Faraway, So Close: Business Cycle Effect of Long-Run Ambiguity

Sara Biadetti* Lorenzo Carbonari† Filippo Maurici‡

Abstract

This paper examines the effects of forward-looking ambiguity, also known as Knightian uncertainty, in a model with homogeneous workers and credit-constrained heterogeneous entrepreneurs. Agents are ambiguity-averse and employ a worst-case criterion to formulate expectations about (future) total factor productivity. By comparing our economy with one that has the same fundamentals but lacks uncertainty, we observe that ambiguity: (i) reduces the productivity threshold to access the market, (ii) pushes the equilibrium interest rate below the workers' discount rate, and (iii) increases workers' consumption while reducing entrepreneurs' consumption. These consequences arise from the persistent misalignment between expectations and realizations, which, combined with ambiguity aversion, induces a hedging strategy among workers.

Keywords: ambiguity, collateral constraints, heterogeneous agents, transition dynamics.

JEL codes: E22, D81, D84, G14.

We thank Katia Colaneri, Gaetano Bloise, Jacopo Orlandi, Guillaume Pommey, and Federico Giorgi for their comments and suggestions.

*Università di Roma "Tor Vergata".

†Università di Roma "Tor Vergata", DEF and CEIS. E-mail address: lorenzo.carbonari@uniroma2.it.

‡Università di Roma "Tor Vergata". E-mail address: filippo.maurici@uniroma2.it.

Count not that far that can be had,
Though sunset lie between –
Nor that adjacent, that beside,
Is further than the sun.
Emily Dickinson.

1 Introduction

This paper proposes a business cycle theory in which exogenous ambiguity shocks drive movements in macroeconomic variables. Agents are ambiguity-averse in the sense that they dislike uncertainty and unquantifiable knowledge about future events. As a result, they make decisions based on worst-case probabilities derived from a spectrum of possible outcomes. We explore how ambiguity interacts with borrowing constraints in a credit constrained economy, and how such effects are amplified by heterogeneity. Specifically, we integrate the concept of ambiguity as described by Epstein and Schneider (2007), Ilut and Schneider (2014), and Bidder and Dew-Becker (2016) into the analytical framework proposed by Itskhoki and Moll (2019).

We consider an economy populated by two types of infinite-lived agents, workers and entrepreneurs. Workers, who are are homogeneous, consume, save by holding corporate bonds, and decide how much labor to supply. Entrepreneurs, who are heterogeneous in terms of assets and productivity, face borrowing constraints that limit the amount of capital they can use. Depending on the realization of a random productivity shock, they decide whether to engage in production (being *active*) or exit the market (being *inactive*). They also consume and accumulate wealth. Due to heterogeneity and borrowing constraints, entrepreneurs make positive profits.

Both workers and entrepreneurs perceive *Knightian uncertainty* and are ambiguity-averse. We assume that they employ a MaxMin expected utility criterion (Gilboa and Schmeidler, 1989). Ambiguity pertains to the future levels of total factor productivity, although its present value is known with certainty. The interaction between certain knowledge of the current TFP and forward-looking ambiguity on its future realizations drives the economy away from the steady state that would be achieved under perfect foresight. By comparing our *ambiguous* economy to the one with perfect foresight we observe the following. To hedge against the worst-case scenario, workers increase their savings, generating an excess capital supply. This leads to a lower long-run equilibrium interest rate. Interestingly, this form of ambiguity makes workers better off while entrepreneurs are worse off. Lastly, we observe that a higher heterogeneity implies a more significant and persistent effects of ambiguity on macroeconomic variables.

To sum up, the paper provides two main contributions. First, it investigates the role of forward-looking ambiguity in business cycle fluctuations (i.e., it proposes a theory on how future ambiguity influences current saving decisions and interacts with back-propagation effects). An ambiguity shock triggers procyclical movements in the interest rate, aggregate wealth, and workers' consumption. Second, it sheds light on the difference between short-term fluctuations driven by resolving ambiguity shocks and long-term trajectories shaped by persistent ambiguity, steering the economy away from the perfect-foresight steady state.

Related literature We build upon the models presented in Buera and Moll (2015) and Itskhoki and Moll (2019). Buera and Moll (2015) examines the aggregate implications of a credit crunch under various specifications of heterogeneity, highlighting the critical role of heterogeneity in understanding the sources of business cycles. Itskhoki and Moll (2019) sets the same model in continuous time to study optimal development policies and dynamic taxation. Aside from incorporating uncertainty, our economy shares the same fundamentals as the one analyzed in Itskhoki and Moll (2019) in its closed-economy version, which can serve as a benchmark. A similar setup is also found in Banerjee and Moll (2010), which introduces the concept of distinct types of capital misallocation – intensive and extensive – distinguishing the quantity of capital from the quality of its usage and attributing persistent inefficiencies to binding financial constraints and heterogeneity.

