In the limit of small productivity shocks, it does not matter if perturbations emerge from the consumption or housing sector. Hence, Proposition 3 applies in this case as well, and consumption productivity drives the same broad-based comovement among aggregates.

The important difference with respect to our baseline model is that, in this case, our mechanism is amplifying an otherwise smaller demand driven fluctuation in the housing market. In other words, under perfect information a shock to consumption productivity would already translate into a smaller, but still correlated, movement in business cycle variables. To see this, rewrite aggregate consumption of residential investment and consumption in the case of perfect information: $c = \zeta_t^c$ and $\delta_t = -\lambda_t - p = (1-\gamma)\zeta_t^c$, which says that residential investment, the price of new housing and consumption move together even under perfect information. Therefore, our baseline of aggregate shocks to housing productivity has the merit of showing that our mechanism can both generate strong amplification of fundamental shocks and dramatically change the qualitative transmission of shocks in the economy.

5.2 Signal extraction with private signals

Here we show that the signal extraction problem, and corresponding equilibria, are not qualitatively affected by the availability of a private signal about the local shock. Instead, the addition of private information maps into our analysis of Section 3.3 as an increase in the relative variance of aggregate shocks.

Let us assume that a household $j \in (0,1)$ in island i has a private signal

$$\omega_{ij} = \mu_i + \eta_{ij} \tag{50}$$

where $\eta_{ij} \sim N(0, \sigma_\eta)$ is identically and independently distributed across households and islands. In this case, households form expectations according to

$$E[\mu_i|p_i,\omega_{ij}] = a\left(\gamma\mu_i + (1-\gamma)\left(\int E[\mu_i|p_i,\omega_{ij}]di - \zeta\right)\right) + b\left(\mu_i + \eta_{ij}\right),$$

where b measures the weight given to the additional private signal. Averaging out the relation above and solving for the aggregate expectation gives

$$\int E[\mu_i|p_i,\omega_{ij}]di = -\frac{a\,(1-\gamma)}{1-a\,(1-\gamma)}\zeta,$$

which is identical to (28). However, now we need two optimality restrictions to determine a

and b. These are

$$E[p_i(\mu_i - E[\mu_i|p_i,\omega_{ij}])] = 0 \;\;\Rightarrow\;\; \gamma\sigma_\mu - a\left(\gamma^2\sigma_\mu + \frac{(1-\gamma)^2}{(1-a(1-\gamma))^2}\sigma_\zeta\right) - b\gamma\sigma_\mu = 0,$$
$$E[\omega_{ij}(\mu_i - E[\mu_i|p_i,\omega_{ij}])] = 0 \;\;\Rightarrow\;\; \sigma_\mu - a\gamma\sigma_\epsilon - b(\sigma_\mu + \sigma_\eta) = 0,$$

which identify the equilibrium a and b such that each piece of information is orthogonal with the forecast error. Solving the system for a, we get a fix point equation written as

$$a = \frac{\gamma}{\gamma^2 + \frac{(1-\gamma)^2}{(1-a(1-\gamma))^2}\frac{\sigma_\mu+\sigma_\eta}{\sigma_\eta}\frac{\sigma_\zeta}{\sigma_\mu}}. \tag{51}$$

For $\sigma_\eta \to \infty$, the right-hand side of the relation above matches (31). In particular, it follows that a lower σ_η in (51) is equivalent to considering a larger σ_ζ in (31). The analysis of the baseline model thus applies directly to this generalization, and small amounts of exogenous private information do not qualitatively change any of our earlier results.

5.3 Relation with sentiments

Authors such as Benhabib et al. (2015) have found that *extrinsic* (non-fundamental) sentiment shocks may emerge in environments with endogenous signals. A natural question, given the results in Proposition 2, is whether any equilibria exist in which errors are driven by extrinsic shocks in addition or instead of productivity. The next proposition states that, in fact, *extrinsic* sentiments are always crowded-out by common shocks to productivity.

Proposition 5. *Suppose that*

$$\int E[\mu_i|p_i]di = \phi_\zeta\zeta + \phi_\varepsilon\varepsilon,$$

where ϕ_ε is the equilibrium effect of an extrinsic sentiment shock, $\varepsilon \sim N(0,\sigma_\varepsilon^2)$, not related to fundamentals. Then, $\phi_\varepsilon = 0$ for any $\sigma^2 > 0$.

Proof. Given in Appendix C. □

Fundamental shocks always dominates extrinsic shocks because the former have two channels — one endogenous and one exogenous — through which they influences people's information. Intuitively, conjecture that the average action reflects a response to both fundamental and extrinsic shocks. In equilibrium, agents respond to the average expectation, and therefore proportionally to the conjectured endogenous coefficients for each shock. But agents also respond to the exogenous component of the fundamental that appears in the price signal. Thus, any equilibrium must depend somewhat more-than-conjectured on the fundamental relative to the extrinsic shock. This guess and update procedure cannot converge unless the weight on the extrinsic shock is zero.

This logic highlights the fragility of the extrinsic version of sentiments, which are coordinated by endogenous signal structures. For, any shock which tends to coordinate actions for exogenous reasons will also benefit from the self-reinforcing nature of learning, thereby absorbing the role of belief shock for itself. Indeed, the same results emerge if *local* shocks μ_i have any common component, as we consider in Section 5.1.

5.4 Stability analysis

Here, we examine the issue of out-of-equilibrium convergence, that is, whether or not an equilibrium is a rest point of a process of revision of beliefs in a repeated version of the static economy. We suppose that agents behave like econometricians. At time t they set a weight $a_{i,t}$ that is estimated from the sample distribution of observables collected from past repetitions of the economy.

Agents learn about the optimal weight according to an optimal adaptive learning scheme:

$$a_{i,t} = a_{i,t-1} + \gamma_t \; S_{i,t-1}^{-1} \; p_{i,t} \left(\mu_{i,t} - a_{i,t-1} p_{i,t}\right) \tag{52}$$

$$S_{i,t} = S_{i,t-1} + \gamma_{t+1} \left(p_{i,t}^2 - S_{i,t-1}\right), \tag{53}$$

where γ_t is a decreasing gain with $\sum \gamma_t = \infty$ and $\sum \gamma_t^2 = 0$, and matrix $S_{i,t}$ is the estimated variance of the signal. A rational expectations equilibrium $\hat{a}$ is a locally learnable equilibrium if and only if there exists a neighborhood $\digamma(\hat{a})$ of $\hat{a}$ such that, given an initial estimate $a_{i,0} \in \digamma(\hat{a})$, then $\lim_{t \to \infty} a_{i,t} \overset{a.s}{=} \hat{a}$; it is a globally learnable equilibrium if convergence happens for any $a_{i,0} \in \mathbb{R}$.

The asymptotic behavior of statistical learning algorithms can be analyzed by stochastic approximation techniques (see Marcet and Sargent, 1989a,b; Evans and Honkapohja, 2001, for details.) Below we show that the relevant condition for stability is $a_i'(a) < 1$, which can easily checked by inspection of Figure 2.

