

Confidence and the Propagation of Demand Shocks^{*}

George-Marios Angeletos[†] Chen Lian[‡]

July 7, 2020

Abstract

We revisit the question of why shifts in consumer spending drive business cycles. Our theory combines intertemporal substitution in production with rational confusion (or bounded rationality) in consumption. The first element allows aggregate supply to respond to shifts in aggregate demand without nominal rigidity. The second introduces a “confidence multiplier,” namely a positive feedback loop between real economic activity, consumer perceptions of permanent income, and investor expectations of returns. This mechanism amplifies the business-cycle fluctuations triggered by demand shocks (but not those triggered by supply shocks); it helps investment to comove with consumption; and it allows front-loaded fiscal stimuli to crowd *in* private spending.

^{*}This paper subsumes earlier versions that were entitled “A (Real) Theory of the Keynesian Multiplier” and “On the Propagation of Demand Shocks” and that contained somewhat different frameworks. For helpful comments and suggestions, we thank the editor, three anonymous referees, and seminar participants at NYU, MIT, the 2018 AEA meeting and the 2018 SED meeting. Angeletos acknowledges the financial support of the National Science Foundation (Award #1757198). Lian acknowledges the financial support from Alfred P. Sloan Foundation Pre-doctoral Fellowship in Behavioral Macroeconomics, awarded through the NBER.

[†]MIT and NBER; angelet@mit.edu.

[‡]MIT; lianchen@mit.edu.

1 Introduction

In the data, fluctuations in consumer spending, such as those triggered by shocks to the net worth of consumers (Mian and Sufi, 2014) or by swings in their expectations (Beaudry and Portier, 2006), appear to cause business cycles. Yet, the theoretical mechanism through which such a drop in aggregate demand can precipitate a recession, or a fiscal stimulus can offset it, is debatable.

To fix ideas, take the textbook New Keynesian model (without capital) and consider an “urge to consume less,” i.e., a negative discount rate shock.¹ Insofar as monetary policy stabilizes inflation (which is the same as replicating flexible-price outcomes), the drop in aggregate demand does *not* generate a recession. It only triggers a perfectly-offsetting movement in the natural rate of interest. How does the model then make room for demand-driven business-cycles? By equating aggregate demand shocks to monetary contractions (relative to flexible-price outcomes) and to movements along a Philips curve (imposing a positive relation between real economic activity and inflation).

This seems unsatisfactory, not only for our tastes but also on empirical grounds. Philips curves are elusive in the data (Mavroudis, Plagborg-Møller, and Stock, 2014) and the principal component of the actual business cycle fits better the template of a non-inflationary, non-monetary demand shock (Angeletos, Collard, and Dellas, 2020). This paper therefore seeks a theory that is both “Keynesian” in the sense of letting demand shocks be the main business-cycle driver and “neoclassical” in the sense of not relying on nominal rigidity.²

The proposed theory combines two key elements. The first element, intertemporal substitution in production via variable utilization, opens the door to demand-driven fluctuations in the natural rate of output. The second element, a certain kind of rational confusion or bounded rationality in the demand side of the economy, introduces a confidence multiplier—a feedback loop between outcomes and expectations that amplifies demand shocks (but not necessarily supply shocks), helps investment co-move with consumption, and allows government spending to crowd in private spending.

A non-vertical AS curve. In the textbook New Keynesian model, there is no room for demand-driven business cycles because Aggregate Supply (AS) is vertical in the sense that the aggregate production of today’s goods is irresponsive to their *relative* price, or the real interest rate.³ Our contribution starts with the elementary observation that this unwanted property disappears once we allow for intertemporal substitution in production via variable capacity utilization.⁴

¹This is a standard proxy for shocks in consumer credit (Eggertsson and Krugman, 2012; Guerrieri and Lorenzoni, 2017; Hall, 2011; Mian and Sufi, 2014) and other shifts in aggregate private spending.

²While our paper is committed to obtaining demand-driven fluctuations *outside* the New Keynesian framework, the mechanisms discussed below are relevant *inside* it, too, for they modify the properties of the natural rate of output.

³Note that here we have in mind an AS curve in the space of real output today vs its *relative* price, namely the real interest rate. This differs from the Keynesian tradition that casts the AS curve as a Philips curve, or as a positive relation between real output and the *nominal* price level, or the nominal inflation rate.

⁴More specifically, our baseline model introduces variable utilization of *land*. This mimics the standard formulation of capital utilization (e.g., Greenwood, Hercowitz, and Huffman, 1988) but abstracts, tentatively, from investment. An

This element of our theory, which does *not* require any departure from full information and full rationality, seems empirical relevant on its own right: capacity utilization is strongly pro-cyclical in the data. But a literal interpretation of this element is not strictly needed: as long as there is *some* neoclassical margin of intertemporal substitution in production, the second part of our contribution, which regards the Aggregate Demand (AD) side of the economy, goes through.

AD and the confidence multiplier. With fully informed and fully rational consumers, the AD side of our model is standard: it boils down to the Euler condition of a representative consumer. In this familiar benchmark, consumers respond very differently to aggregate income fluctuations than to idiosyncratic ones, not only because they can perfectly tell the two apart but also because they fully comprehend the latter's distinct equilibrium implications.

We instead let consumers confuse aggregate and idiosyncratic income fluctuations when deciding how much to spend. In our main analysis, such confusion is rational, resulting from lack of information about, or inattention to, aggregate shocks.⁵ But the same basic behavioral pattern obtains if we replace the lack of information with lack of sophistication, or bounded rationality.⁶ Our main contribution is to show that this friction introduces a feedback loop of a Keynesian flavor.

Whenever a negative AD shock hits the economy, firms cut down on production and employment. Consumers experience a drop in their employment opportunities and income but can't tell whether this was because of a negative AD shock or an adverse idiosyncratic shock. In our environment, AD shocks have *no* effect on aggregate permanent income, regardless of how large or persistent they are; this follows from a dynamic translation of Hulten's theorem.⁷ By contrast, a consumer's permanent income naturally varies with idiosyncratic shocks. It follows that, after an adverse aggregate demand shock, consumers rationally mis-perceive a reduction in their permanent income.

Consumers thus lose confidence, in the sense that they become excessively pessimistic relative to the case of full-information, full-rationality. But as they do, they spend less, amplifying the original drop in aggregate demand. This feeds into a further cut in employment and output, a further drop in confidence, and so on. It is this feedback loop that we refer to as the "confidence multiplier."

extension adds investment and clarifies that the requisite margin for intertemporal substitution in production remains active as long as there is variable utilization along an adjustment cost to capital as in standard Q theory (e.g., [Abel and Blanchard, 1983](#)).

⁵This follows the traditions of [Lucas \(1973\)](#) and [Sims \(2003\)](#). See especially [Mackowiak and Wiederholt \(2009\)](#) for why agents may optimally choose to be little informed about aggregate shocks, and instead allocate most of their attention and cognitive capacities on tracking idiosyncratic conditions, in line with what we assume in our main analysis.

⁶As discussed later, such bounded rationality could take the form of simple extrapolative beliefs as in [Greenwood and Shleifer \(2014\)](#) and [Gennaioli, Ma, and Shleifer \(2016\)](#), or a low-dimensional representation of the world as in [Molavi \(2019\)](#).

⁷This is a subtle point. In the New Keynesian model, permanent income covaries with demand shocks insofar as the latter trigger procyclical gaps from the natural rate of output. But as long as monetary policy replicates flexible-price outcomes, demand shocks do *not* affect permanent income, regardless of how they affect the natural rate of output itself.

Discounting the neoclassical GE adjustment. Our confidence multiplier is active even if the AD shock lasts only one period. But when the shock lasts multiple periods, an additional mechanism may emerge: consumers may discount the future GE adjustment in the real interest rate, either because they perceive the current adjustment as “noise” or because they under-estimate how much *others* are reducing their spending.

This mechanism is akin to the forms of GE attenuation articulated in [Angeletos and Lian \(2018\)](#), [Farhi and Werning \(2019\)](#) and [Gabaix \(2020\)](#). But whereas in the New Keynesian context of those papers GE attenuation translates to *under*-reaction of the output gap to news about future monetary policy, in our neoclassical context it maps to *over*-reaction of the natural rate of output to the underlying aggregate demand shock.⁸ This mechanism thus complements our confidence multiplier in amplifying the business cycles triggered by demand shocks.

Supply shocks. Consider supply shocks of the RBC type, namely aggregate TFP shocks. Because such shocks *do* affect permanent income, their confusion with idiosyncratic shocks has an ambiguous effect on consumer confidence: consumers under-estimate the aggregate fluctuations in permanent income at the same time that they over-estimate the idiosyncratic ones. In an empirically plausible case, these effects wash out and our confidence multiplier is switched off.

Furthermore, because the GE adjustment in the real interest rate is the reverse of that triggered by demand shocks, the effect of the second mechanism is also the reverse: when consumers under-estimate the adjustment in the real interest rate, they respond more to demand shocks but less to supply shocks. All in all, supply shocks are thus dampened at the same time that demand shocks are amplified. In combination with the assumption of flexible prices, or the absence of a Philips curve, this allows our theory to match the disconnect of the business cycle from *both* inflation and TFP.⁹

Comovement in consumption and investment. To simplify the exposition, our baseline analysis allows for variable utilization of capital (or “land”) but abstracts from investment. If we add investment but let it be decided with full information and full rationality, we run on a familiar co-movement “puzzle.” Whenever consumers spend less, investment picks up some of the slack, simply because the cost of investment (the real interest rate) goes down in equilibrium. Consumption and investment thus move in opposite directions, which is at odds with the data.¹⁰

⁸This reflects a difference in the underlying form of strategic interaction, or the way higher-order beliefs matter. Whereas the kind of NK models studied in the aforementioned works correspond to games of strategic complementarity, the kind of RBC model we study here corresponds to a game of strategic substitutability. Under-estimating the responses of others leads to under-reaction in the first type of games and to over-reaction in the latter.

⁹Consistent with our theory, here we have in mind *utilization-adjusted* TFP in the data, as in [Basu, Fernald, and Kimball \(2006\)](#). See also [Blanchard and Quah \(1989\)](#), [Gali \(1999\)](#), [Barsky and Sims \(2011\)](#) and [Angeletos, Lorenzoni, and Pavan \(2010\)](#) for additional facets of the disconnect between technology and business cycles.

¹⁰This puzzle extends from the baseline RBC framework to state-of-the-art DSGE models despite their use of flat Philips curves, accommodative monetary policies, and various bells and whistles that bolster the empirical performance of the New Keynesian framework. For instance, [Justiniano, Primiceri, and Tambalotti \(2010\)](#) and [Christiano, Motto, and Rostagno \(2014\)](#)

A version of our confidence multiplier applied on investment decisions helps resolve this puzzle. Insofar as investors are subject to a similar confusion as consumers,¹¹ a temporary drop in *either* the household or the business component of aggregate demand can trigger a drop in *both* consumer confidence (expectations of income) and investor sentiment (expectations of returns). This allows consumption and investment to co-move.

All in all, our theory thus fits all the main qualitative properties of actual business cycles: positive comovement between employment, output, consumption and investment, without commensurate comovement in either TFP or inflation.

Fiscal Policy. Commentators often claim that fiscal policy can help “boost confidence” but fail to clarify what they mean by it.¹² Our theory provides a way to think about this. On the one hand, it accommodates the idea that fiscal stimuli can boost consumer confidence in the sense we have defined. On the other hand, it qualifies that this is possible only when the stimulus is front-loaded.

Such a stimulus raises the aggregate demand for today’s goods, triggering our confidence multiplier. When this multiplier is large enough, the stimulus can *crowd in* private spending, despite the higher tax burden.¹³ By contrast, because a backloaded stimulus induces the opposite intertemporal shift in production, it fails to raise consumer confidence and necessarily crowds out private spending.

Discussion and related literature. All in all, the picture painted by our theory is consistent with the Keynesian view of demand-driven business cycles. But it does not equate such fluctuations to monetary shocks, or to movements along a Philips curve. It therefore fits the template of the “main business cycle shock” provided by [Angeletos, Collard, and Dellas \(2020\)](#): the principal component of the business-cycle fluctuations in unemployment, GDP, hours, worked, consumption and investment in the data appears to be disconnected from both technology and inflation, calling for a theory of non-inflationary, non-monetary, demand-driven fluctuations.

This brings to mind the older literature on coordination failures ([Benhabib and Farmer, 1999](#); [Cooper and John, 1988](#); [Diamond, 1982](#)), as well as a few more recent works that share our motivation but our insights ([Angeletos, Collard, and Dellas, 2018](#); [Angeletos and La’O, 2013](#); [Bai, Rios-Rull, and](#)

fail to generate positive co-movement between consumption and investment in response not only to consumer discount-rate shocks but also to the kind of investment-specific shocks (or “risk shocks”) that are the main drivers of the business cycle in those models.

¹¹In our model this is simply because investors and consumers are the same entities.

¹²For instance, Robert Shiller made such a claim during the Great Recession, prompting the following response by N. Gregory Mankiw (<http://gregmankiw.blogspot.com/2009/01/>):

“Yale’s Bob Shiller argues that confidence is the key to getting the economy back on track. I think a lot of economists would agree with that. The question is what it would take to make people more confident. ... Until we figure it out, it is best to be suspicious of any policy whose benefits are supposed to work through the amorphous channel of ‘confidence.’”

¹³What is more, unlike the New Keynesian model, such crowding-in is possible without accommodative monetary policies and inflationary pressures.

Storesletten, 2017; Beaudry and Portier, 2013, 2018; Huo and Rios-Rull, 2020; Ravn, Schmitt-Grohe, and Uribe, 2006). What distinguishes our contribution are its two novel propagation mechanisms: the margin for intertemporal substitution on the supply side, and the confidence multiplier on the demand side.

To the extent that the aforementioned works make room for fluctuations in “confidence,” they do so by equating such fluctuations to exogenous shifts across multiple equilibria or other extrinsic shocks to higher-order beliefs. We instead allow confidence to vary endogenously, and in response to intrinsic shocks (such as shocks to consumer discount rates or government spending).

Closely related in this respect are [Chahrour and Gaballo \(2018\)](#) and [Ilut and Saijo \(2018\)](#). These works focus on different economics—the former focuses on housing price fluctuations, the latter on ambiguity and learning. But they share with us the broader theme of letting intrinsic shocks drive certain “wedges” in consumer or producer beliefs.

[Chahrour and Gaballo \(2018\)](#) further share with us the property that a signal-extraction problem is the source of macroeconomic complementarity. But they do not share our specific feedback mechanism, its context, and its policy implications. The same point applies to [Angeletos, Lorenzoni, and Pavan \(2010\)](#), [Benhabib, Wang, and Wen \(2015\)](#) and [Gaballo \(2017\)](#).