Proposition 6. *For $\gamma > 1/2$ the unique equilibrium a_u is globally learnable. For $\gamma < 1/2$ the "low" and the "high" equilibrium, a_- and a_+, respectively, are always locally learnable, whereas the middle equilibrium $a_\circ$ is never.*

Proof. Given in Appendix C. □

It turns out that the unique equilibrium is globally learnable: revisions will lead agents to coordinate on the equilibrium regardless of initial beliefs. With multiplicity, the "high" and "low" equilibrium are locally learnable, whereas the middle equilibrium is not. Instead, the middle equilibrium works as a frontier between the basins of attraction of the "low" and "high" equilibria.

6 Conclusion

Learning from prices has played an important role in our understanding of financial markets since at least Grossman and Stiglitz (1980). Yet, learning from prices appeared even earlier in the macroeconomics literature, including in the seminal paper of Lucas (1972). Nevertheless, that channel gradually disappeared from models of the business cycle, in large part because people concluded that fundamental shocks would be effectively revealed before incomplete knowledge about them could influence relatively slow-moving macroeconomic aggregates.

In this paper we have shown that, even if aggregate shocks are *nearly* common knowledge, learning from prices may still play a crucial role driving fluctuations in beliefs. In fact, the feedback mechanism we described may be strongest precisely when the aggregate shock is almost, but not-quite-fully, revealed. Endogenous information structures can deliver strong multipliers on small common disturbances, and thus offer a foundation for coordinated, expectations-driven economic fluctuations that are entirely rational. Moreover, the key feature of our theory is also a feature of reality: agents observe and draw inference from prices that are, themselves, influenced by aggregate conditions.

We have applied this idea to house prices, because these are among the most salient prices in the economy. Even if the economy is driven only by productivity shocks, this mechanism captures several salient features of business cycles and its close correlation with the housing market while remaining consistent with the evidence that productivity and endogenous outcomes are weakly correlated. Hence, our results suggest that the relationship between supply and demand shocks is more subtle than typically assumed in the empirical literature and future empirical work may wish to take in account the implications of price-based learning.

7 Acknowledgments

We are grateful to the Editor, Veronica Guerrieri, and three anonymous referees for their insightful guidance in revising the paper. We also thank Elena Afanasyeva, Jess Benhabib, Edouard Challe, Mehmet Ekmekci, Roger Guesnerie, Christian Hellwig, Peter Ireland, Pierre-Olivier Wiell, Jianjun Miao, Patrick Pintus, Kristoffer Nimark, Richard Tresch, Robert Ulbricht, Rosen Valchev, Laura Veldkamp, Venky Venkateswaran, Xavier Vives, and seminar participants at the Toulouse School of Economics, UCLA, Cornell University, Barcelona GSE Summer Forum, Society for Economic Dynamics, and Boston Green Line Macro conferences for valuable suggestions and comments. Laura Veronika Gáti and Serge Bechara provided excellent research assistance.

8 Funding

The research leading to these results has received financial support from the European Research Council under the European Community's Seventh Framework Program FP7/2007-2013 grant agreement No.263790 and the *Laboratoire d'excellence* in Economics and Decision Sciences (LabEx EDODEC).

9 Data Availability

The data underlying this article are available in Zenodo at https://doi.org/10.5281/zenodo.4110947.

References

Adam, K., P. Kuang, and A. Marcet (2011): "House Price Booms and the Current Account," in *NBER Macroeconomics Annual 2011, Volume 26*, National Bureau of Economic Research, Inc, NBER Chapters, 77–122.

Albagli, E., C. Hellwig, and A. Tsyvinski (2014): "Dynamic Dispersed Information and the Credit Spread Puzzle," NBER Working Papers 19788, National Bureau of Economic Research, Inc.

Amador, M. and P.-O. Weill (2010): "Learning From Prices: Public Communication and Welfare," *The Journal of Political Economy*, 118, pp. 866–907.

Angeletos, G.-M. (2018): "Frictional Coordination," *Journal of the European Economic Association*, 16, 563–603.

Angeletos, G.-M., F. Collard, and H. Dellas (2018): "Quantifying confidence," *Econometrica*, 86, 1689–1726.

——— (2020): "Business-Cycle Anatomy," *American Economic Review*, 110, 3030–70.

Angeletos, G.-M. and J. La'O (2009): "Noisy Business Cycles," in *NBER Macroeconomics Annual 2009, Volume 24*, National Bureau of Economic Research, Inc, NBER Chapters, 319–378.

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

Angeletos, G.-M. and C. Lian (2020): "Confidence and the Propagation of Demand Shocks," Working Paper 27702, National Bureau of Economic Research.

Attanasio, O. P., L. Blow, R. Hamilton, and A. Leicester (2009): "Booms and Busts: Consumption, House Prices and Expectations," *Economica*, 76, 20–50.

Attanasio, O. P. and G. Weber (1994): "The UK Consumption Boom of the Late 1980s: Aggregate Implications of Microeconomic Evidence," *The Economic Journal*, 104, 1269–1302.

Barlevy, G. and P. Veronesi (2000): "Information Acquisition in Financial Markets," *The Review of Economic Studies*, 67, 79–90.

Barro, R. J. and R. G. King (1984): "Time-Separable Prefrences and Intertemporal-Substitution Models of Business Cycles," *Quarterly Journal of Economics*, 99, 817 – 839.

Basu, S., J. Fernald, J. Fisher, and M. Kimball (2014): "Sector-specific Technical Change," Working paper.

Basu, S., J. G. Fernald, and M. S. Kimball (2006): "Are Technology Improvements Contractionary?" *The American Economic Review*, 96, pp. 1418–1448.

Benhabib, J., P. Wang, and Y. Wen (2015): "Sentiments and Aggregate Demand Fluctuations," *Econometrica*, 83, 549–585.

Benhima, K. (2019): "Booms and busts with dispersed information," *Journal of Monetary Economics*, 107, 32 – 47.

Benhima, K. and I. Blengini (2020): "Optimal Monetary Policy when Information is Market-Generated," *The Economic Journal*, 130, 956–975.

Berger, D., V. Guerrieri, G. Lorenzoni, and J. Vavra (2017): "House Prices and Consumer Spending," *The Review of Economic Studies*, 85, 1502–1542.

Buiter, W. H. (2010): "Housing wealth isn't wealth," *Economics - The Open-Access, Open-Assessment E-Journal*, 4, 1–29.

Burguet, R. and X. Vives (2000): "Social Learning and Costly Information Acquisition," *Economic Theory*, 15, 185–205.

Burnside, C., M. Eichenbaum, and S. Rebelo (2016): "Understanding Booms and Busts in Housing Markets," *Journal of Political Economy*, 124, 1088–1147.

Campbell, J. Y. and J. F. Cocco (2007): "How do house prices affect consumption? Evidence from micro data," *Journal of Monetary Economics*, 54, 591 – 621.

Cass, D. and K. Shell (1983): "Do Sunspots Matter?" *Journal of Political Economy*, 91, pp. 193–227.

Chahrour, R. and G. Gaballo (2017): "Learning from Prices: Amplification and Business Fluctuations," Working Paper Series 2053, European Central Bank.