Our confidence multiplier originates from an informational friction à la [Lucas \(1972\)](#), [Sims \(2003\)](#), and [Mackowiak and Wiederholt \(2009\)](#). Empirical evidence in support of such a friction abounds; see in particular the evidence in [Andrade et al. \(2020\)](#) about the confusion of idiosyncratic (industry-specific) and aggregate shocks.¹⁴ But a similar feedback loop can emerge in variants in which agents are fully informed but form simple extrapolative beliefs as in [Greenwood and Shleifer \(2014\)](#) and [Gennaioli, Ma, and Shleifer \(2016\)](#), or form a sparse representation of the world as in [Molavi \(2019\)](#). This allows for a behavioral re-interpretation of our confidence multiplier, which we welcome.

A similar point applies to the discounting of the GE adjustment in the real interest rate: while our model captures this with the help of incomplete information, as in [Angeletos and Lian \(2018\)](#), it is possible to recast it with two plausible departures from rationality, cognitive discounting à la [Gabaix \(2020\)](#) and level-k thinking à la [Garcia-Schmidt and Woodford \(2019\)](#) and [Farhi and Werning \(2019\)](#). What distinguishes our contribution vis-a-vis these works is not only the change in the context and the empirical implications of the shared mechanism (recall footnote 8 and the surrounding discussion), but also our confidence multiplier, which is entirely novel.

Last but not least, our combination of informational frictions and demand shocks is reminiscent of [Lorenzoni \(2009\)](#). That paper proposes a new micro-foundation of demand shocks, in terms of

¹⁴Less direct but complementary is also the evidence in [Coibion and Gorodnichenko \(2012, 2015\)](#) and [Angeletos, Huo and Sastry](#), which points towards incomplete information about aggregate shocks, and that in [Coibion, Gorodnichenko, and Ropele \(2018\)](#) and [Cavallo, Cruces, and Perez-Truglia \(2017\)](#), which suggest inattention to and/or sub-optimal use of publicly available information about the aggregate state of the economy. See also [Gabaix \(2019\)](#) for a review of the broader empirical and theoretical literatures on inattention.

noisy expectations about future productivity and income, but maintains the propagation mechanism of the New Keynesian model, using nominal rigidity (along with appropriate monetary policy) to translate such shocks to procyclical gaps from flexible-price outcomes. Our paper instead revisits the propagation mechanism itself, dropping nominal rigidity altogether and allowing for demand-driven fluctuations in the natural rate of output itself.

2 The Model

We consider a neoclassical economy, with flexible prices. There is a single aggregate shock, which proxies for fluctuations in aggregate demand, and various idiosyncratic shocks, which play the dual role of introducing idiosyncratic income fluctuations and of allowing for an informational friction. This informational friction is the core element of the demand block of our model. The core element of our model's supply block, on the other hand, is a margin for intertemporal substitution in production.

Islands. Following the tradition of [Lucas \(1973\)](#), we represent the economy as a continuum of islands, indexed by $i \in [0, 1]$. On each island, there is a continuum of firms, each being a monopolistically competitive producer of a differentiated good. We use (i, j) to identify both the j -th variety produced on island i and the firm producing it. On each island, there is also a representative household, indexed by $h = i$. In each period $t \in \{0, \dots, +\infty\}$, the household is employed in the island the household resides in, but “visits” and purchases the goods of a (non-representative) sample of the other islands.

Household preferences. The preferences of any household h are represented by

$$\mathcal{U}^h = U(c_0^h, n_0^h) + \beta_0^h U(c_1^h, n_1^h) + \beta_0^h \beta_1^h U(c_2^h, n_2^h) + \dots, \quad (1)$$

where c_t^h and n_t^h is its consumption and labor supply at t , β_t^h is its discount factor between t and $t + 1$ (with $\beta_{-1} \equiv 1$), and U is the per-period utility function. The latter is given by

$$U(c, n) = \frac{c^{1-\frac{1}{\sigma}}}{1-\frac{1}{\sigma}} - \frac{n^{1+\frac{1}{\nu}}}{1+\frac{1}{\nu}}, \quad (2)$$

where $\sigma > 0$ and $\nu > 0$ are, respectively, the elasticity of intertemporal substitution in consumption and the Frisch elasticity of labor supply.

Consumption is given by

$$c_t^h = F\left(\left\{c_{i,t}^h, \xi_i\right\}_{i \in \mathcal{C}_t^h}\right) \quad \text{and} \quad c_{i,t}^h = H\left(\left\{c_{i,j,t}^h\right\}_{j \in [0,1]}\right), \quad (3)$$

where the functions F and H are CES aggregators, $c_{i,j,t}^h$ is the consumption of variety j from island i

in period t , $c_{i,t}^h$ is a consumption index for all the varieties consumed from island i in period t , C_t^h is the set of islands which household h “visits” (i.e., consumes the products of) in period t , and $\xi_{i,t}$ is an island-specific taste shock. The latter follows an AR(1) process with persistence $\rho_\xi \in [0, 1)$:

$$\log \xi_{i,t} = \rho_\xi \log \xi_{i,t-1} + \log \epsilon_{i,t}^\xi, \quad (4)$$

where $\log \epsilon_{i,t}^\xi \sim \mathcal{N}(0, \sigma_\xi^2)$ is i.i.d. across i and t and independent from all other shocks (and $\xi_{i,-1} = 1$).

One can think of the islands as different categories of expenditure. To simplify the exposition, we fix the elasticity of substitution across them to 1, which means a constant expenditure share per island. On the other hand, to make sure that the monopolist’s problem is well defined and the markup is finite, we let the elasticity of substitution across the different varieties of the same island be $1 + \frac{1}{\mu}$, where $\mu > 0$ identifies in equilibrium the monopoly markup. Perfect competition is nested as $\mu \rightarrow 0$.

Demand shocks. We let an aggregate discount-rate shock proxy for shifts in aggregate demand. Each household’s discount factor follows an AR(1) process with persistence $\rho_\beta \in [0, 1)$:

$$\log \beta_t^h = (1 - \rho_\beta) \log \beta + \rho_\beta \log \beta_{t-1}^h - \log \eta_t + \log \epsilon_t^{\beta,h} \quad (5)$$

where $\beta \in (0, 1)$ is the steady-state discount factor, $\log \eta_t \sim \mathcal{N}(0, \sigma_{AD}^2)$ is the aforementioned *aggregate* shock, i.i.d. across t and independent of all other shocks in the economy, and $\log \epsilon_t^{\beta,h} \sim \mathcal{N}(0, \sigma_\beta^2)$ is an *idiosyncratic* shock, i.i.d. across both h and t and independent of all other shocks.¹⁵ The latter’s sole modeling role is to make sure that the household does not know the aggregate shock even though it knows its own discount factor. Finally, note that η_t enters (5) with a minus, so that a positive realization for η_t corresponds to an urge to consume more, or a positive AD shock.

Inter-temporal trading. Let R_t denote the *aggregate*, real, gross, interest rate between t and $t+1$. The corresponding rate faced by the households in island i is given $R_{i,t} = R_t \epsilon_{i,t}^R$, where $\log \epsilon_{i,t}^R \sim \mathcal{N}(0, \sigma_R^2)$ is an island-specific shock, i.i.d. across time and islands, and independent from all other shocks. This island-specific, interest-rate shock can be interpreted as the product of a random intermediation cost.¹⁶ Its sole modeling role is to limit the information households can extract from their personal borrowing and lending conditions, or from the local interest rate.¹⁷

¹⁵Throughout, we denote idiosyncratic innovations by $\epsilon_t^{x,h}$ and $\epsilon_{i,t}^x$ for, respectively, households and islands, where x stands for the various random variables.

¹⁶To see this, let borrowing and lending takes place via a two-tier market. On the bottom, there is a local market for one-period bonds, or IOUs, for each island, where households trade with financial intermediaries. On the top, there is a centralized market where only intermediaries trade. Intermediaries act competitively in both markets, but intermediation is costly: there is an “iceberg cost” for clearing claims between the local and the centralized market.

¹⁷Rational inattention (Sims, 2003; Mackowiak and Wiederholt, 2009) and related forms of bounded rationality can play a similar modeling role. See Vives and Yang (2018) for an example where rational inattention limits the information extracted from prices in an asset-market context in the tradition of Grossman and Stiglitz (1980).

Household budget. The household's total income in period t is $w_{h,t}n_t^h + e_{h,t}$, where $w_{h,t}$ is the real wage on the island $i = h$ the household lives in, and $e_{h,t}$ is the sum of the local profits, which the household receives as dividends. Its budget constraint can therefore be written as follows:

$$\int_{i \in \mathcal{C}_t^h} \int_{j \in [0,1]} p_{i,j,t} c_{i,j,t}^h dj di + R_{h,t}^{-1} b_{t+1}^h = w_{h,t} n_t^h + e_{h,t} + b_t^h \quad (6)$$

where b_t^h is the net amount of bonds held at the start of period t , with $b_0^h = 0$.

Firms and production. Consider firm (i, j) on island i . We assume that production is given by a Cobb-Douglas form between labor and land services, so that the firm's output in period t is given by

$$q_{i,j,t} = (l_{i,j,t})^\alpha (u_{i,j,t} k_{i,j,t})^{1-\alpha}, \quad (7)$$

where $l_{i,j,t}$ is the labor input and $u_{i,j,t} k_{i,j,t}$ are land services. The latter are given by the product of effective land quantity, $k_{i,j,t}$, and utilization rate, $u_{i,j,t}$.

The firm's operating profit in period t is given by $\pi_{i,j,t} = p_{i,j,t} q_{i,j,t} - w_{i,t} l_{i,j,t}$, where $p_{i,j,t}$ is the price of good (i, j) and $w_{i,t}$ is the wage of the local labor market in island i . The firm owns its own land, and is itself owned by the local household. It follows that the firm's objective is the local valuation of the stream of its profits:

$$\mathbb{E}_t \left[\sum_{s=0}^{\infty} m_{i,t,t+s} \pi_{i,j,t+s} \right] \quad (8)$$

where $m_{i,t,t+s} \equiv E_t^i [(\prod_{l=0}^s \beta_{l-1}^h) U_c(c_{t+s}^i, n_{t+s}^i)]$ is the local pricing kernel, or the marginal utility of consumption for household $h = i$.

Each firm is a monopolist in the market for the particular commodity she produces. But since there are many firms in each island competing for the local labor, each firm is a price taker in the labor market of the island it operates in (and so is the island's representative household). Finally, all the firms of any given island is owned by the local household and distribute their profits to it as dividends. It follows that the dividend income received by household h (also the total dividends in island $i = h$) in period t is given by $e_{h,t} = \int \pi_{h,j,t} dj$.

Land and utilization. The effective quantity of land obeys the following law of motion:

$$k_{i,j,t+1} = (1 - \delta(u_{i,j,t})) k_{i,j,t}, \quad (9)$$

where $\delta(\cdot)$ is an increasing and convex function and $k_{i,j,0} = \bar{k}_0 > 0$. The term $1 - \delta(u_{i,j,t})$ captures the negative effect that higher land utilization today has on land productivity tomorrow. This resembles models where the depreciation of capital increases with its utilization, except that there is no

investment in land.¹⁸ Accordingly, we require that $\delta(u) = 0$ when u takes its steady-state value.¹⁹

Matching. Each household is randomly matched to, and visits, a non-representative sample of the islands in the first period. This guarantees that the household would not learn the aggregate shock from observing the prices of the islands it visits *even if* these prices were to co-move with that shock.²⁰ At the same time, we assume that the set of the consumers visiting any given island i is a representative sample of all the consumers. Relaxing this assumption would only introduce a second local demand shock, which would play the same role as $\xi_{i,t}$ and would not alter any of the results.

Let $p_{i,t}$ be the ideal price index of the goods produced on island i in period t , p_t^h the ideal price index of the basket of goods consumed by household h in period t , $y_{i,t}$ the total spending for the goods produced on island i in period t , p_t and y_t the corresponding aggregates, and c_t real aggregate spending.²¹ By our choice of numeraire, $p_t = 1$ and $y_t = c_t$. By the assumption that each household visits only a non-representative, random sample of islands,

$$p_t^h = p_t \exp(\epsilon_t^{p,h}),$$

where $\log \epsilon_t^{p,h} \sim \mathcal{N}(0, \sigma_p^2)$ is i.i.d. across h and t , independent from the rest of the economy, and captures the idiosyncratic variations of the price of the basket of goods consumed by household h . Finally, by the assumption that each island receives a representative sample of the consumers,

$$y_{i,t} = \xi_{i,t} y_t = \xi_{i,t} c_t. \quad (10)$$

Information. We assume away any informational friction among the firms so as to isolate the role of the informational friction on consumers, or the demand side of the economy.²² We then specify the information set of household h in each period t as follows:

$$\mathcal{I}_t^h = \mathcal{I}_{t-1}^h \cup \left\{ \beta_t^h \right\} \cup \left\{ w_{h,t}, e_{h,t}, R_{h,t}, (p_{i,l,t})_{i \in \mathcal{C}_t^h, l \in [0,1]} \right\} \cup \left\{ \epsilon_{t-1}^\beta \right\}. \quad (11)$$

¹⁸We will consider an extension with investment, and capital instead of land, in Section 5.3.

¹⁹Without this restriction, the economy features a balanced-growth path, with positive growth if $\delta(u) < 0$ and negative growth if $\delta(u) > 0$. Although these possibilities make little sense under our interpretation of k as land, our results are robust to them. Furthermore, this "nuisance" disappears once we add investment (Section 5.3). The economy then admits a unique and globally stable steady state, as in the standard RBC model.

²⁰This possibility does *not* occur in our main analysis because we focus on equilibria in which the aggregate price level is fixed at one. Accordingly, we can even drop entirely the random matching and let households observe *all* the commodity prices in the economy. However, the random matching guarantees that our main lessons survive in situations in which monetary policy fails to stabilize the price level. See Section XX for further discussion of this issue.

²¹More specifically, we define $p_{i,t} \equiv \left(\int_{j \in [0,1]} p_{i,j,t}^{1-\epsilon} dj \right)^{1/(1-\epsilon)}$, $p_t^h \equiv \left(\int_{i \in \mathcal{C}_t^h} \xi_i di \right) \exp \left(\int_{i \in \mathcal{C}_t^h} \xi_i \ln(p_i/\xi_i) di / \int_{i \in \mathcal{C}_t^h} \xi_i di \right)$, $y_{i,t} \equiv \int_{\{h: i \in \mathcal{C}_t^h\}} \int_{j \in [0,1]} p_{i,j,t} c_{i,j,t}^h dj dh$, $P_t \equiv \int_{h \in [0,1]} \frac{c_t^h}{c_t} p_t^h dh$, $y_t \equiv \int_{i \in [0,1]} y_{i,t} di$, and $c_t \equiv \int c_t^h$.