Chahrour, R. and K. Jurado (2018): "News or Noise? The Missing Link," *American Economic Review*, 108, 1702–36.

——— (2019): "Recoverability," Working paper.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (2005): "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 113, 1–45.

Davis, M. A. and J. Heathcote (2005): "Housing and the Business Cycle," *International Economic Review*, 46, 751–784.

——— (2007): "The Price and Quantity of Residential Land in the United States," *Journal of Monetary Economics*, 54, 2595 – 2620.

Eusepi, S. and B. Preston (2011): "Expectations, Learning, and Business Cycle Fluctuations," *American Economic Review*, 101, 2844–72.

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

Fehrle, D. (2019): "Housing and the Business Cycle Revisited," *Journal of Economic Dynamics and Control*, 99, 103 – 115.

Gaballo, G. (2016): "Rational Inattention to News: The Perils of Forward Guidance," *American Economic Journal: Macroeconomics*, 8, 42–97.

——— (2018): "Price Dispersion, Private Uncertainty, and Endogenous Nominal Rigidities," *The Review of Economic Studies*, 85, 1070–1110.

Galí, J. (1999): "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review*, 89, 249–271.

Gazzani, A. (2019): "News and Noise Bubbles in the Housing Market," *Review of Economic Dynamics*.

Grossman, S. J. and J. E. Stiglitz (1976): "Information and Competitive Price Systems," *The American Economic Review*, 66, 246–253.

——— (1980): "On the Impossibility of Informationally Efficient Markets," *The American Economic Review*, 70, pp. 393–408.

Hassan, T. A. and T. M. Mertens (2017): "The Social Cost of Near-Rational Investment," *American Economic Review*, 107, 1059–1103.

Hellwig, M. F. (1980): "On the Aggregation of Information in Competitive Markets," *Journal of Economic Theory*, 22, 477–498.

Iacoviello, M. (2005): "House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle," *American Economic Review*, 95, 739–764.

Iacoviello, M. and S. Neri (2010): "Housing Market Spillovers: Evidence from an Estimated DSGE Model," *American Economic Journal: Macroeconomics*, 2, 125–64.

Ilut, C. L. and M. Schneider (2014): "Ambiguous Business Cycles," *American Economic Review*, 104, 2368–99.

Jaimovich, N. and S. Rebelo (2009): "Can News About the Future Drive the Business Cycle?" *American Economic Review*, 99, 1097 – 1118.

Jorgensen, D. W., M. S. Ho, and J. Samuels (2012): "A Prototype Industry-level Production Account for the United States, 1947-2010," in *Second World KLEMS Conference, Harvard University.*, vol. 9.

Kaplan, G., K. Mitman, and G. L. Violante (2017): "The Housing Boom and Bust: Model Meets Evidence," Working Paper 23694, National Bureau of Economic Research.

Karabarbounis, L. (2014): "The Labor Wedge: MRS Vs. MPN," *Review of Economic Dynamics*, 17, 206 – 223.

King, M. (1990): "(Discussion of) Is the UK Balance of Payments Sustainable?" *Economic Policy*, 5, 348–395.

Leamer, E. E. (2007): "Housing IS the Business Cycle," Working Paper 13428, National Bureau of Economic Research.

——— (2015): "Housing Really is the Business Cycle: What Survives the Lessons of 2008–09?" *Journal of Money, Credit and Banking*, 47, 43–50.

L'Huillier, J.-P. (2020): "Consumer Imperfect Information and Endogenous Price Rigidity," *American Economic Journal: Macroeconomics*, 12, 94–123.

Lichter, A., A. Peichl, and S. Siegloch (2015): "The own-wage elasticity of labor demand: A meta-regression analysis," *European Economic Review*, 80, 94 – 119.

Liu, Z., J. Miao, and T. Zha (2016): "Land prices and unemployment," *Journal of Monetary Economics*, 80, 86–105.

Lorenzoni, G. (2009): "A Theory of Demand Shocks," *American Economic Review*, 99.

Lucas, R. E. (1972): "Expectations and the Neutrality of Money," *Journal of Economic Theory*, 4, 103 – 124.

Manzano, C. and X. Vives (2011): "Public and Private Learning from Prices, Strategic Substitutability and Complementarity, and Equilibrium Multiplicity," *Journal of Mathematical Economics*, 47, 346–369.

Marcet, A. and T. J. Sargent (1989a): "Convergence of Least-Squares Learning in Environments with Hidden State Variables and Private Information," *Journal of Political Economy*, 97, pp. 1306–1322.

——— (1989b): "Convergence of Least Squares Learning Mechanisms in Self-referential Linear Stochastic Models," *Journal of Economic Theory*, 48, 337 – 368.

MIAN, A., K. RAO, AND A. SUFI (2013): "Household Balance Sheets, Consumption, and the Economic Slump," *The Quarterly Journal of Economics*, 128, 1687–1726.

MIAN, A. AND A. SUFI (2014): "What Explains the 2007–2009 Drop in Employment?" *Econometrica*, 82, 2197–2223.

MILLER, N., L. PENG, AND M. SKLARZ (2011): "House Prices and Economic Growth," *The Journal of Real Estate Finance and Economics*, 42, 522–541.

MUELLBAUER, J. AND A. MURPHY (1990): "Is the UK Balance of Payments Sustainable?" *Economic Policy*, 5, 348–395.

NGUYEN, Q. H. (2018): "Revisiting Housing and the Business Cycle," *Journal of Housing Economics*, 41, 85 – 92.

NIMARK, K. (2008): "Monetary Policy With Signal Extraction from the Bond Market," *Journal of Monetary Economics*, 55, 1389–1400.

PAGANO, M. (1990): "(Discussion of) Is the UK Balance of Payments Sustainable?" *Economic Policy*, 5, 348–395.

SHILLER, R. J. (2007): "Understanding Recent Trends in House Prices and Homeownership," *Proceedings - Economic Policy Symposium - Jackson Hole*, 89–123.

VENKATESWARAN, V. (2013): "Heterogeneous Information and Labor Market Fluctuations," Working Paper.

VIVES, X. (2014): "On The Possibility Of Informationally Efficient Markets," *Journal of the European Economic Association*, 12, 1200–1239.

A Model: general case

We present our baseline model, extended to include convex costs of supplying labor, preference shocks, and shocks to productivity in the consumption sector.

The representative household living on island i solves

$$\max_{C_{it},\Delta_{it},B_{it},N_{it}^c,N_{it}^h} E_0 \sum_{t=0}^{\infty} \beta^t \left(\log \left(C_{it}^{\phi} \mathcal{H}_{it}^{1-\phi} \right)^{\Theta_{it}} - v^c \frac{\left(N_{it}^c\right)^{1+\chi_c}}{1+\chi_c} - v^h \frac{\left(N_{it}^h\right)^{1+\chi_h}}{1+\chi_h} \right)$$

subject to the constraints that (i) $\{C_{it}, \Delta_{it}\}$ can only depend on the information set $\{P_{it}, \Omega_{it-1}\}$; (ii) $\{B_{it}, N_{it}^c, N_{it}^h\}$ can only depend on Ω_{it}; and (iii) the budget constraint,

$$\frac{1}{R_t} B_{it} + C_{it} + P_{it}\Delta_{it} \leq W_{it}^c N_{it}^c + W_{it}^h N_{it}^h + B_{it-1} + \Pi_{it}^c + \Pi_t^h,$$

must hold state-by-state.