²²See, however, Section 5.3 for an extension that lets the friction impact not only consumption but also investment.

That is, the household learns its current discount factor, all the objects that enter its current budget (its two sources of income, its interest rate, and the prices of the goods it purchases), and the past aggregate shock.

These properties, the entire structure of the economy, and the fact that agents have rational expectations are common knowledge. It follows that, in each period, households have common knowledge of the *past* aggregate outcomes.²³

This buys us a lot of tractability, as there is no need to keep track of the dynamics of the hierarchy of beliefs. But it does not drive our results. In fact, it only weakens the ferocity of the mechanisms identified in this paper, because it forces learning to be very fast.

What is essential is only that we have removed common knowledge of the concurrent aggregate demand shock and have allowed this shock to be confounded with idiosyncratic income fluctuations. These properties and their implications will become clear in the subsequent analysis.

Equilibrium . The solution concept is (Noisy) Rational Expectations Equilibrium. Because of the CES specification and the symmetry of the firms within each island i , the demand faced by each firm (i, j) in period t can always be written as

$$\int_{\{h:i \in \mathcal{C}_t^h\}} c_{i,j,t}^h dh = y_{i,t} (p_{i,j,t}/p_{i,t})^{-\epsilon} \quad \forall i, j, t. \quad (12)$$

We can then define an equilibrium as follows:

1. Each household h chooses contingent plans for consumption and labor supply so as to maximize its expected utility subject to the budget constraint given in (6), the informational constraints embedded in (11).
2. Each firm (i, j) chooses contingent plans for inputs, production and prices so as to maximize its expectation of (8) subject to (9) and (12).
3. The goods and labor markets clear:²⁴ for every island i , every period t , and every realization of uncertainty,

$$\int_{\{h:i \in \mathcal{C}_t^h\}} c_{i,j,t}^h dh = q_{i,j,t} \quad \forall j \quad \text{and} \quad n_t^h = \int l_{i,j,t} dj. \quad (13)$$

Steady state, log-linearization and notation. To keep the analysis tractable, we henceforth work with the log-linearization of the model around its steady state.²⁵ With abuse of notation, we henceforth

²³Because aggregate outcomes are measurable in the aggregate shocks, common knowledge of the history of the exogenous shocks, which we have assumed via (11), implies along any rational-expectations equilibrium common knowledge of the history of the endogenous aggregate outcomes.

²⁴We do not explicitly require that the bond market clear, because this follows from Walras' law once we clear that markets for good and labor.

²⁵For the existence and characterization of the steady state, see Appendix XX.

reinterpret all the variables as the log-deviations from their steady-state counterparts. Finally, because firms are symmetric within islands, we can drop the firm index from all the relevant variables: $l_{i,j,t} = l_{i,t}$, $u_{i,j,t} = u_{i,t}$, $q_{i,j,t} = q_{i,t}$ and $p_{i,j,t} = p_{i,t}$ for all firms j within any given island i .

3 Aggregate Supply, Aggregate Demand, and Beliefs

In this section we derive the equilibrium conditions of the economy. We organize them in two blocks, one representing aggregate supply (AS) and another representing aggregate demand (AD). We further show how consumer beliefs are determined and how they enter the demand block.²⁶ We thus set the stage for the next section, which completes the equilibrium characterization by studying the fixed point between aggregate supply, aggregate demand, and beliefs.

3.1 Aggregate Supply

We characterize the supply side in two steps. First, we momentarily take the utilization decisions as given and obtain aggregate employment and aggregate output as functions of the aggregate utilization and the aggregate stock of land. Second, we work out the optimal utilization decisions.

By the Cobb-Douglas specification of the production function, local output is

$$q_{i,t} = (1 - \alpha)(u_{i,t} + k_{i,t}) + \alpha l_{i,t}.$$

By the FOC of the firm's problem with respect to its labor input, the local demand for labor is

$$l_{i,t} = p_{i,t} + q_{i,t} - w_{i,t}. \quad (14)$$

By the corresponding FOC for household $h = i$, the local supply of labor is

$$n_t^i = \nu (w_{i,t} - p_t^i) - \frac{\nu}{\sigma} c_t^i, \quad (15)$$

By market clearing in the local labor market,

$$l_{i,t} = n_t^i.$$

Finally, by (16) and the fact that $c_t = y_t$ in equilibrium, local demand, or local firm revenue, can be expressed as follows:

$$p_{i,t} + q_{i,t} = y_{i,t} = y_t + \xi_{i,t}, \quad (16)$$

²⁶Keep in mind that, although the supply side of our model is forward-looking, the informational friction enters the equilibrium only via the demand side, because it is consumers, not firms, who are imperfectly informed.

Aggregating all these conditions, and combining them, we can solve for aggregate output/spending as functions of u_t and k_t , the aggregate utilization rate and the aggregate land stock. In particular,

$$y_t = (1 - \tilde{\alpha})(u_t + k_t) \quad (17)$$

where $\tilde{\alpha} \equiv 1 - \frac{(1-\alpha)(1+\frac{1}{\nu})}{1+\frac{1}{\nu}-\alpha+\frac{\alpha}{\sigma}} \in (0, 1)$.

Consider now a firm's optimal choice of utilization. Because more utilization today increases output today at the expense of degrading the effective stock of land, the optimal utilization choice is naturally forward-looking. Let $\vartheta_{i,t}$ denote the shadow value of land at the start of period t . The typical firm's FOC for utilization can then be expressed as follows:

$$p_{i,t} + q_{i,t} - u_{i,t} = -R_{i,t} + \mathbb{E}_t[\vartheta_{i,t+1}] + \phi u_{i,t} + k_{i,t}, \quad (18)$$

where $\phi > 0$ is a scalar that parameterizes the elasticity of utilization²⁷ and $E_t[\cdot]$ is the full-information rational expectation in period t . The left-hand side measures the benefit of more utilization in terms of higher current production. The right-hand side measures the cost in terms of future productivity losses, appropriately discounted. The evolution of the shadow value of land is finally given by the following asset-pricing-like condition:²⁸

$$\vartheta_{i,t} = (1 - \beta)(p_{i,t} + q_{i,t} - k_{i,t}) + \beta E_t[\vartheta_{i,t+1} - R_{i,t} - \kappa u_{i,t}], \quad (19)$$

Combining and aggregating (18)-(19), replacing aggregate output from (17), using the law of motion for land, and solving for u_t , we reach the following Euler condition for aggregate utilization:

$$u_t = \frac{\beta}{\tilde{\alpha} + \beta\phi} R_t + \beta \mathbb{E}_t[u_{t+1}]. \quad (20)$$

This epitomizes the core supply-side element of our analysis: other things equal, an increase in the relative price of today's goods (i.e., the real interest rate) stimulates aggregate utilization, thus also increasing the aggregate supply of today's goods. As anticipated, the elasticity of this margin is inversely related to ϕ , because a higher ϕ means a more convex cost of utilization (a more convex δ).

The supply block is completed with the law of motion for the effective stock of land, which can be expressed in log-linearized form as $k_{t+1} = k_t - \kappa u_t$, for the same scalar κ as that appearing in (19). Putting everything together, we reach the following result.

²⁷This scalar is defined by $\phi \equiv \frac{\delta''(u^*)u^*}{\delta'(u^*)}$, where u^* denotes the steady-state level of utilization. A higher ϕ therefore means a more convex cost for utilization, which translates in equilibrium to a more inelastic utilization margin.

²⁸The scalar κ appearing in this condition is given by $\kappa \equiv \frac{\delta'(u^*)u^*}{1-\delta(u^*)}$, where u^* denotes the steady-state level of utilization. As explained in the Appendix, κ ends up coinciding with $\frac{1-\beta}{\beta}$ in the steady state.

Proposition 1 (Aggregate Supply). *The optimal behavior of the firms, the optimal labor supply of the households, and market clearing impose the following equilibrium restrictions for all t :*

$$y_t = (1 - \tilde{\alpha})(u_t + k_t), \quad (21)$$

$$u_t = \frac{\beta}{\tilde{\alpha} + \beta\phi} R_t + \beta \mathbb{E}_t [u_{t+1}], \quad (22)$$

$$k_{t+1} = k_t - \kappa u_t, \quad (23)$$

along with k_0 exogenously fixed and $\lim_{t \rightarrow \infty} \beta^t u_t = 0$.

Because the above result does *not* use optimality of consumer spending, it can be interpreted as a description of the “AS curve” of our model. To grasp how this works, consider a temporary increase in the real interest rate: $R_t > 0$ for $t = 0, \dots, T$ and $R_t = 0$ for $t > T$. Condition 22 implies that utilization is also temporarily increased: $u_t > 0$ for $t = 0, \dots, T$ and $u_t = 0$ for $t > T$. Why? because a higher-than-normal R_t means a higher-than-normal relative price for goods today, and hence a higher value for “transferring” land services from tomorrow to today via a higher rate of utilization.

This is the crunch of why aggregate supply is “upward sloping” in our model. But note that here we have expressed aggregate supply as an increasing function of the real interest rate, or the *relative* price of goods today vis-a-vis goods tomorrow. This contrasts the Keynesian tradition, which expresses aggregate supply as an increasing function of either the *nominal* price level or the rate inflation.

Importantly, this is not just about graphs. In the New Keynesian model, it is possible—although unusual—to express aggregate supply as a positive function of the real interest rate by appropriately combining the Philips curve with the Taylor rule for monetary policy and the Fisher equation. Still, this possibility hinges exclusively on the presence of nominal rigidity and a monetary policy that fails to replicate flexible prices: in the flexible-price core of that model, AS is invariant to the real interest rate. By contrast, the reason that AS responds to the real interest rate in our model has nothing to do with either nominal rigidity or monetary policy: it is a feature of the “natural” level of output and it reflects the accommodation of intertemporal substitution in production via variable utilization.

3.2 Aggregate Demand

We now turn to optimal consumption and aggregate demand. Consider household h . As standard, the Euler equation and the relevant transversality condition are necessary and sufficient for optimality. Combining these conditions with the household’s budget constraint, we arrive at the following characterization of the household’s (log-linearized) consumption function.

Lemma 1. For every h and t , household h 's optimal consumption in period t is given by:

$$c_t^h = (1 - \beta)b_t^h - \beta\sigma \left\{ \sum_{k=0}^{+\infty} \beta^k E_t^h [R_{h,t+k} + \beta_{t+k}^h] \right\} + (1 - \beta) \left\{ \sum_{k=0}^{+\infty} \beta^k E_t^h [y_{h,t+k}] \right\} \quad (24)$$

where $E_t^h [\cdot] = E_t [\cdot | \mathcal{I}_t^h]$ is the household's expectation at period t .

The first term captures initial wealth. The second term combines two intertemporal-substitution effects: that for interest rates and that for intertemporal preference shocks. The last terms captures permanent income. In this last part, we have used the fact the household's income coincides with $y_{h,t} = p_{h,t} + q_{h,t} = y_t + \xi_{h,t}$, or the total revenue of the firms operating in the island where the household lives, works, and receives profits from.

To obtain a representation for aggregate demand, we aggregate (24) and use the facts that average bond holdings are zero, that each household knows its current interest rate, that the future idiosyncratic shocks to interest rates are unpredictable, and that local revenue equals aggregate income plus the local taste shock. This yields the following expression for aggregate demand:

$$c_t = -\beta(R_t + \beta_t) - \beta\sigma \left\{ \int \sum_{k=1}^{+\infty} \beta^k E_t^h [R_{t+k}] dh + \int \sum_{k=1}^{+\infty} \beta^k E_t^h [\beta_{t+k}^h] dh \right\} \\ + (1 - \beta) \left\{ \int \sum_{k=0}^{+\infty} \beta^k E_t^h [y_{h,t+k}] dh \right\}. \quad (25)$$

We next manipulate this condition as follows. We rewrite the first integral in it as

$$\int \sum_{k=1}^{+\infty} \beta^k E_t^h [R_{t+k}] dh = \sum_{k=1}^{+\infty} \beta^k \mathbb{E}_t [R_{t+k}] - \frac{1}{\sigma} \mathcal{G}_t \quad (26)$$

where

$$\mathcal{G}_t \equiv -\sigma \int \sum_{k=1}^{+\infty} \beta^k \int \left\{ E_t^h [R_{t+k}] - \mathbb{E}_t [R_{t+k}] \right\} dh \quad (27)$$

captures the average misperception of the future adjustment in the real interest rate. Similarly, we rewrite the second integral as

$$\int \sum_{k=1}^{+\infty} \beta^k E_t^h [\beta_{t+k}^h] dh = \int \left\{ \frac{\beta\rho_\beta}{1 - \beta\rho_\beta} \beta_t \right\} dh = \frac{\beta\rho_\beta}{1 - \beta\rho_\beta} \beta_t = \sum_{k=1}^{+\infty} \beta^k \mathbb{E}_t [\beta_{t+k}], \quad (28)$$

and the third integral as

$$\int \sum_{k=0}^{+\infty} \beta^k E_t^h [y_{h,t+k}] dh = \sum_{k=0}^{+\infty} \beta^k \mathbb{E}_t [y_{t+k}] + \frac{\beta}{1 - \beta} \mathcal{B}_t, \quad (29)$$

where

$$\mathcal{B}_t \equiv \frac{1-\beta}{\beta} \sum_{k=0}^{+\infty} \beta^k \int \left\{ E_t^h [y_{h,t+k}] - \mathbb{E}_t [y_{t+k}] \right\} dh, \quad (30)$$

captures average misperception of permanent income. Combining, and replacing c_t by y_t , we obtain the following expression for aggregate demand:

$$y_t = -\beta(R_t + \beta_t) - \beta\sigma \sum_{k=1}^{+\infty} \beta^k \mathbb{E}_t [R_{t+k}] + \sum_{k=0}^{+\infty} \beta^k \mathbb{E}_t [y_{t+k}] + \beta\mathcal{B}_t + \beta\mathcal{G}_t$$

Finally, rewriting this condition in recursive form, we reach the following result.

Proposition 2 (Aggregate Demand). *Aggregate spending satisfies the following Euler condition:*

$$y_t = -\sigma \{R_t + \beta_t\} + \mathbb{E}_t [y_{t+1}] + (\mathcal{B}_t + \mathcal{G}_t), \quad (31)$$

where \mathcal{B}_t and \mathcal{G}_t are defined in conditions (27) and (30).