In the above problem, the household may be hit by an island-specific "taste" shock Θ_{it}, which we use to demonstrate that consumers' forecast errors are equivalent to taste shocks so defined. The convex disutility of labor in each sector is parametrized by χ_c and χ_h. Since workers are imperfectly mobile across sectors, we track sector-specific wages W_{it}^c and W_{it}^h and drop the labor market clearing condition.

In the competitive consumption sector we allow for an aggregate productivity shock and decreasing return to scale. The production function is

$$Y_t = e^{\tilde{\zeta}_t^c} (N_t^c)^{\alpha_c} \quad \text{with} \quad N_t^c \equiv \left(\int e^{\tilde{\mu}_{it}/\eta} {N_{it}^c}^{1-\frac{1}{\eta}} di \right)^{\frac{1}{1-\frac{1}{\eta}}}.$$

Produtivity $\tilde{\zeta}_t^c = \tilde{\zeta}_{t-1}^c + \zeta_t^c$ where ζ_t^c is an iid innovation drawn from a normal distribution $N\left(0, \sigma_{\zeta^c}\right)$ and $\alpha_c \in (0,1)$ measures returns to scale. We denote by W_t^c the price of N_t^c such that $W_t^c N_t^c = \int W_{it}^c N_{it}^c di$.

We allow for exogenous variation in the supply Z_t to clarify that housing productivity shocks are isomorphic to changes in the supply the input. Market clearing for the traded input is requires

$$Z_t = \int Z_{it} di.$$

Remaining portions of the model are the same as in the main text. The model in the main text is nested here by setting $\chi_c = \chi_h = 0$, $v^c = v^h$, $\sigma_{\zeta^c} = 0, \alpha_c = 1$, and $Z_t = Z$.

A.1 Complete list of equilibrium conditions

We list here all of the necessary equilibrium conditions at a given time t. The first order conditions for the household are:

$$
\begin{aligned}
\Lambda_{it} &= \beta E_t[\Lambda_{i,t+1} R_t | \Omega_{it}] \\
W_{it}^c &= \Lambda_{it}^{-1} (N_{it}^c)_c^{\chi} \\
W_{it}^h &= \Lambda_{it}^{-1} \left(N_{it}^h\right)^{\chi_h} \\
\Theta_{it} \phi C_{it}^{-1} &= E_t[\Lambda_{it} | P_{it}, \Omega_{it-1}], \\
\Theta_{it} \frac{(1-\psi)(1-\phi)}{(1-(1-d)\beta\psi)^{-1}} \Delta_{it}^{-1} &= E_t[\Lambda_{it} P_{it} | P_{it}, \Omega_{it-1}]
\end{aligned}
$$

where we use the fact that

$$
\frac{\partial U_{i0}}{\partial \Delta_{it}} = (1-\psi)(1-\phi)\Theta_{it} \sum_{\tau=t}^{\infty} ((1-d)\beta\psi)^{\tau-t} \Delta_{it}^{-1} = \Theta_{it} \frac{(1-\psi)(1-\phi)}{1-(1-d)\beta\psi} \Delta_{it}^{-1}.
$$

The budget constraint

$$
\frac{1}{R_t} B_{it} + C_{it} + P_{it}\Delta_{it} = W_{it}^c N_{it}^c + W_{it}^h N_{it}^h + B_{it-1} + \Pi_t^c + \Pi_{it}^h
$$

holds with equality and the transversality condition

$$
\lim_{\tau \to \infty} E\left[\prod_{\kappa=0}^{\tau} R_{t+\kappa} B_{it+\tau} | \Omega_{it}\right] = 0,
$$

must hold at the individual level.

The conditions for optimality in the consumption sector are:

$$
\begin{aligned}
N_{it}^c &= e^{\tilde{\mu}_{it}} \left(\frac{W_{it}^c}{W_t^c}\right)^{-\eta} N_t^c \\
N_t^c W_t^c &= \alpha_c Y_t \\
Y_t &= e^{\zeta_t^c} (N_t^c)^{\alpha_c}.
\end{aligned}
$$

The conditions for optimality in the housing sector are:

$$
\begin{aligned}
Z_{it} Q_t &= \alpha(1-\gamma) P_{it} \Delta_{it} \\
N_{it}^h W_{it}^h &= \gamma \alpha P_{it} \Delta_{it} \\
V_{it} L_{it} &= (1-\alpha) P_{it} \Delta_{it} \\
\Delta_{it} &= L_{it}^{1-\alpha} \left(\left(N_{it}^h\right)^{\phi} \left(e^{-\tilde{\zeta}_t} Z_{it}\right)^{1-\phi} \right)^{\alpha}.
\end{aligned}
$$

Finally, market clearing requires

$$
L_{it} = 1, \qquad \int C_{it} di = Y_t, \qquad \int B_{it} di = 0, \quad \text{and} \quad \int Z_{it} di = Z_t.
$$

A.2 Log-linearized model

We now provide the log-linear relations that describe an approximate equilibrium. The first order conditions for the household are:

$$\lambda_{it} = E[\lambda_{it+1}|\Omega_{it}] + r_t \tag{A.1}$$

$$\chi_c n^c_{it} = \lambda_{it} + w^c_{it} \tag{A.2}$$

$$\chi_h n^h_{it} = \lambda_{it} + w^h_{it} \tag{A.3}$$

$$\theta_{it} - c_{it} = E[\lambda_{it}|p_{it}] \tag{A.4}$$

$$\theta_{it} - \delta_{it} = E[\lambda_{it}|p_{it}] + p_{it}, \tag{A.5}$$

The conditions for the consumption sector are:

$$n^c_{it} = \tilde{\mu}_{it} - \eta\left(w^c_{it} - w^c_t\right) + n^c_t \tag{A.6}$$

$$n^c_t + w^c_t = y_t \tag{A.7}$$

$$y_t = \tilde{\zeta}^c_t + \alpha_c n^c_t. \tag{A.8}$$

The conditions for the housing sector are:

$$z_{it} + q_t = p_{it} + \delta_{it}, \tag{A.9}$$

$$n^h_{it} + w^h_{it} = p_{it} + \delta_{it} \tag{A.10}$$

$$v_{it} = p_{it} + \delta_{it} \tag{A.11}$$

$$\delta_{it} = (1-\alpha)l_i + (\alpha\gamma)n^h_{it} + \alpha\left(1-\gamma\right)\left(z_{it} - \tilde{\zeta}_t\right). \tag{A.12}$$