This is the “AD curve” of our model. It is the same as the Euler condition of a textbook, representative-agent economy, except for the inclusion of the terms \mathcal{B}_t and \mathcal{G}_t . When households share the same information, $E_t^h [\cdot] = \mathbb{E}_t [\cdot]$ for all h and these terms vanish. *Both* terms thus originate from the informational friction. But each one captures a *different* mechanism.

3.3 The Two Mechanisms

Consider first \mathcal{B}_t . This term captures the average error in households’ beliefs of their permanent income. When $\mathcal{B}_t > 0$, households on average over-estimate their permanent income, and this contributes, other things equal, to higher spending today. (The converse is of course true when $\mathcal{B}_t < 0$.)

To understand under what conditions \mathcal{B}_t is positive or negative, let us first let us first study $\sum_{k=0}^{+\infty} \beta^k \mathbb{E}_t [y_{t+k}]$, the *true* aggregate permanent income. In the RBC framework and our economy alike, this object is invariant to preference shocks (up to the first order).

Proposition 3. *The discounted present value of aggregate income satisfies*

$$\sum_{k=0}^{+\infty} \beta^k y_{t+k} = \frac{1-\tilde{\alpha}}{1-\beta} k_t \quad (32)$$

and is invariant to η_t , the aggregate demand shock.

As anticipated in the Introduction, this is a corollary of an intertemporal version of Hulten’s theorem. The present discounted value of aggregate income in our dynamic economy is the equivalent of aggregate GDP in a static economy. In the latter context, Hulten’s theorem, which only requires

production efficiency, implies that aggregate GDP is invariant to taste shocks. The translation of this result to our context, where production efficiency is preserved even though the demand side is ridden with an informational friction, gives Proposition 3 above.

This result does not mean that demand shocks have no real effects: as we will see, they *do* drive fluctuations in current output by shifting the *intertemporal* pattern in the equilibrium use of the available resources (land and labor) in the economy. But because the total resources themselves are fixed and there are no “wedges” in production (unlike what happens in the New Keynesian model), any demand-driven movements in aggregate output today are necessarily offset by opposite movements in aggregate output in the future.

Using the above result, we can reduce \mathcal{B}_t to

$$\mathcal{B}_t = \frac{1-\beta}{\beta} \sum_{k=0}^{+\infty} \beta^k \int E_t^h [\xi_{h,t+k}] dh.$$

That is, the misperceptions of *total* permanent income coincide with the misperceptions of the *idiosyncratic* permanent income. Furthermore, because $E_t^h[\xi_{h,t+k}] = \rho_\xi^k E_t^h[\xi_{h,t}]$ and $E_t^h[\xi_{h,t}] = y_{h,t} - E_t^h[y_t]$, we have

$$\mathcal{B}_t = \frac{1-\beta}{\beta(1-\beta\rho_\xi)} (y_t - \bar{E}_t[y_t])$$

where $\bar{E}_t[\cdot] \equiv \int E_t^h[\cdot] dh$. That is, whenever consumers under-estimate the increase in current aggregate income, they also mis-perceive an increase in the permanent income.

Consider next \mathcal{G}_t . This term captures the average error in beliefs of future interest rates. When $\mathcal{G}_t > 0$, households on average expect the aggregate discount-factor to be higher than theirs, or equivalently the aggregate interest rate to be lower than theirs, and this, too, contributes towards higher spending today. (The converse is of course true when $\mathcal{G}_t < 0$.)

The characterization of \mathcal{G}_t is a bit more complicated, because the true present discounted value of the real interest rate, unlike that of income, does vary with the demand shock. However, thank to the simple structure of the economy, it can be shown that the equilibrium value of R_t is proportional to β_t . It follows that the misperceptions of the future interest rates are themselves proportional to the misperceptions of the current demand shock. In particular,

$$\mathcal{G}_t = -\frac{\sigma^2}{\sigma+\varsigma} \frac{\beta\rho_\beta}{1-\beta\rho_\beta} (\beta_t - \bar{E}_t[\beta_t]).$$

We summarize the above properties in the following result.

Lemma 2. *With incomplete information,*

$$\mathcal{B}_t = \frac{1-\beta}{\beta(1-\beta\rho_\xi)} (y_t - \bar{E}_t[y_t]) \quad \text{and} \quad \mathcal{G}_t = -\frac{\sigma^2}{\sigma+\varsigma} \frac{\beta\rho_\beta}{1-\beta\rho_\beta} (\beta_t - \bar{E}_t[\beta_t]),$$

where $\bar{E}_t[\cdot] \equiv \int E_t^h[\cdot] dh$ is the average household expectation.

Consider now how \mathcal{B}_t behaves in response to a positive demand shock (a negative realization for β_t). Because the AS is positively sloped, the demand shock is expansionary, that is, y_t goes up. This is true even with complete information. But as long as information is incomplete, expectations under-react, that is, $\bar{E}_t[y_t]$ increases less than y_t . It follows that \mathcal{B}_t is positive, which contributes to a further increase in y_t . But as y_t increases because of this reason, \mathcal{B}_t further increases, and so on. This is the crux of the mechanism we refer to as “confidence multiplier.”

Note that this effect hinges on $\rho_\xi > 0$, meaning that idiosyncratic income shocks are persistent. If instead $\rho_\xi = 0$, today’s income experiences do not feed into expectations of future income, and $\mathcal{B}_t = 0$. But as long as $\rho_\xi > 0$, the effect is present regardless of whether the demand shock itself is transitory ($\rho_\beta = 0$) or persistent ($\rho_\beta > 0$).

By contrast, ρ_β is a critical determinant of \mathcal{G}_t . When demand shock is transitory ($\rho_\beta = 0$), agents make no mistake in predicting the future path of either the aggregate interest rate or the aggregate level of income, guaranteeing that $\mathcal{G}_t = 0$. When instead the demand shock is persistent ($\rho_\beta > 0$), a positive innovation today causes agents to underestimate the future adjustment in the aggregate real interest rate and the corresponding adjustment in aggregate income. This is the crux of the mechanism we refer to as “discounting the GE adjustment in the real interest rate.”

In the next section, we fill in the details and formally establish how both of these mechanisms help amplify the response of the economy to demand shocks. This involves the solution of not only the fixed point between AS and AD, but also of the inference problem of the household.

4 Equilibrium Characterization

In this section we complete the characterization of the equilibrium. We proceed in four steps. First, we recast AS and AD in a more convenient form. Second, we analyze the inference problem of the household. Third, we isolate the role of the “confidence multiplier” by letting, momentarily, the demand shock to be transitory. Finally, we study the complementary role of the “dampening GE” that emerges when, and only when, the demand shock is persistent.

4.1 Normalized AS and AD

The complete equilibrium is given by the fixed point between equations (21) - (23), which characterize AS, and equation (31), which characterizes AD. To analyze this fixed point, it is useful to drop the backward-looking variable k_t from the first set of equations and re-express the equilibrium as a purely

forward-looking system. To this goal, we introduce the following transformation:

$$\tilde{y}_t \equiv \frac{1}{\beta} (y_t - (1 - \tilde{\alpha})k_t) = \frac{1 - \tilde{\alpha}}{\beta} u_t.$$

The first part defines \tilde{y}_t as aggregate output appropriately “normalized.” The second part uses the production function and the equilibrium determination of employment to establish that \tilde{y}_t is proportional to utilization, which itself is a purely forward-looking variable. We can then combine Propositions 2 and 1 to obtain the following representation of the equilibrium in our economy.

Proposition 4 (AS and AD). *A path for \tilde{y}_t and R_t is part of an equilibrium if and only if it satisfies the terminal conditions $\lim_{t \rightarrow \infty} \beta^t \tilde{y}_t = \lim_{t \rightarrow \infty} \beta^t \tilde{R}_t = 0$ along with the following two dynamic equations:*

$$\tilde{y}_t = \varsigma R_t + \beta \mathbb{E}_t [\tilde{y}_{t+1}] \quad (33)$$

$$\tilde{y}_t = -\sigma R_t + \beta \mathbb{E}_t [\tilde{y}_{t+1}] + \mathcal{B}_t + \mathcal{G}_t - \sigma \beta_t, \quad (34)$$

with $\varsigma \equiv \frac{1 - \tilde{\alpha}}{\tilde{\alpha} + \beta \phi}$ and with \mathcal{B}_t and \mathcal{G}_t given by Lemma 2.

Conditions (33) and (34) are the “normalized” versions of, respectively, aggregate supply (AS) and aggregate demand (AD). Both equations are forward-looking. In the case of AS, it is because utilization is forward-looking. In the case of AD, it is because consumption is forward-looking. Their slopes with respect to the real interest rate are, respectively, ς and $-\sigma$. Finally, the terms \mathcal{B}_t and \mathcal{G}_t , which are present in the AD equation and encapsulate the two mechanisms of interest, have now been restated in terms of the transformed variable \tilde{y}_t rather than the original variable y_t .

The system provided above is easier to work with than the original given earlier in Propositions 2 and 1 because the new system is *purely* forward looking whereas the original contains a backward-looking variable.²⁹ But one has to keep in mind that any solution for \tilde{y}_t must be transformed back to a solution for y_t . This implies that, even though \tilde{y}_t is purely forward looking, y_t is both forward and backward looking. Intuitively, any movement in \tilde{y}_t maps to a moment in the concurrent level of utilization, which does not affect \tilde{y}_{t+1} but *does* affect y_{t+1} , the actual future income, via its effect on the effective stock of land.

More succinctly, an increase in \tilde{y}_t maps to an increase in current aggregate income but also a decrease in future aggregate income. This underscores that an increase in \tilde{y}_t represents an intertemporal shift in resources—which is exactly what a “shift in aggregate demand” means in a neoclassical context.

²⁹The relevant backward-looking state variable, the effective stock of land, has dropped out of the new system thanks to the adopted transformation of variables. This transformation also explains why $\mathbb{E}_t [\tilde{y}_{t+1}]$ is discounted by β in the normalized AD curve (34), whereas there is no such discounting in the original AD curve (31).

4.2 Beliefs: a Simple Representation

We now put more structure on the expectation terms that show up in aggregate demand. Recall that, in any given period t , all past shocks and all past outcomes are common knowledge and the only concurrent exogenous shock is η_t . It follows that, conditional on the commonly known past, the innovations in the average expectations of all current and future aggregate outcomes, as well in the aggregate outcomes themselves, are all proportional to η_t . That is, for any aggregate variable x in the set $\{\beta, y, \tilde{y}, R\}$ and any $s \geq 0$,

$$\mathbb{E}_t[x_{t+s}] = \mathbb{E}_{t-1}[x_{t+s}] + \gamma_s^x \eta_t, \quad (35)$$

for some scalar γ_s^x .³⁰ Note that this condition uses the full-information expectation operator, not the expectation of the typical household in the economy. But because the information in E_t is a *superset* of the information in E_t^h , the expectation of the typical household, and because the latter itself contains the information in E_{t-1} , the following is also true:

$$E_t^h[x_{t+s}] = E_t^h[\mathbb{E}_t[x_{t+s}]] = \mathbb{E}_{t-1}[x_{t+s}] + \gamma_s^x E_t^h[\eta_t]$$

Aggregating the above across h and letting λ denote the coefficient of the projection of $\bar{E}_t[\eta_t]$ on η_t , we get

$$\bar{E}_t[x_{t+s}] = \mathbb{E}_{t-1}[x_{t+s}] + \gamma_s^x \bar{E}_t[\eta_t] = \mathbb{E}_{t-1}[x_{t+s}] + \gamma_s^x \lambda \eta_t \quad (36)$$

Combining (35) and (36), we finally reach the following result.

Proposition 5. *In any equilibrium, there exist a scalar $\lambda \in (0, 1)$ such that, for any aggregate variable x in the set $\{\beta, y, \tilde{y}, R\}$,*

$$\bar{E}_t[x_{t+s}] = (1 - \lambda) \mathbb{E}_{t-1}[x_{t+s}] + \lambda \mathbb{E}_t[x_{t+s}] \quad \forall t, s \geq 0$$

That is, it is as if the economy is populated by two representative agents: one that only knows the past, and another that also knows the present, with respective weights $1 - \lambda$ and λ . By the same token, λ is an inverse measure of the severity of the informational friction: the smaller λ is, the further away the economy is from the complete-information benchmark.

The scalar λ itself is an equilibrium outcome: it is determined from the fixed point between the aggregate outcomes, which drive the market signals observed by the typical household, and the latter's inference problem, which feeds into actual behavior. But the endogeneity of λ is not essential for understanding the response of the economy to the demand shock. We thus postpone the solution

³⁰The vector $\{\gamma_s^x\}_{s \geq 0}$ identifies the Impulse Response Function of variable x with respect to the demand shock: γ_s^x is the impact of today's innovation ϵ_t^β on variable x after s lags. Clearly, this object is endogenously determined in equilibrium. For the present argument, however, this endogeneity is irrelevant. All that matters is the mere existence of such a scalar, and this follows directly from the linearity of economy and the Gaussianity of the shock.

of the aforementioned fixed point for later and focus first on completing the characterization of the equilibrium dynamics for given λ .

4.3 Isolating the Confidence Multiplier

We first consider the case $\rho_\beta = 0$, meaning an aggregate demand shock lasting only one period. This guarantees that $\mathbb{E}_t[\tilde{y}_{t+1}] = \mathbb{E}_{t-1}[\tilde{y}_{t+1}] = 0$ and $\tilde{\mathcal{G}}_t = 0$, and therefore that the AS and AD equations take the following, simpler forms:

$$\begin{aligned}\tilde{y}_t &= \varsigma R_t \\ \tilde{y}_t &= -\sigma R_t + \underbrace{(1 - \lambda) \left(\frac{1 - \beta}{1 - \beta \rho_\xi} \right)}_{\mathcal{B}_t} \tilde{y}_t - \sigma \beta_t.\end{aligned}$$

This system, which is essentially statics, allows a straightforward characterization of how the economy responds a demand shock. As mentioned before, one only needs to keep in mind that any “static” solution for \tilde{y}_t must be transformed back to a “dynamic” solution for y_t .

Consider a negative demand shock (a positive realization for η_t) and start with the complete-information benchmark (herein nested as $\lambda = 1$). In this case, $\mathcal{B}_t = 0$ and the AD curve shifts down by an amount equal to $-\sigma\eta_t$. This is represented in Figure 1 by the shift of the AD curve from AD^{old} to AD^1 . The intersection of AD^1 with AS thus identifies the equilibrium pair of \tilde{y} and R that obtain under complete information.

Why is output falling in response to the demand shock? As consumers try to substitute consumption from today to tomorrow, the real interest rate falls. As this happens, it signals to firms that its good time to substitute production intertemporally, which amounts to reducing utilization today. This explains why aggregate output *today* falls in response to the negative demand shock—but also why aggregate output *tomorrow* increases, indeed in a way leaves the permanent income of the household unchanged (i.e. Proposition 32).