Only the budget constraint must be approximated. Using expressions for profits, we have

$$\frac{B_{it}}{R_t} + C_{it} + P_{it}\Delta_{it} + P_{it}H_{it-1} = W^c_{it}N^c_{it} + W^h_{it}N^h_{it} + P_{it}H_{it-1} + B_{it-1} + \underbrace{Y_t - W^c_t N^c_t}_{\Pi^c_t} + \underbrace{P_{it}H_{it} - W^h_{it}N^h_{it} - Q_t\left(Z_{it} - Z_t\right) + V_{it}}_{\Pi^h_{it}}.$$

This simplifies to

$$\frac{B_{it}}{R_t} + (C_{it} - Y_t) - (W^c_{it}N^c_{it} - W^c_t N^c_t) = -Q_t\left(Z_{it} - Z_t\right) + B_{it-1}.$$

We consider a linearization around a non-stochastic steady-state in which $B_{it} = 0$ for all i, hence we linearize around B_{it} and log-linearize for other variables. In such a steady-state, the terms in parenthesis above are zero, so that the linearization is

$$\beta b_{it} + C(c_{it} - c_t) = C(w^c_{it} - w^c_t) + C(n^c_{it} - n^c_t) - Q(z_{it} - z_t) + b_{it-1}, \tag{A.13}$$

where capital letters denote steady states values.

Finally, market clearing conditions are:

$$0 = \int b_{it}di \tag{A.14}$$

$$z_t = \int z_{it}di \tag{A.15}$$

$$c_t = y_t. \tag{A.16}$$

Remark 1. *Inspection of the first order conditions (A.1)-(A.16) shows that consumers' forecast errors are equivalent to a shock to the intratemporal margin, in particular, an individual consumption-housing taste shock. In particular, note that any equilibrium in the incomplete information economy without taste shocks can be implemented in a fictitious full information economy in which taste shocks are equal to the forecast errors of the corresponding incomplete information economy, i.e.* $\theta_{it} \equiv E[\lambda_{it}|p_{it}] - \lambda_{it}$.

B Equilibrium

This section shows the analytical solution of the extended model. In light of Remark 1, we ignore taste shocks going forward. We also generalize our information structure by introducing news about future aggregate productivity, in line with the extend model in section 5.1. We focus on time t and we assume that in the second stage the worker-saver i knows the current housing productivity, current and future consumption productivity and local productivity, i.e. $\Omega_{it} = \{\tilde{\mu}_{it+1}, \zeta_t, \zeta_t^c, \zeta_{t+1}^c\} \cup \Omega_{it-1}$. We continue to assume that shoppers only observe $\{p_{it}, \Omega_{it-1}\}$ at time t.

B.1 Expectations of the saver-worker from $t+1$ onwards

Here we characterize the equilibrium of the economy from $t+1$ onwards, conditional on the information set of the worker-savers at time t. Throughout, we make extensive use of the law of iterated expectations, especially the result

$$E[E[\lambda_{it+j}|p_{it+1}, \Omega_{it}]|\Omega_{it}] = E[\lambda_{it+j}|\Omega_{it}] \quad \text{for all} \quad j \geq 1.$$

More generally, we denote $E_t[x_{it+j}] \equiv E[x_{it+j}|\Omega_{it}]$ to capture worker expectations of any future variable x_{it+j}, letting the integer j span future horizons $j \geq 1$.

Expectations of aggregates

Equation (A.7) combined with market clearing in the consumption market implies

$$c_t - w_t^c = n_t^c \quad \text{for all } t,$$

while combining and aggregating (A.2) and (A.4) implies that in expectation

$$E_t[c_{t+j} - w_{t+j}^c] = -\chi_c E_t[n_{t+j}^c].$$

Comparing the two equations just above, we can conclude that

$$E_t[n_{t+j}^c] = 0.$$

Combining $E_t[n_{t+j}^c] = 0$ with (A.7) and (A.8), we can then establish that

$$E_t[c_{t+j}] = E_t[w_{t+j}^c] = E_t[\tilde{\zeta}_{t+j}^c]. \tag{B.1}$$

Local budget constraint

Combining first order conditions (A.5) and (A.9) one finds that:

$$z_{it} - z_t = c_{it} - c_t. \tag{B.2}$$

Plugging (A.6) and (B.2) into the budget constraint we have:

$$(C+Q)(c_{it}-c_t)+\beta b_{it}=C\tilde{\mu}_{it}+(1-\eta)C(w^c_{it}-w^c_t)+b_{it-1}. \tag{B.3}$$

Similarly, we can use equations (A.2) and (A.6) to relate island and aggregate labor,

$$\chi_c(n_{it}-n^c_t)=(\lambda_{it}-\lambda_t)+(w^c_{it}-w^c_t)$$
$$\chi_c(n_{it}-n^c_t)=\chi_c\tilde{\mu}_{it}-\eta\chi_c(w^c_{it}-w^c_t)$$

Use the consumption demand condition in (A.4) to eliminate Lagrange multipliers, combine the above two equations and take expectations to get

$$E_t[w^c_{it+j}-w^c_{t+j}]=\frac{\chi_c}{1+\eta\chi_c}\tilde{\mu}_{it+j}+\frac{1}{1+\eta\chi_c}(c_{it+j}-c_{t+j}). \tag{B.4}$$

Taking expectations of (B.3), substituting in expression (B.4), and simplifying yields

$$\underbrace{\left(C\frac{\eta(1+\chi_c)}{1+\eta\chi_c}+Q\right)E_t[c_{it+j}-c_{t+j}]-C\left(\frac{1+\chi_c}{1+\eta\chi_c}\right)E_t[\tilde{\mu}_{it+j}]}_{\equiv\Delta\tilde{c}_{it+j}}+\beta E_t[b_{it+j}]=E_t[b_{it+j-1}]. \tag{B.5}$$

Using the definition of $\Delta\tilde{c}_{it+j}$ above, this reduces to

$$\Delta\tilde{c}_{it+j}+\beta E_t[b_{it+j}]=E_t[b_{it+j-1}]. \tag{B.6}$$

Use of the Euler equation and transversality

A first observation involves the Euler equation (A.1). Subtracting (A.1) from its aggregated version establishes that

$$E_t[c_{it+j+1}-c_{t+j+1}]=E_t[c_{it+1}-c_{t+1}] \tag{B.7}$$

Moreover, since local productivity is a random walk, we also have that

$$E[\tilde{\mu}_{it+j}|\Omega_{it}]=\tilde{\mu}_{it+1}$$

for for any $\tau\geq t$. Hence, the $\Delta\tilde{c}_{it+j}$ term in (B.6) is constant across all horizons $j\geq 1$. Calling this constant value Δc, equation (B.6) can be solved forward to find

$$E_t[b_{it+j-1}]=\frac{1}{1-\beta}\Delta c.$$

Since this equation holds for all $j\geq 1$, bonds holdings must be expected to be constant going forward, i.e.

$$b_{it+j}=b_{it}.$$

This is the unique equilibrium path for bonds, since any other solution satisfying (B.6) implies expected bond holdings grow unboundedly over time, violating transversality.

Derivation of $E[\lambda_{it+1}|\Omega_{it}]$

Using $E_t[b_{it+j}]=b_{it}$ in equation (B.5) and $E_t[c_{t+1}]=\tilde{\zeta}^c_{t+1}$, and solving for $E_t[c_{it+1}]$ we get

$$E_t[\lambda_{it+1}]=-E_t[c_{it+1}]=-\omega_\mu\tilde{\mu}_{it+1}-\omega_b b_{it}-\tilde{\zeta}^c_{t+1},$$

with

$$\omega_\mu = \frac{C\left(\frac{1+\chi_c}{1+\eta\chi_c}\right)}{C\frac{\eta(1+\chi_c)}{1+\eta\chi_c}+Q} > 0, \quad \text{and} \quad \omega_b = \frac{1-\beta}{C\frac{\eta(1+\chi_c)}{1+\eta\chi_c}+Q} > 0.$$

As stated in the main text, notice that $\lim_{\beta\to 1}\omega_b = 0$.

B.2 Equilibrium at time t

Derivation of λ_{it} (Lemma 1)

Our first objective is finding the equilibrium mapping from fundamentals to b_{it}. Let us find their common component λ_t. One can use the aggregate version of (A.2),(A.7) and (A.8) to get $\chi_c n_t^c = \lambda_t + w_t^c$ and $w_t^c = \tilde{\zeta}_t^c + (\alpha_c - 1)n_t^c$ to get

$$(1-\alpha_c+\chi_c)n_t^c \quad = \quad \lambda_t + \tilde{\zeta}_t^c.$$

Combining this with (A.8) gives a relation between the realized aggregate lambda and shoppers' expectations

$$\lambda_t = -\frac{1-\alpha_c+\chi_c}{\alpha_c}\int E[\lambda_{it}|p_{it}]di - \frac{1+\chi_c}{\alpha_c}\zeta_t^c.$$

Note that this expression is valid also for future times. Using the law of iterated epectations, we have that $E_t[\lambda_{t+1}] = -\zeta_{t+1}^c$ which is consistent with what we have found above. In this case, the aggregate version of the Euler equation (A.1) implies,

$$r_t = \lambda_t - E[\lambda_{t+1}|\Omega_{it}] = -\frac{1-\alpha_c+\chi_c}{\alpha_c}\int E[\lambda_{it}|p_{it}]di - \frac{1+\chi_c}{\alpha_c}\zeta_t^c + \zeta_{t+1}^c.$$

Given that (A.1) must also hold at the local level, we then have the following

$$\begin{aligned}\lambda_{it} &= \underbrace{-\omega_\mu\tilde{\mu}_{it+1} - \omega_b b_{it} - \zeta_{t+1}^c}_{=E[\lambda_{it+1}|\Omega_{it}]} + r_t \\ &= -\omega_\mu\tilde{\mu}_{it+1} - \omega_b b_{it} + \frac{1-\alpha_c+\chi_c}{\alpha_c}c_t - \frac{1+\chi_c}{\alpha_c}\zeta_t^c, \qquad \text{(B.8)}\end{aligned}$$

which reduces to expression (21) of Lemma 1 under our baseline assumptions that $\alpha_c = 1$, $\chi_c = 0$, and $\zeta_t^c = 0$.

Remark 2. *Equation (B.8) shows that the anticipation of future consumption productivity does not affect the marginal valuation of current consumption. This is a standard finding in real business cycle models, since real interest rates neutralize the effect of anticipated productivity changes. By contrast, current productivity does move the marginal valuation of current consumption.*

Price of new housing

Here we derive the expression for the equilibrium price of new housing. By using (A.3), (A.5), (A.9), and (A.10) we get

$$\begin{aligned} p_{it} + \delta_{it} &= -E[\lambda_{it}|p_{it}], \\ z_{it} &= -E[\lambda_{it}|p_{it}] - q_t, \\ n^h_{it} &= \frac{1}{1+\chi_h}\left(\lambda_{it} - E[\lambda_{it}|p_{it}]\right). \end{aligned}$$

The housing price is then

$$\begin{aligned} p_{it} &= -E[\lambda_{it}|p_{it}] - \alpha\gamma n_{it} - \alpha\left(1-\gamma\right)\left(-\zeta_t + z_{it}\right) \\ &= -E[\lambda_{it}|p_{it}] - \alpha\gamma\left(\frac{1}{1+\chi_h}\left(\lambda_{it} - E[\lambda_{it}|p_{it}]\right)\right) - \alpha\left(1-\gamma\right)\left(-\zeta_t - q_t - E[\lambda_{it}|p_{it}]\right) \\ &= \left(1 - \alpha(1-\gamma) - \frac{\alpha\gamma}{1+\chi_h}\right) E[-\lambda_{it}|p_{it}] + \alpha \underbrace{\left(\left(\frac{-\gamma}{1+\chi_h}\right)\lambda_{it} + \left(1-\gamma\right)\left(\zeta_t + q_t\right)\right)}_{=s_i}. \qquad \text{(B.9)} \end{aligned}$$

The case in the text obtains for $\chi_h = 0$.

Remark 3. *By substituting $q_t = -\int E[\lambda_{it}|p_{it}]di - z_t$ in the equation above we clearly see that an increase in productivity (negative ζ_t) in the housing sector is isomorphic to an increase in the endowment in raw capital (positive z_t).*

Price of housing stock

In analogy with (A.1) we can derive the price in consumption units $P^h_{it|v}$ of housing vintage $\Delta_{it|v}$ as

$$P^h_{it|v} = E[\Lambda_{it}|p_{it}]^{-1}\frac{\partial U_{i0}}{\partial \Delta_{it|v}} = \frac{(1-\psi)(1-\phi)\psi^{t-v}}{1-(1-d)\beta\psi}\Delta^{-1}_{it|v}E[\Lambda_{it}|p_{it}]^{-1},$$

for any $v \le t$. Therefore the price of the total stock of housing is given as

$$P^h_{it} = \sum_{v=-\infty}^{t}\frac{P^h_{it|v}\Delta_{it|v}}{H_{it}} = \frac{1-\phi}{1-(1-d)\psi\beta}E[\Lambda_{it}|p_{it}]^{-1}H^{-1}_{it}$$

which in log-terms gives

$$p^h_{it} = -E[\lambda_{it}|p_{it}] - \kappa\delta_{it} = (1-\kappa)E[-\lambda_{it}|p_{it}] + \kappa p_{it}$$

where $\kappa = \bar{\Delta}/\bar{H}$ defines the steady state share of residential investment over existing housing stock.