Let us now see what changes once information is incomplete. Had \mathcal{B}_t been zero, the shock would have triggered the same shift in AD as under complete information. But because this “first-round” shift in AD causes current aggregate income to fall, and because households confuse aggregate and idiosyncratic movements in current income, they incorrectly perceive a decrease in their permanent income—even though, as already argued above, there is no actual change in their permanent income, only an intertemporal shift. As this happens, the AD curve shifts further down, from AD^1 to AD^2 in Figure 1, triggering a further reduction in aggregate income, which in turn feeds to additional misperception (a further reduction in \mathcal{B}), a further downward shift in the AD curve, and so on.

This series of feedbacks, whose limits corresponds to the “final-round” curve AD^{new} in the figure, and explains the economics behind our “confidence multiplier.” The formal result is offered below.




Figure 1: Feedback Mechanism

Proposition 6 (The i.i.d. Case). *Suppose the demand shock is entirely transitory ($\rho_\beta = 0$). The equilibrium response of real aggregate output is given by*

$$\frac{\partial y_t}{\partial \eta_t} = \gamma_0 \cdot m^{conf}(\lambda, \rho_\xi),$$

where

$$\gamma_0 \equiv \frac{\varsigma \sigma \beta}{\varsigma + \sigma}$$

is the complete-information response of output to i.i.d demand shocks and

$$m^{conf}(\lambda, \rho_\xi) \equiv \frac{\varsigma + \sigma}{\varsigma + \sigma - \varsigma \frac{1-\beta}{1-\beta \rho_\xi} (1-\lambda)} > 1 \quad (37)$$

is the “confidence multiplier,” namely a multiplier that captures the mechanism described above.

The following is then immediate:

Fact 1. $m^{conf}(\lambda, \rho_\xi)$ is decreasing in λ and increasing in ρ_ξ . That is, the confidence multiplier increases with both the severity of the informational friction and the persistence of idiosyncratic income shocks.

The intuition is obvious: a lower λ means more confusion of the aggregate fluctuations for idiosyncratic income fluctuations, and a higher ρ_ξ means more extrapolation of the present to the future and therefore a larger change in perceived permanent income for a given demand shock.

4.4 The General Case: Confidence Plus GE Dampening

We now allow the demand shock to be persistent. This does not affect the logic that the demand shock has no actual effect on aggregate permanent income: the demand shock still represents an intertemporal shift between “the long run” and “the short run,” although the “short run” is now longer.

This also does not affect the size of our confidence multiplier: the ferocity of the feedback loops described above depend only on the magnitude of the aggregate shock, the degree of confusion, and the persistence of the idiosyncratic income.

What changes, though, are the expectations that households form about *future* real interest rates and the effect this has on aggregate demand under both complete and incomplete information. This amounts to an additional multiplier, which we formalize below.

Proposition 7 (The General Case). *The equilibrium response of real aggregate output is given by*

$$\frac{\partial y_t}{\partial \eta_t} = \gamma \cdot m^{\text{conf}}(\lambda, \rho_\xi) \cdot m^{\text{GE}}(\lambda, \rho_\beta)$$

where

$$\gamma \equiv \frac{\beta\sigma\varsigma}{\sigma + \varsigma} \frac{1}{1 - \rho_\beta\beta}$$

is the complete-information response of output to persistent demand shocks, $m^{\text{conf}}(\lambda, \rho_\xi)$ is the same confidence multiplier as that defined in Proposition 6 above, and

$$m^{\text{GE}}(\lambda, \rho_\beta) \equiv 1 + \frac{\rho_\beta\beta\sigma}{\sigma + \varsigma} (1 - \lambda) \geq 1 \quad (38)$$

is another multiplier, which captures the mechanism we refer to as “dampening the neoclassical GE adjustment.”

What exactly is this mechanism? And why does it hinge on the demand shocks being persistent?

When a negative demand shock hits the economy, it triggers *in equilibrium* a reduction in the real interest rate, which in turn moderates the original reduction in aggregate demand. This represents a “neoclassical GE adjustment” that partially offsets the direct, PE effect of the shock on aggregate demand. Furthermore, the GE adjustment in the real interest rate is as persistent as the demand shock itself. For example, if the shock lasts 5 periods, the real interest rate falls for 5 periods as well.

With complete information, households not only see the current drop in the real interest rate, but also perfectly foresee how much the real interest rate will fall during the next 4 periods. With incomplete information, households still observe the current drop in real interest rates but underestimate their future drop. This arrests the forward-looking component of the aforementioned GE adjustment, leaving aggregate demand to move more strongly with the exogenous shock than what it would have done under complete information.

We refer to this mechanism as “dampening the neoclassical GE adjustment.” This mechanism is conceptually distinct from, although complementary to, our confidence multiplier. That mechanism hinges on the confusion of aggregate and idiosyncratic income fluctuations. The present one instead hinges on lack of common knowledge of the aggregate shock, and in particular of the news it contains

about the future: because this news is not common knowledge, households under-estimate the future responses of others, which herein means that they under-estimate the future drop in aggregate spending and the resulting drop in the real interest rate. In short, households “discount” the future GE adjustment in the real interest rate.

At a high level, this mechanism is the same as that articulated in [Angeletos and Lian \(2018\)](#) in the context of forward guidance for monetary policy at the zero lower bound. But there is a subtle difference, on top of the obvious difference in contexts. In the model of that paper, the complete-information GE adjustment *amplifies* the PE effect of the relevant exogenous shock (a change in monetary policy). In our model, instead, the complete-information GE adjustment *moderates* the effect of the relevant exogenous shock (a change in consumer spending). This in turn explains why the same basic mechanism ends up generating seemingly opposite results. In both cases, lack of common knowledge arrests the GE adjustment. But whereas this means arresting GE amplification in [Angeletos and Lian \(2018\)](#), it means arresting GE attenuation in the present paper.³¹

Putting aside this clarification, the key observation for our purposes is that this mechanism works hand-in-hand with our confidence multiplier: they both help sustain larger movements in AD, and thereby in equilibrium output and employment, than under complete information.

We conclude with the following observation regarding comparative statics.

Fact 2. $m^{GE}(\lambda, \rho_\beta)$ is decreasing in λ and increasing in ρ_β . That is, the multiplier due to the dampening of the neoclassical GE adjustment increases with both the severity of the informational friction and the persistence of the aggregate demand shock.

Combining Facts 1 and 2, we thus infer that the *overall* multiplier,

$$m(\lambda, \rho_\xi, \rho_\beta) \equiv m^{\text{conf}}(\lambda, \rho_\xi)m^{GE}(\lambda, \rho_\beta), \quad (39)$$

is increasing in the severity of the informational friction and the persistence of both kinds of shocks.³²

4.5 Signal Extraction and the Fixed Point

In the preceding analysis, we have treated the value of λ as given. However, the severity of the informational friction depends on how much households can extract from the available market signals. This implies that the values of λ in our economy is the solution to a fixed point problem that relates the informativeness of these signals to the responsiveness of aggregate output and the real interest rate

³¹Another way to think about it is that the setting studied in [Angeletos and Lian \(2018\)](#) maps to a game of strategic complementarity, whereas the one studied here maps to a game of strategic substitutability. In both cases, incomplete information amounts to reducing the absolute magnitude of the strategic interaction. But the ultimate effect on behavior switches signs between the two games.

³²Thus, $\bar{m} \equiv m(1, 1, 1) = 1 + \beta + \frac{\xi}{\sigma}$ represent a tight upper bound for the overall multiplier.

to the aggregate demand shock: the greater this response, the more precise the information revealed by the aforementioned signals, and the smaller the friction. This fixed point is characterized below.

Proposition 8. *There exist a unique $\lambda \in (0, 1)$ in equilibrium and is given by the solution of the following fixed point problem:*

$$\lambda = \frac{\sigma_{\beta}^{-2} + \sigma_R^{-2} \{s^{-1}\gamma [m(\lambda, \rho_{\xi}, \rho_{\beta}) - \rho_{\beta}]\}^2 + \sigma_{\xi}^{-2} \{\gamma m(\lambda, \rho_{\xi}, \rho_{\beta})\}^2}{\sigma_{AD}^{-2} + \sigma_{\beta}^{-2} + \sigma_R^{-2} \{s^{-1}\gamma [m(\lambda, \rho_{\xi}, \rho_{\beta}) - \rho_{\beta}]\}^2 + \sigma_{\xi}^{-2} \{\gamma m(\lambda, \rho_{\xi}, \rho_{\beta})\}^2}, \quad (40)$$

where $m(\lambda, \rho_{\xi}, \rho_{\beta})$ is the overall multiplier defined in (39).

The above condition, which follows from the standard formula for combining multiple Gaussian signals, reflects the fact that the consumer has three such signals about the underlying AD shock: her own discount rate, β_t^h ; her local interest rate $R_{h,t}$; and her own income, or the local demand, $y_{h,t}$. And whereas the informativeness of the first signal is exogenous, that of the other two signals is endogenous. In particular, their informativeness increase with the multiplier: the larger the multiplier, the more these market signals moves with the AD shock, and hence the more precise the information households can extract from them about the underlying shock. This explains why the right-hand side of the above condition decreases with λ , which in turn guarantees the uniqueness of the equilibrium.

It is also immediate to verify that, for given ρ_{ξ} and ρ_{β} , λ is a decreasing function of the ratio σ/σ_{AD} for any $\sigma \in \{\sigma_{\beta}, \sigma_R, \sigma_{\xi}\}$. By varying these “noise” parameters, we can indeed induce any value for λ in the $(0, 1)$ range. This, together with the uniqueness of λ , justifies our earlier treatment of it as a “free” parameter.³³ Finally, by combining the comparative statics of m with the comparative statics of the fixed point in condition (40) we reach the follow result.

Proposition 9. *Let m^* denote the overall multiplier evaluated at the equilibrium value of λ . This is necessarily increasing in the ratio σ/σ_{AD} for any $\sigma \in \{\sigma_{\beta}, \sigma_R, \sigma_{\xi}\}$, as well as in both ρ_{ξ} and ρ_{β} .*

That is, our statement “a larger friction implies a larger multiplier,” which was previously formalized with the monotonicity of $m(\lambda, \rho_{\xi}, \rho_{\beta})$ in λ , is now recast as the monotonicity of m^* in the relevant noise parameters. And the insights about how higher persistence reinforces the two mechanisms at work, which were previously articulated holding λ constant, are robust to taking into account the endogeneity of λ .

5 Extensions

In this section we first discuss how the mechanisms identified above can be recast as the product of bounded rationality, or lack of sophistication, instead of informational friction. We then consider

³³Of course, one could try discipline the triplet $(\sigma_{\beta}, \sigma_R, \sigma_{\xi})$, and thereby λ , with micro-level data on spending, interest rates, and income and/or with surveys of expectations. But this is beyond the (theoretical) scope of this paper.

three extensions. The first two replace the aggregate discount-factor shock with, respectively, a shock in government spending and a shock in aggregate technology. This allows us to draw lessons for fiscal policy and to explain why the friction works asymmetrically over aggregate demand and aggregate supply shocks. The third variant adds investment, letting us show how our theory helps resolve the “co-movement puzzle.”

5.1 Bounded Rationality

Our confidence multiplier originates from an informational friction in the tradition of [Lucas \(1972\)](#) and [Sims \(2003\)](#). Empirical evidence in support of such a friction abounds.³⁴ But a similar feedback loop may emerge in a plausible variant in which agents are fully informed but not fully rational.

Suppose, in particular, that agents use a simple, random-walk model to extrapolate current income to future income:

$$E_t^h [y_{h,t+k}] = y_{h,t}. \quad (41)$$

This resembles the form of simple, “extrapolative” beliefs found in [Barberis et al. \(2015\)](#) and [Gennaioli, Ma, and Shleifer \(2016\)](#). It also proxies an optimally “sparse” representation of the individual income process, along the lines of [Molavi \(2019\)](#): in a world where individual income is driven by both random-walk idiosyncratic shocks (like $\xi_{h,t}$ in our model) and transitory aggregate shocks (like β_t in our model) but agents can only use a lower-dimension, one-factor, statistical model to track current income and form expectations about future income, the optimal such model converges to the model that governs the idiosyncratic fluctuations alone as these fluctuations get larger and larger relative to the aggregate ones.

In this variant, aggregate demand can still be expressed as in [Proposition 2](#). What changes is only the characterization of the confidence term \mathcal{B}_t : whereas in the original model we have

$$\mathcal{B}_t = \frac{(1-\beta)(1-\lambda)}{\beta(1-\beta\rho_\xi)} \frac{\partial y_t}{\partial \eta_t} \eta_t,$$

in the present variant we have

$$\mathcal{B}_t = \frac{1}{\beta} \frac{\partial y_t}{\partial \eta_t} \eta_t.$$

It is therefore as if we are back in our original model but have set $\lambda = 0$ and $\rho_\xi = 1$ (maximal informational friction and random-walk idiosyncratic shocks). This illustrates the point anticipated in the Introduction: our confidence multiplier can be recast as the product of a plausible form of

³⁴For example, [Coibion, Gorodnichenko, and Ropele \(2018\)](#) show, using a survey by the Bank of Italy, that firms ignore relevant, public information about the state of the economy: they change both their expectations and their behavior when such information is made salient to them. [Cavallo, Cruces, and Perez-Truglia \(2017\)](#), on the other hand, use an experiment to show that consumers do not process information optimally even when such information is readily available to them. See also [Gabaix \(2019\)](#) for a review of the broader empirical and theoretical literatures on inattention.

bounded rationality, or lack of sophistication, instead of lack of information.³⁵

A similar point applies to the discounting of the GE adjustment in the real interest rate, or the \mathcal{G}_t wedge. While our model generates this mechanism with the help of incomplete information as in Angeletos and Lian (2018), a few other papers (Gabaix, 2020; Farhi and Werning, 2019; Iovino and Sergeyev, 2017) capture a similar mechanism with plausible departures from rationality.

Suppose, in particular, that, while a fraction of $\ell \in (0, 1)$ of the consumers are both fully informed and fully rational, the remaining fraction is “level-1 thinkers” in the sense of Farhi and Werning (2019): although they themselves experience an urge to consume less, they don’t expect *others* to spend less in the future. These agents expect future aggregate demand, and hence also the future real interest rate, to remain unchanged. It follows that the relevant term is given given by

$$\mathcal{G}_t = (1 - \ell)\sigma \sum_{k=1}^{+\infty} \beta^k \frac{\partial \mathbb{E}_t[R_{t+k}]}{\partial \eta_t} \eta_t.$$

By comparison, in our baseline model this term satisfies

$$\mathcal{G}_t = (1 - \lambda)\sigma \sum_{k=1}^{+\infty} \beta^k \frac{\partial \mathbb{E}_t[R_{t+k}]}{\partial \eta_t} \eta_t.$$

Clearly, the two terms coincide when $\lambda = \ell$. The GE attenuation in our original model is therefore exactly the same as that in the present variant. The only twist is that $1 - \lambda$ can now be re-interpreted as the lack of sophistication rather than the lack of information.

5.2 Government Spending

We now return to our original model, with full rationality but incomplete information, and consider a different kind of aggregate demand shocks, namely shocks to government spending. Unsurprisingly, these shocks have similar effects as those featured in our baseline analysis, both with and without common knowledge. But there two interesting twists, one regarding the distinction between front-loaded and back-loaded fiscal stimuli and another regarding Ricardian equivalence.

The model is the same, except for two modifications. First, we shut down the wealth effect on labor supply.³⁶ And second, we drop the shocks to the household’s discount factor and, instead, add

³⁵Suppose that, instead of (41), the subjective expectations of income satisfy

$$E_t^h[y_{h,t+k}] = \hat{\rho}^k y_{h,t},$$

for some $\hat{\rho} \in (0, 1)$. This would let $\hat{\rho}$ govern the degree of over-confidence during booms (and that of under-confidence during recession), similarly to what ρ_ξ and $1 - \lambda$ and do in our main analysis. The essence thus remain the same. But the math gets more complicated, which explains why we have herein opted to fix $\hat{\rho} = 1$.

³⁶It is well understood both that an increase in government spending can stimulate employment via this channel in theory, and that this channel is probably very weak in practice. Formally, we set $\sigma = 0$ in the optimal labor supply (14), while

shocks to government spending.

In each period, the government purchases a basket of all the goods in the economy and uses that to produce a public good, whose quantity we denote by G_t . We treat G_t as exogenous and random. Specifically, we assume that G_t is drawn from an AR(1) process with persistence $\rho_G \in [0, 1)$ and independent of any other shock:

$$G_t = \rho_G G_{t-1} + \eta_t^G, \quad (42)$$

where $\eta_t^G \sim N(0, \sigma_G^2)$ and the initial G_{-1} is set to be 0. We also let the composition of government spending mirror that of private spending.³⁷

Turning to taxes, we let government spending be financed by lump-sum taxation and, without further loss of generality, impose budget balance in each period, i.e. $G_t = T_t$. At the same time, we prevent one's own tax from being a perfect signal of the aggregate level of government spending by introducing idiosyncratic tax shocks. Specifically, we let the lump-sum tax levied to household h in period t be given by $T_t^h = T_t + \Delta_t^h$, where T_t is the average tax and $\Delta_t^h \sim N(0, \sigma_\Delta^2)$ also follows an AR(1) process with persistence ρ_G , and independent from any other shock.³⁸

We close the model by letting the information structure be the same as in the baseline model, modulo, of course, that the knowledge of η_{t-1} and β_t^h is now replaced with the knowledge of η_{t-1}^G and T_t^h . That is, $\mathcal{I}_t^h = \mathcal{I}_{t-1}^h \cup \{T_t^h\} \cup \{w_{h,t}, e_{h,t}, R_{h,t}, (p_{i,l,t})_{i \in \mathcal{C}_t^h, l \in [0,1]}\} \cup \{\eta_{t-1}^G\}$.

As in the baseline model, we log-linearize the equilibrium conditions, re-interpret all the variables as log-deviations from their steady-state counterparts.³⁹ Following similar steps as in Section 4, we reach the following result.

Proposition 10. (i) *Aggregate supply remains the same as in Proposition 1.*

preserving $\sigma > 0$ in the optimal consumption function (24). This modifies the definition of $\bar{\alpha}$ but does not otherwise affect any of the preceding results. The same is true if we let the real wage be determined in a non-competitive fashion, as a function of employment and income but not of government spending or taxes. Finally, similar results obtain if we assume GHH preferences.

³⁷That is, we impose, for all i, j , and t ,

$$\frac{p_{i,j,t} g_{i,j,t}}{G_t} = \frac{\int_{\{h:i \in \mathcal{C}_t^h\}} p_{i,j,t} c_{i,j,t}^h dh}{c_t},$$

where $g_{i,j,t}$ is the government's purchase of variety j from island i in period t . This restriction involves equilibrium outcomes in the right hand side, but can be derived from first principles by interpreting G_t as the quantity of a public good produced with a CES composite of all the varieties, letting this CES composite mirror that describing private preferences, and finally having the government optimally choose its purchases of all the variates so as to minimize the cost of producing the exogenously given quantity, G_t , of the public good. Also note that the above condition, together with the definition of C_t , implies $\int_{i \in [0,1]} \int_{j \in [0,1]} p_{i,j,t} g_{i,j,t} dj di = G_t$.

³⁸By equating the persistence of the idiosyncratic tax shock to that of aggregate government spending, we make sure that households correctly perceive the tax burden of any fiscal stimulus. Relaxing this property may be interesting on its own right—see, for example, Gabaix (2020) for an example that goes in this direction with the help of bounded rationality instead of incomplete information—but it is not the issue we wish to study here.

³⁹The following exception applies: $G_t, T_t, g_{i,j,t}$ and T_t^h henceforth represent, respectively, $G_t/y^*, T_t/y^*, g_{i,j,t}/y^*$, and T_t^h/y^* , where y^* is the steady-state (also, complete-information) value of aggregate output. This is a standard trick in the literature on fiscal multipliers (e.g., Woodford, 2011) and it simply takes care of the issue that the log-deviation of the government spending is not well defined when its steady-state value is 0.

(ii) Aggregate demand satisfies

$$y_t = -\sigma R_t + (1 - \rho_G) G_t + \mathbb{E}_t [y_{t+1}] + (\mathcal{B}_t + \mathcal{G}_t), \quad (43)$$

and \mathcal{B}_t and \mathcal{G}_t are defined the same as in (30) and (27), modulo the replacement of ρ_β with ρ_G .

(iii) The equilibrium response of real aggregate output is given by

$$\frac{\partial y_t}{\partial \eta_t^G} = \gamma^G \cdot m^{\text{conf}}(\lambda, \rho_\xi) \cdot m^{\text{GE}}(\lambda, \rho_G),$$

where $\gamma^G = \frac{(1-\rho_G)\beta}{1+\sigma\zeta^{-1}}$ is the complete-information counterpart and $m^{\text{conf}}(\cdot, \cdot)$ and $m^{\text{GE}}(\cdot, \cdot)$ are the same multipliers as those in Propositions 6–7.

In a nutshell, $(1 - \rho_G) G_t$ replaces $-\sigma\beta_t$ as the AD shock in Proposition 2. The rest is the same. And the following is then immediate.

Fact. With complete information, $\frac{\partial y_t}{\partial \eta_t^G} < 1$ necessarily. With incomplete information, instead, $\frac{\partial y_t}{\partial \eta_t^G} > 1$ is possible insofar as $\bar{m}\gamma^G > 1$, where the maximum of the overall multiplier defined in (32). That is, whereas an increase in government spending crowds out private consumption under complete information, it can crowd in under complete information.

Our theory therefore helps accommodate the Keynesian narrative of how a fiscal stimulus can help boost consumer spending despite the increase in the tax burden. But unlike the Keynesian framework, our theory does not rely on nominal rigidity and not require the boom to be inflationary.

The fiscal-policy predictions of our theory are even more distinct from those of the Keynesian paradigm if we shift the focus from “front-loaded” to “back-loaded” fiscal stimuli. By front-load stimuli we mean unanticipated shocks to *current* government spending, such as those modeled above. By back-load stimuli, we instead have news about *future* government spending.

To capture the latter, let $G_t = \eta_{t-1}^{\text{news}}$, where η_{t-1}^{news} is i.i.d. across time and realized at $t - 1$. In this case, the AS curve remains again the same, whereas the AD curve can be written as

$$Y_t = -\sigma R_t - \eta_t^{\text{news}} + \mathbb{E}_t [y_{t+1}] + \mathcal{B}_t + \mathcal{G}_t.$$

News of a future fiscal stimulus are therefore akin to a *negative* demand shock today. The following is then immediate.

Fact. While an increase in today’s government spending is expansionary, news of an increase in tomorrow’s government spending are contractionary, and the more so the larger the informational friction.

A back-loaded fiscal stimulus is therefore contractionary in our economy. This is true even in the absence of informational friction, thanks to optimal intertemporal reallocation in production: when

the economy expects an increase in the demand for goods tomorrow, it responds by economizing on its use of resources today. But when information is imperfect, the reduction in current income is mis-perceived by the consumers as a reduction in permanent income. Our confidence multiplier thus kicks in, amplifying the contraction trigger by news of future increases in government spending.

5.3 Investment

As noted earlier, our baseline analysis allows for variable utilization of capital (or “land”) but abstracts from investment. If we add investment but let it be decided under complete information, we run on a familiar problem: whenever consumers spend less, the cost of investment (the real interest rate) goes down in equilibrium, causing investment to move in the opposite direction than consumption, which is at odds with the data. We now show how this negative co-movement problem is resolved when our confidence multiplier extends from consumption to investment.

To keep the analysis tractable, we work with a two-period version of our model. Accordingly, household’s preferences are thus given by

$$\mathcal{U}^h = U(c_0^h, n_0^h) + \beta^h U(c_1^h, n_1^h),$$

where U is as in (2) and

$$\log \beta^h = \log \eta + \log \epsilon^{\beta, h}$$

where $\eta \sim \mathcal{N}(0, \sigma_{AD}^2)$ and $\epsilon^{\beta, h} \sim \mathcal{N}(0, \sigma_\beta^2)$ are, respectively, aggregate and idiosyncratic discount rate shocks. As in the case of government spending, we also shut down the wealth effect on labor supply.

Next, to apply the exact same informational friction to investment as that applied to consumption, we let both choices be made by the household and under the same informational friction. We accordingly re-interpret land as capital (and firm earnings net of labor costs are returns to capital) and modify its law of motion as follows:

$$k_2^h = \left[1 - \delta(u_h) + \Psi(\iota^h) \right] k_1^h \epsilon_k^h,$$

where u_h is the utilization rate, ι^h is the investment rate (per unit of capital), ϵ_k^h is an idiosyncratic shock, δ is an increasing and strictly convex function (as before), $\Psi(\iota) = \iota - \frac{\psi}{2}(\iota - \iota^*)^2$ is an increasing and concave function, and ι^* denotes the no-shock (“steady-state”) investment rate.

The concavity of Ψ introduces adjustments costs to capital, as in standard Q theory, and ψ parameterize the elasticity of investment with respect to net returns, As for the idiosyncratic shock ϵ_k^h , this plays an inessential, auxiliary role: the household’s knowledge of her own capital stock k_2^h will not

reveal the aggregate demand shock.⁴⁰

Because the technology for producing the final good is the same as in our baseline model and because utilization is still made under complete information, the aggregate supply is also the same. What changes is only aggregate demand, which now has two components: a consumption component determined as before, and an investment component.

As shown in the Appendix, the equilibrium rate of investment can be expressed as a decreasing function of the real interest rate and an increasing function of the same confidence term \mathcal{B} as that appearing in consumption. The negative dependence of investment on the real interest rate captures the standard, neoclassical, cost channel. Its positive dependence on \mathcal{B} captures the “investor sentiment” translation of our mechanism: during a demand-driven recession, households rationally mis-perceive not only an increase in their permanent income but also an increase in the returns to capital.

The first channel contributes to negative co-movement between consumption and investment. The second contributes to positive co-movement. And since the importance of the second channel is regulated by the informational friction, one may hope that this channel dominates when the informational friction is large enough. The next result verifies this logic subject to a small qualify.

Proposition 11 (Investment). *There exists $\underline{\kappa} > 0$ and $\underline{\rho} > 0$, such that whenever $\kappa > \underline{\kappa}$ and $\rho_\xi > \underline{\rho}$, the following is true: the equilibrium levels of employment, output, investment, and consumption positively co-move in respond to demand shocks if and only if the informational friction is large enough, namely $\lambda < \hat{\lambda}$, for some $\hat{\lambda} \in (0, 1)$ that is itself increasing in κ and ρ_ξ .*

5.4 Technology Shocks

We now turn to aggregate supply shocks, in the sense of aggregate technology shocks. The main lesson delivered below is that the information friction works asymmetrically between demand and supply shocks, favoring the former as the main source of business cycles.

The model is the same as in our baseline analysis (with an infinite horizon and no investment), except that we shut down the shocks to the household’s discount factor and, instead, introduce shocks to the technology of each firm:

$$y_{i,j,t} = A_t (l_{i,j,t})^\alpha (u_{i,j,t} k_{i,j,t})^{1-\alpha},$$

⁴⁰That is, we specify the information set of household/consumer/investor h in period 1 as follows:

$$\mathcal{I}_1^h = \left\{ \beta_1^h \right\} \cup \left\{ w_{h,1}, e_{h,1}, R_{h,1}, (p_{i,l,1})_{i \in \mathcal{C}_1^h, l \in [0,1]} \right\} \cup \left\{ k_2^h \right\}. \quad (44)$$

This is essentially the same as that in (11), except that we have added knowledge of the end-of-period, local capital stock.

where A_t is an aggregate TFP shock. This follows an AR(1) process with persistence $\rho_A \in [0, 1)$:

$$\log A_t = \rho_A \log A_{t-1} + \eta_t^A, \quad (45)$$

where $\eta_t^A \sim N(0, \sigma_A^a)$ and A_{-1} is normalized at 1. Finally, the household's information set in period t is given by $\mathcal{I}_t^h = \mathcal{I}_{t-1}^h \cup \left\{ w_{h,t}, e_{h,t}, R_{h,t}, (p_{i,l,t})_{i \in \mathcal{C}_t^h, l \in [0,1]} \right\} \cup \{A_{t-1}\}$. It follows that, although the identity of the aggregate shock is different, local income and local interest rates serve again as noisy signals of the underlying aggregate shock.

As in the baseline model, we log-linearize the equilibrium conditions, re-interpret all the variables as log-deviations from their steady-state counterparts,. Following similar steps as in Section 4, we reach the following characterization of the aggregate supply and aggregate demand curves.