Savings choice

Finally, to compute b_{it}, we re-consider the same steps leading to (B.3) at time t where we now impose $\mu_{it} = b_{it-1} = 0$ to get:

$$(C+Q)(c_{it} - c_t) + \beta b_{it} = (1-\eta)C(w^c_{it} - w^c_t)$$

which can we rewritten as

$$(C+Q)\left(-E[\lambda_{it}|p_{it}]+\int E[\lambda_{it}|p_{it}]di\right)+\beta b_{it}=(1-\eta)C(-\lambda_{it}+\lambda_t).$$

Use the fact $-E[\lambda_{it}|p_{it}]=as_i$, where s_i is defined as above, and (B.8) to get

$$(C+Q)\left(a\frac{\gamma}{1+\chi_h}\left(\omega_\mu\hat{\mu}_{it+1}+\omega_b b_{it}\right)\right)+\beta b_{it}=(1-\eta)C(\omega_\mu\hat{\mu}_{it+1}+\omega_b b_{it}).$$

Rearranging this expression yields

$$\left((C+Q)a\frac{\gamma}{1+\chi_h}\omega_b-(1-\eta)C\omega_b+\beta\right)b_{it}=\left((1-\eta)C-(C+Q)a\frac{\gamma}{1+\chi_h}\right)\omega_\mu\hat{\mu}_{it+1},$$

which can be solved for b_{it}:

$$b_{it}=\frac{-(1+\chi_h)(\eta-1)C-(C+Q)a\gamma}{(C+Q)a\gamma\omega_b+(1+\chi_h)(\eta-1)C\omega_b+(1+\chi_h)\beta}\omega_\mu\hat{\mu}_{it+1}. \tag{B.10}$$

Using this result, we have

$$\begin{aligned}\lambda_t-\lambda_{it} &= \omega_\mu\hat{\mu}_{it+1}+\omega_b b_{it}\\ &= \omega_\mu\hat{\mu}_{it+1}+\omega_b\frac{-(1+\chi_h)(\eta-1)C-(C+Q)a\gamma}{(C+Q)a\gamma\omega_b+(\eta-1)(1+\chi_h)C\omega_b+(1+\chi_h)\beta}\omega_\mu\hat{\mu}_{it+1}\\ &= \frac{(1+\chi_h)\beta}{\underbrace{(C+Q)a\gamma\omega_b+(1+\chi_h)(\eta-1)C\omega_b}_{\equiv f(a,\beta)}+(1+\chi_h)\beta}\omega_\mu\hat{\mu}_{it+1}. \end{aligned} \tag{B.11}$$

Remark 4. *Given that ω_b is a decreasing function of β, we can conclude that a higher a or lower β strictly increases $f(a,\beta)$, and so it strictly decreases the volatility of the idiosyncratic component of λ_{it}, namely $Var(\lambda_t-\lambda_{it})$. This remark will be useful in the following proofs.*

C Proofs of propositions

Proof of Proposition 1. Let us first solve the case for which σ is exogenous and fixed which corresponds to the limit case $\beta\to 1$. The fix point equation reads as

$$a^*(a)=\frac{1}{\gamma}\frac{\tau(a)}{1+\tau(a)}=\frac{1}{\gamma}\frac{1}{1+\left(\frac{(1-\gamma)\sigma}{\gamma(1-a(1-\gamma))}\right)^2}=\frac{\gamma\left(1-a\left(1-\gamma\right)\right)^2}{\gamma^2\left(1-a\left(1-\gamma\right)\right)^2+(1-\gamma)^2\sigma^2} \tag{C.1}$$

To prove uniqueness for $\gamma\geq 1/2$, observe that the function $a^*(a)$ is continuous, bounded above by γ^{-1}, and monotonically decreasing in the range $(0,(1-\gamma)^{-1})$. From $\gamma\geq 1/2$, we have $(1-\gamma)^{-1}>\gamma^{-1}$. Thus $a^*(a)$ intersects the 45-degree line a single time.

To prove the existence of a_-, notice that $\lim_{a\to-\infty}a^*=\gamma^{-1}$ and $a^*((1-\gamma)^{-1})=0$. By continuity, an equilibrium $a_-\in(0,(1-\gamma)^{-1})$ must always exist. Moreover a_- must be monotonically decreasing in σ^2 as a^* is monotonically decreasing in σ^2.

We now assess the conditions under which additional equilibria may also exist. Because $\lim_{a\to\infty}a^*=\gamma^{-1}$, the existence of a second equilibrium (crossing the 45-degree line in Figure 2) implies the existence of a third. Thus, we must determine whether the difference $a^*(a)-a$

is positive anywhere in the range $a > (1-\gamma)^{-1}$. Such a difference is positive if and only if

$$\Phi(\sigma) \equiv \gamma (1 - a(1-\gamma))^2 (1 - \gamma a) - a (1-\gamma)^2 \sigma^2 > 0, \tag{C.2}$$

which requires $a < \gamma^{-1}$ as a necessary condition. Therefore, if two other equilibria exist they must lie in $((1-\gamma)^{-1}, \gamma^{-1})$. Fixing $a \in ((1-\gamma)^{-1}, \gamma^{-1})$, $\lim_{\sigma \to 0} \Phi(\sigma)$ is positive, implying that there always exists a threshold $\bar{\sigma}$, and so a threshold $\bar{\sigma_\zeta}$, such that two equilibria $a_+, a_\circ \in ((1-\gamma)^{-1}, \gamma^{-1})$ exist with $a_+ \geq a_\circ$ for $\sigma^2 \in (0, \bar{\sigma}^2)$.

Let us now consider β less than one. In this case, $\omega_b \neq 0$ and the variance of the idiosyncratic portion of λ_{it} is also endogenous to a, as captured by the function $f(a,\beta)$ in equation (B.11) and noted in Remark 4. Since $\eta > 1$ and $\omega_b > 0$, it follows that $f(a,\beta)$ is strictly positive and increasing in a for all $\beta < 1$. In this case, we must replace σ with the endogenous variance $\sigma(a,\beta)$ in the fixed-point equation (C.1). Since the $\sigma(a,\beta) > \sigma$ and is increasing in a, $a^*(a,\beta)$ is weakly below $a^*(a)$ and any intersection (fixed point) $a^*(\beta)$ must lie strictly to the left of the value a^* for the model with $\beta \to 1$. Hence, if the economy has a unique equilibrium when $\beta \to 1$ it must also have a unique equilibrium $\beta < 1$. Moreover, since $\sigma(a,\beta)$ increases with β, it must be true that the threshold $\bar{\sigma}$ for a multiplicity falls along with β.

□

Proof of Proposition 2. To prove the limiting statement for $\gamma \geq 1/2$, consider any point $a_\delta = \frac{1-\delta}{1-\gamma}$ such that $\delta > 0$. We then have

$$a^*(a_\delta) = \frac{\gamma\delta^2}{\gamma^2\delta^2 + \sigma^2(1-\gamma)^2.} \tag{C.3}$$

Since $\lim_{\sigma^2 \to 0} a^*(a_\delta) = \frac{1}{\gamma}$ for any δ, the unique equilibrium must converge to the same point. That the variance of this equilibrium approaches zero follows from equation (28).

To prove the limiting statement for $\gamma < 1/2$, recall the monotonicity of $a^*(a)$ on the range $(0, (1-\gamma)^{-1})$. Following the logic of Proposition 1, for any point a_δ in that range, $\lim_{\sigma^2 \to 0} a^*(a_\delta) = \gamma^{-1}$, while $a^*((1-\gamma)^{-1}) = 0$. Thus, the intersection defining a_- must approach $(1-\gamma)^{-1}$. An analogous argument for the point just to the right of $(1-\gamma)^{-1}$ establishes that a_- converges to the same value. Finally, the bounded monotonic behavior of $a^*(a)$ establishes that $\lim_{\sigma^2 \to 0} a_+ = \gamma^{-1}$ for the "high" equilibrium.