Proposition 12. (i) *Aggregate supply is given by*

$$y_t = A_t + (1 - \tilde{\alpha})(u_t + k_t), \quad (46)$$

$$u_t = \frac{\beta}{\tilde{\alpha} + \beta\phi} (A_t - \mathbb{E}_t[A_{t+1}]) + \frac{\beta}{\tilde{\alpha} + \beta\phi} R_t + \beta \mathbb{E}_t[u_{t+1}], \quad (47)$$

$$k_{t+1} = k_t - \psi u_t, \quad (48)$$

(ii) *Aggregate demand is given by*

$$y_t = -\sigma R_t + \mathbb{E}_t[y_{t+1}] + (\mathcal{B}_t + \mathcal{G}_t), \quad (49)$$

where \mathcal{B}_t and \mathcal{G}_t are defined as (27) and (30).

The characterization of AS is similar as before, except, of course, for the accommodation of the TFP shock. This shows up, not only in (46), the production function, but also in (47), the optimality condition for utilization. Intuitively, the optimal intertemporal pattern of production depends, not only the relative price of today's goods, but also on their relative cost. The former is captured in condition (47) by R_t , the latter by $A_t - \mathbb{E}_t[A_{t+1}]$.

Turning to AD, \mathcal{B}_t and \mathcal{G}_t have the same mathematical definition. However, \mathcal{B}_t here not only contains misperception of idiosyncratic permanent income but also misperception of aggregate permanent income.

In the case of demand shocks, the latter effect was ruled out thanks to Proposition 3: because the "neoclassical" structure of the economy guarantees that demand shocks do not influence aggregate permanent income (and because this fact is common knowledge), demand shocks can never trigger rational mistakes in the calculation of aggregate permanent income. With supply shocks, things are different: because a positive aggregate TFP shock naturally increases aggregate permanent income (and because this fact, too, is common knowledge), imperfect knowledge of the aggregate TFP shock

naturally translates to under-estimation of the corresponding movement in aggregate permanent income.

Proposition 13 (Technology Shocks). *The equilibrium response of aggregate output to an aggregate technology shock is given by*

$$\frac{\partial y_t}{\partial \eta_t^A} = \gamma^A \cdot m_A^{\text{conf}}(\lambda, \rho_\xi, \rho_A) \cdot m_A^{\text{GE}}(\lambda, \rho_A)$$

where $\gamma^A > 0$ is the complete-information counterpart, $m_A^{\text{GE}}(\lambda, \rho_A)$ is a multiplier that captures the effect of \mathcal{G}_t , or the dampening of the GE movement in the real interest rate, and $m_A^{\text{conf}}(\lambda, \rho_\xi, \rho_A)$ is a multiplier that captures the combined effect of \mathcal{B}_t , or the overall confusion of idiosyncratic and aggregate income fluctuations. Furthermore,

$$m_A^{\text{conf}}(\lambda, \rho_\xi, \rho_A) < 1 \text{ if and only if } \rho_\xi < \bar{\rho}_\xi$$

for some $\bar{\rho}_\xi = \bar{\rho}_\xi(\lambda, \rho_A)$ that is itself increasing in ρ_A , and

$$m_A^{\text{GE}}(\lambda, \rho_A) < 1 \text{ necessarily}$$

The exact characterization of the multipliers can be found in the Appendix. The main lessons are two. First, the confusion of aggregate and idiosyncratic shocks now has an ambiguous effect: m_A^{conf} can be either higher or lower than one, depending on the persistences of the two shocks. Second, the effect of the as-if discounting of the GE adjustment in the real interest rate is both unambiguous and of the opposite direction than that in the case of demand shocks: m_A^{GE} is necessarily *less* than one.

To understand the first property, ignore momentarily the second mechanism, and consider how “confidence” responds to a positive TFP shock. With incomplete information, consumer perceives an increase in the idiosyncratic component of their permanent income, which pushes \mathcal{B}_t positive, but also under-estimate the increase in the aggregate permanent income, which pushes \mathcal{B}_t negative.

Clearly, these two effects work in opposite directions: the one increases aggregate demand, the other decreases it. As explained in the Appendix, their relative strength depends subtly on the signal-extraction problem, but the following basic intuition is valid: the lower the persistence of the idiosyncratic shock relative to that of the aggregate shock, the more likely it is that $\mathcal{B}_t < 0$, i.e., a positive TFP shock leads to a *reduction* in confidence. This explains why $m_A^{\text{conf}} < 1$ as long as ρ_ξ is small enough.

Consider now the second mechanism, the discounting of the GE adjustment in the real interest rate. A persistent positive TFP shock triggers a drop in real interest rate both now and in the future. When information is incomplete, consumers under-estimate the future drop in the real interest rate, and this *reduces* aggregate demand relative to the complete-information case. This explains why

$m_A^{\text{GE}} < 1$ necessarily.

To sum up, as long as idiosyncratic and aggregate income fluctuations have comparable persistence, the combination of the two mechanisms dampens the economy's response to supply shocks at the same time that it amplifies its response to demand shocks. This completes both the “Keynesian flavor” of our theory and its fits vis-a-vis a key empirical regularity: the majority of the business cycle in the data is disconnected from both technology and inflation.

5.5 Sticky Prices (or Other Wedges)

The starting point of our paper was the desire to accommodate the Keynesian narrative of demand-driven fluctuations outside the nexus of sticky prices and Philips curves. As explained in the Introduction, this desire was grounded on both “philosophical” predispositions (a priori that the Keynesian narrative does not hinge on nominal rigidity) and empirical considerations (the lack of significant co-movement between inflation and real economic activity).

But the mechanisms we have identified do not hinge on the absence of nominal rigidity. They naturally extend to a New Keynesian context. The only twist is that the interaction of informational friction with nominal rigidity allows for an *additional* mechanism, one that regards the expectations of future output gaps.

To illustrate these points, consider a modification of our baseline model that allows for sticky prices. As well known (e.g., [Correia, Nicolini, and Teles, 2008](#)), this is equivalent to maintaining flexible prices but allowing for a time-varying tax, or wedge, in production, which is effectively under the control of monetary policy. Under this representation, a “hawkish” monetary policy that stabilizes inflation maps to a constant wedge, whereas an “accommodative” monetary policy that lets positive demand shocks trigger inflation and positive output gaps maps to a counter-cyclical wedge: it is *as if* there is a subsidy on production whenever consumers spend more.

With this in mind, we bypass the details of how monetary policy is conducted (via a Taylor rule or in some other way) and instead work directly with the following modification of Proposition 4, which is valid whether the relevant wedge, denoted below by τ_t , is the product of nominal rigidity or of an actual tax on revenue.

Proposition 14 (AS and AD with sticky prices). *Let τ_t denote the production wedge. Aggregate supply is given by*

$$y_t = (1 - \tilde{\alpha})(u_t + k_t) - \tau_t, \quad (50)$$

$$u_t = \frac{\beta}{\tilde{\alpha} + \beta\phi} R_t - \frac{\beta}{\tilde{\alpha} + \beta\phi} \frac{v + \sigma + \tilde{\alpha}\sigma}{\tilde{\alpha}\sigma} (\tau_t - \mathbb{E}_t[\tau_{t+1}]) + \beta \mathbb{E}_t[u_{t+1}], \quad (51)$$

$$k_{t+1} = k_t - \psi u_t, \quad (52)$$

Aggregate demand is given by

$$y_t = -\sigma \{R_t + \beta_t\} + \mathbb{E}_t[y_{t+1}] + (\mathcal{B}_t + \mathcal{G}_t), \quad (53)$$

where \mathcal{B}_t and \mathcal{G}_t are defined as (27) and (30).

The presence of τ_t in (50) and (51) captures the effects of the tax on, respectively, employment and utilization—or, equivalently, the power of monetary policy over aggregate supply once prices are sticky. The textbook version of the New Keynesian model, which abstracts from variable utilization, corresponds to either $\tilde{\alpha} = 1$ (utilization is unproductive) or $\phi \rightarrow \infty$ (variation in utilization is prohibitively costly). Aggregate supply then reduces to $y_t = 0 - \tau_t$, where 0 stands for the natural rate of output and τ_t for the wedge, or the output gap, induced by any monetary policy that does not replicate flexible prices. Relative to this familiar case, the key supply-side novelty of our analysis is to let the natural rate of output to be sensitive to the real interest rate, in the manner explained in Section 3.

Let us now turn to aggregate demand, or equation (53) above. In the textbook version of the New Keynesian, this equation holds with $\mathcal{B}_t = \mathcal{G}_t = 0$. Relative to this case, we see that the informational friction continues to give rise to our two mechanisms, captured by the exact same terms \mathcal{B}_t and \mathcal{G}_t as in our baseline analysis.

However, because the GE adjustment in the real interest rate is now modulated by monetary policy, the magnitude of \mathcal{G}_t now depends on monetary policy. Similarly, and more crucially for our narrative about confidence, \mathcal{B}_t here contains not only misperceptions of the “natural” level of permanent income but also misperception of the output gaps induced by monetary policy. In particular, \mathcal{B}_t can be decomposed as follows:

$$\mathcal{B}_t = \mathcal{B}_t^{\text{natural}} + \mathcal{M}_t - \bar{\mathbb{E}}_t[\mathcal{M}_t],$$

where

$$\mathcal{B}_t^{\text{natural}} \equiv \frac{1-\beta}{\beta} \sum_{k=0}^{+\infty} \beta^k \int E_t^h[\xi_{h,t+k}] dh$$

is the value of \mathcal{B}_t that obtains when monetary policy replicates flexible prices (equivalently, the value

of \mathcal{B}_t in our baseline analysis),

$$\mathcal{M}_t \equiv \sum_{k=0}^{+\infty} \beta^k \mathbb{E}_t[\tau_{t+k}]$$

is a measure of how much monetary policy deviates from that benchmark, and $\mathcal{M}_t - \bar{\mathbb{E}}_t[\mathcal{M}_t]$ is the corresponding average misperception.

To put more structure on the new term, let us henceforth assume that monetary policy is such that

$$\tau_t = -\varphi\eta_t + \rho_\tau\tau_{t-1} \quad (54)$$

where $\varphi \geq 0$ indexes the degree of policy accommodation, or the size of the associated output gaps, and $\rho_\tau \in [0, 1)$ indexes their persistence. Under this representation, a “hawkish” monetary policy that stabilizes inflation and replicates flexible-price outcomes corresponds to $\varphi = \rho_\tau = 0$, whereas a monetary policy that smooths the real interest rate and lets positive demand shocks trigger inflation and persistent output gaps maps to $\varphi > 0$ and $\rho_\tau > 0$.⁴¹

By direct analogy to what was shown earlier for \mathcal{G}_t , we now have that the gap between the aforementioned PDV and its incomplete-information counterpart is given by

$$\mathcal{M}_t - \bar{\mathbb{E}}_t[\mathcal{M}_t] = -\frac{\varphi}{1 - \beta\rho_\tau}(1 - \lambda)\eta_t. \quad (55)$$

That is, as long as $\lambda < 1$ and $\varphi > 0$, a positive demand shock decreases negative \mathcal{B}_t .

What does this mean? As long as $\varphi > 0$, monetary policy lets output expand beyond its natural rate in response to a positive demand shock. This translates to an increase in *true* aggregate permanent income, which is perfectly forecasted under complete information ($\lambda = 1$) but imperfectly so under incomplete information ($\lambda < 1$). It follows that, as long as information is incomplete, consumers underestimate the increase in aggregate permanent income sustained by an accommodative monetary policy. And because the true increase in aggregate permanent income is larger when the output gaps induced by monetary policy are themselves larger (higher φ) or more persistent (higher ρ_τ), the size of belief mistake is larger under the same circumstances.

This effect is reminiscent of the one regarding aggregate technology shocks in Section 7. As in that case, the underestimation of the variation in aggregate permanent income works in the opposite direction than our confidence multiplier. But whereas in that context the strength of this countervailing

⁴¹ It is possible to translate these parameters in the language of Taylor rules and Philips curves: a higher φ corresponds to a flatter Taylor rule and/or a flatter Philips curve (because such flatness regulates the size of the output gaps), and a higher ρ_τ corresponds to a higher persistence in the underlying demand shock, a larger smoothness parameter in the Taylor rule, or a larger backward-looking component in the Philips curve (because all these parameters regulate the persistence of the output gaps). This translation may be useful for quantitative explorations (which are left outside the paper), but adds little insight. At the end of the day, the critical question vis-a-vis the theory is whether and by how much monetary departs from the benchmark of replicating flexible-price outcomes—and this is what we shall explain below.

effect was pinned down by exogenous primitives, here its is regulated by monetary policy.

In the sense just described, an accommodative monetary policy goes *against* our confidence multiplier. But such a policy also *complements* our confidence multiplier by helping aggregate supply be more responsive to aggregate demand under sticky prices than under flexible prices. To see what we mean by this, let us shut down variable utilization. In this case, the flexible-price AS curve is vertical and the natural rate of output is invariant to AD shocks. It follows that, as long as monetary policy replicates flexible prices ($\varphi = 0$), our confidence multiplier is switched off regardless of how large the informational friction is. But as soon as monetary policy is accommodative ($\varphi > 0$), our confidence multiplier is on under sticky prices, even though it is off under flexible prices.

Perhaps more interestingly, our confidence multiplier helps amplify the power of monetary policy itself. To see this, abstract from the exogenous shock to consumer spending and, instead, modify (56) as follows:

$$\tau_t = \rho_\tau \tau_{t-1} - \eta_t^{MP} \quad (56)$$

where η_t^{MP} represents a pure policy shock, independent of any other shock in the economy. It is then immediate to show that, while the informational friction dampens the effect of this shock via (55), it amplifies it via \mathcal{B}_t . That is, our confidence multiplier helps monetary policy to have a large effect on aggregate demand even the expansion is short-lived and consumers are not liquidity constrained.

We conclude with the following comment, which circles back to our starting point about the apparent prevalence of non-monetary, non-inflationary, demand-driven fluctuations in the data. Such fluctuations are captured here when monetary policy is non-accommodative, and correspond to setting $\tau_t = 0$. Conversely, a more accommodative monetary policy maps, not only to a larger absolute value for $\mathcal{M}_t - \bar{\mathbb{E}}_t[\mathcal{M}_t]$, but also to a stronger pro-cyclicality between inflation and real economic activity. Indeed, in the textbook version of the New Keynesian model, inflation is proportional to the PDV of the output gaps, which is the object behind $\mathcal{M}_t - \bar{\mathbb{E}}_t[\mathcal{M}_t]$.

From this perspective, and unless the Philips curve is extremely flat, the lack of significant co-movement between inflation and real economic activity suggests that τ_t may be small in practice. This is consistent both with the empirical template provided in [Angeletos, Collard, and Dellas \(2020\)](#), which we mentioned before and we review in the concluding section, and with the view that real-world policymakers have done a good job in stabilizing inflation (e.g., [McLeay and Tenreyro, 2020](#)).

6 Conclusion

We revisited the question of why shifts in consumer spending may trigger a recession. Unlike the Keynesian framework, our theory did not rely on nominal rigidity or the failure of monetary policy to replicate flexible prices. Instead, it combined two “neoclassical” elements: variable utilization (or

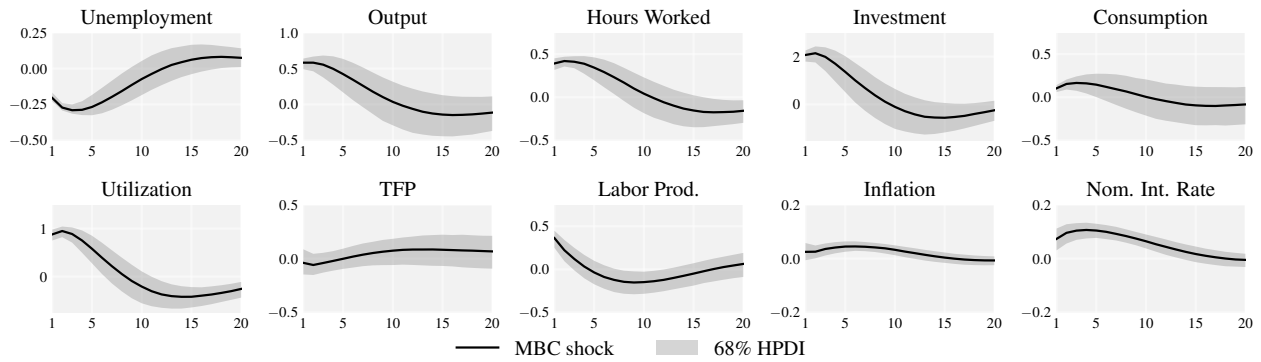


Figure 2: The “Main Business Cycle Shock” in the Data

some other margin of intertemporal substitution) on the supply side, and an informational friction (or bounded rationality) on the demand side.

The first element allowed the aggregate production of today’s goods to be responsive to intertemporal preference shocks, or other shifts in aggregate demand. The second element amplified these shifts, and also helped investment and consumption to comove, by introducing a “confidence multiplier,” namely a positive feedback between loop between current economic activity, consumer perceptions of permanent income, and investor expectations of returns.

The same mechanisms were shown to dampen the fluctuations caused by TFP shocks at the same time that they amplified the fluctuations caused by shocks to consumers spending. And both of these properties were true without accommodative monetary policy and commensurate movements in inflation. Our theory was thus consistent not only with the Keynesian view that the majority of business cycles are demand-driven but also with the empirical template that [Angeletos, Collard, and Dellas \(2020\)](#) provide for the “main business cycle shock” in the data.

Figure 6 illustrates the IRFs of a few key macroeconomic variables to the aforementioned shock. The particular version of this shock considered here is identified by running a VAR on these and a few additional variables and by maximizing the shock’s contribution to the fluctuations of unemployment at the business-cycle frequencies. But as shown in [Angeletos, Collard, and Dellas \(2020\)](#), the picture is basically the same, both in terms of IRFs and in terms of variance contributions, if the shock is identified by maximizing its contribution to the business-cycle fluctuations of output, hours works, investment, consumption, or utilization.

For our purposes, the key take-home lessons from this “anatomy” of the data is the following.

First, there is significant co-movement between unemployment, output, hours worked, investment, and consumption, without commensurate co-movement in either TFP or inflation. As emphasized earlier, this exactly what our theory has delivered (qualitatively, of course).

Second, there is significant pro-cyclical movement in utilization, which in turn appears to account for the pro-cyclical movement in labor productivity despite the absence of pro-cyclical movements

in TFP. This is again in line with our theory, and in particular with the property that labor demand increases in response to a positive demand shock without the intermediation of nominal rigidity and accommodative monetary policy.

Third, there is pro-cyclical movement in both the nominal interest rate and the real interest rate.⁴² This is in line with both a non-accommodative monetary policy and an increase in the relative price of today's goods.

Last but not least, the response of utilization (and to a lesser extent that of unemployment and output, too) reverses sign after a few quarters. This is in line with a margin for intertemporal substitution in production, as in our theory.

This is all good news for our theory. It also indicates the potential value of quantifying our theory, a task we leave for future work. What, however, the above facts do not directly speak to is the magnitude of the confidence multiplier.

Ultimately, the sign and size of the kind of mis-perceptions—rational or irrational—modeled in this paper is an empirical question. We leave this question for future work. But by explaining both the macroeconomic implications and the precise nature of the relevant mis-perceptions, we hope to have provided useful guidance for what exactly future empirical work should explore.

Let us expand on this last point by relating to the empirical findings of [Rozsypal and Schlafmann \(2018\)](#), [Gennaioli, Ma, and Shleifer \(2016\)](#) and [Greenwood and Shleifer \(2014\)](#). These works are supportive of our confidence mechanism in the sense that they point out in the direction of expectations of income and returns being excessively optimistic in good times and excessively pessimistic at bad times.⁴³ They do not, however, distinguish on whether such good and bad times are driven by the kind of aggregate demand shocks that are the focus of our paper or by other forces, such as TFP shocks. The litmus test of our theory is therefore *conditional* evidence for how expectations of income and interest rates responds to different kinds of shocks—and it is this specific kind of evidence we invite for future work.

⁴²The IRF of the real interest rate can be readily inferred from the figure by subtracting the nearly flat IRF of inflation from the more strongly procyclical IRF of the nominal interest rate.

⁴³[Rozsypal and Schlafmann \(2018\)](#) use micro data on household income expectations, and find that consumers over-extrapolate from their current income to expectations of future income, as they overestimate the persistence of their income process. Households with currently high income turn out to be too optimistic about their future income, while households with currently low income turn out to be too pessimistic about their future income. [Gennaioli, Ma, and Shleifer \(2016\)](#) turn the focus to the firm side. They find that firm CFOs over-extrapolate from current earnings to expectations of future earnings. [Greenwood and Shleifer \(2014\)](#) find that investors over-extrapolate from current stock returns to expectations of future stock returns.

References

- Abel, Andrew B. and Olivier J. Blanchard. 1983. "An Intertemporal Model of Saving and Investment." *Econometrica* 51 (3):675–692.
- Andrade, Philippe, Olivier Coibion, Erwan Gautier, and Yuriy Gorodnichenko. 2020. "No Firm is an Island? How Industry Conditions Shape Firms' Aggregate Expectations." *NBER Working Paper No. 27317* .
- Angeletos, George-Marios, Fabrice Collard, and Harris Dellas. 2018. "Quantifying Confidence." *Econometrica* 86 (5):1689–1726.
- . 2020. "Business Cycle Anatomy." *American Economic Review* forthcoming.
- Angeletos, George-Marios and Jennifer La'O. 2013. "Sentiments." *Econometrica* 81 (2):739–779.
- Angeletos, George-Marios and Chen Lian. 2018. "Forward Guidance without Common Knowledge." *American Economic Review* 108 (9):2477–2512.
- Angeletos, George-Marios, Guido Lorenzoni, and Alessandro Pavan. 2010. "Beauty contests and irrational exuberance: A neoclassical approach." *NBER Working Paper No. 15883* .
- Bai, Yan, José-Victor Rios-Rull, and Kjetil Storesletten. 2017. "Demand Shocks as Productivity Shocks." *University of Pennsylvania mimeo* .
- Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer. 2015. "X-CAPM: An extrapolative capital asset pricing model." *Journal of Financial Economics* 115 (1):1–24.
- Barsky, Robert B and Eric R Sims. 2011. "News shocks and business cycles." *Journal of Monetary Economics* 58 (3):273–289.
- Basu, Susanto, John G Fernald, and Miles S Kimball. 2006. "Are technology improvements contractionary?" *American Economic Review* 96 (5):1418–1448.
- Beaudry, Paul and Franck Portier. 2006. "Stock Prices, News, and Economic Fluctuations." *American Economic Review* 96 (4):1293–1307.
- . 2013. "Understanding Noninflationary Demand-Driven Business Cycles." In *NBER Macroeconomics Annual 2013, Volume 28*. 69–130.
- . 2018. "Real Keynesian Models and Sticky Prices." *NBER Working Paper No. 24223* .
- Benhabib, Jess and Roger EA Farmer. 1999. "Indeterminacy and sunspots in macroeconomics." *Handbook of Macroeconomics* 1:387–448.

- Benhabib, Jess, Pengfei Wang, and Yi Wen. 2015. "Sentiments and Aggregate Demand Fluctuations." *Econometrica* 83 (2):549–585.
- Blanchard, Olivier J. and Danny Quah. 1989. "The Dynamic Effects Of Aggregate Demand And Supply Disturbances." *American Economic Review* 79 (4):655–673.
- Cavallo, Alberto, Guillermo Cruces, and Ricardo Perez-Truglia. 2017. "Inflation Expectations, Learning, and Supermarket Prices: Evidence from Survey Experiments." *American Economic Journal: Macroeconomics* 9 (3):1–35.
- Chahrour, Ryan and Gaetano Gaballo. 2018. "Learning from Prices: Amplification and Business Fluctuations." *Boston College/ECB mimeo* .
- Christiano, Lawrence, Roberto Motto, and Massimo Rostagno. 2014. "Risk shocks." *American Economic Review* 104 (1):27–65.
- Coibion, Olivier and Yuriy Gorodnichenko. 2012. "What Can Survey Forecasts Tell Us about Information Rigidities?" *Journal of Political Economy* 120 (1):116–159.
- . 2015. "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts." *American Economic Review* 105 (8):2644–78.
- Coibion, Olivier, Yuriy Gorodnichenko, and Tiziano Ropele. 2018. "Inflation Expectations and Firm Decisions: New Causal Evidence." *NBER Working Paper No. 25412* .
- Cooper, Russell and Andrew John. 1988. "Coordinating coordination failures in Keynesian models." *Quarterly Journal of Economics* :441–463.
- Correia, Isabel, Juan Pablo Nicolini, and Pedro Teles. 2008. "Optimal Fiscal and Monetary Policy: Equivalence Results." *Journal of Political Economy* 116 (1):141–170.
- Diamond, Peter A. 1982. "Aggregate Demand Management in Search Equilibrium." *The Journal of Political Economy* :881–894.
- Eggertsson, Gauti B and Paul Krugman. 2012. "Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach." *Quarterly Journal of Economics* 127 (3):1469–1513.
- Farhi, Emmanuel and Iván Werning. 2019. "Monetary Policy, Bounded Rationality, and Incomplete Markets." *American Economic Review* 109 (11):3887–3928.
- Gabaix, Xavier. 2019. *Handbook of Behavioral Economics* :261 – 343.
- . 2020. "A Behavioral New Keynesian Model." *American Economic Review* forthcoming.

- Gaballo, Gaetano. 2017. "Price Dispersion, Private Uncertainty, and Endogenous Nominal Rigidities." *The Review of Economic Studies* 85 (2):1070–1110.
- Gali, Jordi. 1999. "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review* 89 (1):249–271.
- Garcia-Schmidt, Mariana and Michael Woodford. 2019. "Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis." *American Economic Review* 109 (1):86–120.
- Gennaioli, Nicola, Yueran Ma, and Andrei Shleifer. 2016. "Expectations and Investment." *NBER Macroeconomics Annual* 30 (1):379–431.
- Greenwood, Jeremy, Zvi Hercowitz, and Gregory W Huffman. 1988. "Investment, Capacity Utilization, and the Real Business Cycle." *The American Economic Review* :402–417.
- Greenwood, Robin and Andrei Shleifer. 2014. "Expectations of returns and expected returns." *The Review of Financial Studies* 27 (3):714–746.
- Grossman, Sanford J and Joseph E Stiglitz. 1980. "On the Impossibility of Informationally Efficient Markets." *The American Economic Review* 70 (3):393–408.
- Guerrieri, Veronica and Guido Lorenzoni. 2017. "Credit Crises, Precautionary Savings, and the Liquidity Trap." *Quarterly Journal of Economics* forthcoming.
- Hall, Robert E. 2011. "The Long Slump." *American Economic Review* 101 (2):431–69.
- Huo, Zhen and Jose-Victor Rios-Rull. 2020. "Demand Induced Fluctuations." *Review of Economic Dynamics* .
- Ilut, Cosmin and Hikaru Saijo. 2018. "Learning, Confidence and Business Cycle." miméo, Duke University.
- Iovino, Luigi and Dmitriy Sergeyev. 2017. "Quantitative Easing without Rational Expectations." *Work in progress* .
- Justiniano, Alejandro, Giorgio Primiceri, and Andrea Tambalotti. 2010. "Investment Shocks and Business Cycles." *Journal of Monetary Economics* 57 (2):132–145.
- Lorenzoni, Guido. 2009. "A Theory of Demand Shocks." *American Economic Review* 99 (5):2050–84.
- Lucas, Robert E. 1972. "Expectations and the Neutrality of Money." *Journal of Economic Theory* 4 (2):103–124.

- . 1973. “Some International Evidence on Output-Inflation Tradeoffs.” *American Economic Review* 63 (3):326–334.
- Mackowiak, Bartosz and Mirko Wiederholt. 2009. “Optimal Sticky Prices under Rational Inattention.” *American Economic Review* 99 (3):769–803.
- Mavroeidis, Sophocles, Mikkel Plagborg-Møller, and James H. Stock. 2014. “Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve.” *Journal of Economic Literature* 52 (1):124–88.
- McLeay, Michael and Silvana Tenreyro. 2020. “Optimal inflation and the identification of the Phillips curve.” *NBER Macroeconomics Annual* 34 (1):199–255.
- Mian, Atif and Amir Sufi. 2014. “What Explains the 2007–2009 Drop in Employment?” *Econometrica* 82 (6):2197–2223.
- Molavi, Pooya. 2019. “Macroeconomics with Learning and Misspecification: A General Theory and Applications.” mimeo, MIT.
- Ravn, Morten, Stephanie Schmitt-Grohe, and Martin Uribe. 2006. “Deep Habits.” *Review of Economic Studies* 73 (1):195–218.
- Rozsypal, Filip and Kathrin Schlafmann. 2018. “Overpersistence Bias in Individual Income Expectations and its Aggregate Implications.” *IIES mimeo* .
- Sims, Christopher A. 2003. “Implications of Rational Inattention.” *Journal of Monetary Economics* 50 (3):665–690.
- Vives, Xavier and Liyan Yang. 2018. “Costly Interpretation of Asset Prices.” mimeo, IESE/University of Toronto mimeo.
- Woodford, Michael. 2011. “Simple Analytics of the Government Expenditure Multiplier.” *American Economic Journal: Macroeconomics* 3 (1):1–35.

Appendix: Proofs

[TBC]