That the output variance of the "high" equilibrium in the limit $\sigma \to 0$ is zero follows from equation (29). The limiting variance of the two other limit equilibria can be established by noticing that (C.1) implies

$$\frac{(1-\gamma)^2 a^2 \sigma^2}{(1 - a(1-\gamma))^2} = a\gamma(1 - a\gamma) \tag{C.4}$$

which gives (32) for $a \to (1-\gamma)^{-1}$. □

Proof of Proposition 3. Equation (40) establishes that consumption and employment in both sectors move together regardless of the source or variance of shocks. Use (37) and (38) to compute the response of prices to a surprise shock:

$$E[p|\zeta^s] = \frac{\alpha + a(1-\alpha)}{1 - a(1-\gamma)}(1-\gamma)\zeta^s.$$

This shows that house price's dependence on ζ^s has the same sign as consumption's (37), i.e. that house prices always move with consumption. Finally, combining equations (37) and (39) reveals that

$$E[\delta|\zeta^s] = \frac{\alpha(a-1)}{1-a(1-\gamma)}(1-\gamma)\zeta^s,$$

implying that housing investment moves with the other variables if $a > 1$. This holds for sufficiently small σ_s^2 since all limiting equilibria ($a \to (1-\gamma)^{-1}$ or γ^{-1}) satisfy this requirement. □

Proof of Proposition 5. Suppose not, i.e. suppose that

$$\int E[\mu_i|p_i]di = \phi_\zeta \zeta + \phi_\varepsilon \varepsilon,$$

where ϕ_ε is the equilibrium effect of an extrinsic sentiment shock, ε, not related to fundamentals. Then, the price signal is equivalent to

$$p_i = \gamma\mu_i + (1-\gamma)\left((\phi_\zeta + 1)\zeta + \phi_\varepsilon\varepsilon\right)$$

Using the conjectured weights a^*, we have

$$\int a^* p_i di = a(1-\gamma)(\phi_\zeta + 1)\zeta + a(1-\gamma)\phi_\varepsilon\varepsilon$$

implying that

$$\phi_\zeta = a(1-\gamma)(\phi_\zeta + 1)$$
$$\phi_\varepsilon = a(1-\gamma)\phi_\varepsilon$$

which cannot both be true unless $\phi_\varepsilon = 0$. Notice that, differently from the case with multiple sources of signals studied by Benhabib et al. (2015) (section 2.8 page 565), in our case an aggregate shock (our productivity shock) shows up directly in the signal, which ensures determinacy of the average expectation. An implication of this theorem is that the analysis in Benhabib et al. (2015) is not robust to the introduction of correlation (no matter how small) in the v_{jt} shocks appearing in their endogenous signals. □

Proof of Proposition 6. To check local learnability of the rational expectations equilibrium, suppose we are already close to the resting point of the system. That is, consider the case $\int \lim_{t\to\infty} a_{i,t}\, di = \hat{a}$, where $\hat{a}$ is one of the equilibrium points $\{a_-, a_\circ, a_+\}$, and so

$$\lim_{t\to\infty} S_{i,t} = \sigma_s^2(\hat{a}) = \gamma^2\sigma_\mu^2 + \frac{(1-\gamma)^2}{(1-\hat{a}(1-\gamma))^2}\sigma_\zeta^2. \tag{C.5}$$

According to stochastic approximation theory, we can write the associated ODE governing the stability around the equilibria as

$$\begin{aligned}\frac{\mathrm{d}a}{\mathrm{d}t} &= \int \lim_{t\to\infty} \mathrm{E}\left[S_{i,t-1}^{-1} p_{i,t}\left(\mu_{i,t} - a_{i,t-1}p_{i,t}\right)\right] di \\ &= \sigma_s^2\left(\hat{a}\right)^{-1} \int \mathrm{E}\left[p_{i,t}\left(\mu_{i,t} - a_{i,t-1}p_{i,t}\right)\right] di \\ &= \sigma_s^2\left(\hat{a}\right)^{-1}\left(\gamma\sigma_\mu^2 - a_{i,t-1}\left(\gamma^2\sigma_\mu^2 + \frac{\left(1-\gamma\right)^2}{\left(1-a_{t-1}\left(1-\gamma\right)\right)^2}\sigma_\zeta^2\right)\right) \\ &= a_i\left(a\right) - a. \end{aligned} \tag{C.6}$$

For asymptotic local stability to hold, the Jacobian of the differential equation in (C.6) must be less than zero at the conjectured equilibrium. The derivative of $a_i(a)$ with respect to a is:

$$a_i'(a) = -\frac{2\gamma\left(1-\gamma\right)^3\left(1-\left(1-\gamma\right)a\right)\sigma^2}{\left(\left(1-\gamma\right)^2\sigma^2 + \left(1-\left(1-\gamma\right)a\right)^2\gamma^2\right)^2}, \tag{C.7}$$

which is positive whenever $a > (1-\gamma)^{-1}$. Then, necessarily, $a_i'(a_\circ) > 1$, $a_i'(a_+) \in (0,1)$, $a_i'(a_-) < 0$ and $a_i'(a_u) < 0$. i.e. the low and unique equilibrium are respectively locally and globally learnable. □

D Data definitions

Unless otherwise noted, variables are download from the FRED database maintained by the Federal Reserve Bank of St. Louis. FRED variable codes are provided in parenthesis when available. Annual variables are averaged over all observations of the given year.

For per-capita variables, our measure of population is the civilian non-institutional population over the age of 16 (CNP16OV). Our measure of real per-capita gross domestic product is seasonally-adjusted nominal quarterly GDP (GDP) deflated by the GDP price deflator (GDPDEF) and population. Our measure of real per-capita consumption is nominal personal consumption expenditure (PCEC) deflated by the GDP price deflator and population. Our measure of per-capita hours in the seasonally-adjusted hours of all persons in the non-farm business sector (HOANBS) divided by population.

We compute real per-capita residential investment using nominal private residential fixed investment (PRFI) deflated by the chain-type price index for the real private fixed investment in the residential sector (B011RG3Q086SBEA) and population. We use the same deflator (B011RG3Q086SBEA) divided by the GDP deflator price index for the real price of residential investment. For real house prices, we used the index of Shiller (downloaded from http://www.econ.yale.edu/ shiller/data.htm). Construction and aggregate productivity are taken from the World KLEMS database (downloaded from http://www.worldklems.net/data.htm).

Finally, for our measures expectations, we use data from the Survey of Consumers at the University of Michigan (downloaded from https://data.sca.isr.umich.edu/charts.php). For real income expectations we use their index generated from question 14, "Expected Change in Real Income During the Next Year." For past house price experience, we use the index generated from question 45, "Change in Home Values During the Past Year." Finally, for the measure of economic news heard, we use the index generated from question 23 "News Heard of Recent Changes in Business Conditions."