Modeling heterogeneity among entrepreneurs as a productivity shock with a Pareto distribution aligns with recent literature – see, e.g., Buera and Moll (2015), Gabaix et al. (2016), Jones and Kim (2018), Itskhoki and Moll (2019), and Achdou et al. (2022) – and empirical observations, which describe the right tail of the wealth distribution as following a Pareto distribution. This approach captures the extreme variability in productivity among entrepreneurs, highlighting that only those with sufficiently high productivity levels engage in production at the aggregate level.

We model ambiguity following Epstein and Schneider (2003), Epstein and Schneider (2007), and Ilut and Schneider (2014). Epstein and Schneider (2007) analyzes an economy where agents' confidence evolves as they acquire new information. Similarly, Ilut and Schneider (2014) examines a New Keynesian business cycle model that incorporates confidence shocks, demonstrating how variations in ambiguity can generate business cycles. Unlike these studies, we do not rely on nominal rigidities. We explore the interplay between ambiguity and heterogeneity, setting the economy within a continuous-time framework and specifically focus on forward-looking ambiguity. This forward-looking approach is similar to that of Bidder and Dew-Becker (2016), where long-run ambiguity influences investors' asset pricing decisions.

In a small open economy with credit constraints, entrepreneurial heterogeneity, and *hand-to-mouth* workers, Carbonari and Maurici (2023) shows that ambiguity has both short- and long-run implications. Similar to what happens in our economy, agents' hedging strategies are at the root of short-run fluctuations. However, the underlying mechanism is entirely different. In their paper, workers play a residual role, and the entrepreneurs' hedging strategy allows them to take advantage of aggregate uncertainty. Additionally, in the long run, ambiguity widens the consumption inequality between agents. In contrast, in our model, workers can accumulate assets, and ambiguity leads to higher steady state consumption levels for them.

Other recent models with ambiguity that have examined business cycle phenomena are Nimark (2014), Altug et al. (2020), and Michelacci and Paciello (2024). Nimark (2014) presents a business cycle model with higher-order beliefs and considers the impact of signals observed after unusual events that increase uncertainty and disagreement among agents. Using a smooth ambiguity model (Klibanoff et al., 2005, 2009), Altug et al. (2020) studies the business cycle implications of ambiguity aversion in a model in a standard neoclassical model with investment irreversibility. Michelacci and Paciello (2024) proposes a New Keynesian model with ambiguity averse agents and associate their worst-case beliefs to the credibility of the monetary authority actions.

Classic references on ambiguity, robustness and model misspecification are Bewley (2002), Hansen and Sargent (2011), and Hansen and Sargent (2022). Chen and Epstein (2002), Gagliardini et al. (2008), Epstein and Ji (2013), Lin and Riedel (2021), and Hansen and Miao (2022) explore ambiguity in continuous-time settings within the context of asset pricing.

Outline The paper is structured as follows. In Section 2, we lay out the structure of the economy and derive the optimal choices, both individual and aggregate, focusing on consumption and investment decisions. In Section 3, we characterize the competitive equilibrium under forward-looking ambiguity. In Section 4, we study the business cycle dynamics. Section 5 concludes.

2 Model

Consider an infinite-horizon closed economy operating in continuous time. As in Itskhoki and Moll (2019), there are two types of agents, each with unit mass. On one side, identical *workers* consume (c) and work during each period. They supply a variable quantity of labor hours (ℓ) and may save a fraction of their income by holding an asset (b). On the other side, heterogeneous *entrepreneurs*, differing in both their productivity (z) and asset holdings (a), run the production process. At each date, entrepreneurs hire workers at a competitive wage w and rent capital at rate r. Firms combine labor (n) and capital (k) to produce a final good, which serves as the numeraire and is traded in a competitive market. In what follows, let $t^+ = \lim_{\tau \to t^+} \tau$ indicate the *next period* in a continuous setting, meaning that $t^+ \mid t$ should be read as the next-period-ahead forecast. For notational convenience, we will omit time subscripts unless they are necessary.

Ambiguity Workers and entrepreneurs are affected by *Knightian uncertainty* and are ambiguity-averse. We assume the existence of market-wide information regarding the set of perceived distributions ($\mathcal{P}_{t^+|t}$) of *instantaneous* future TFP, based on the information available at time t.[1] Hence, ambiguity is represented by all possible elements $P \in \mathcal{P}_{t^+|t}$. Ambiguity influences choices through expectations about the aggregate total factor productivity (TFP). We model ambiguity aversion by considering a utility function which is the minimum of the expected utilities over the collection $\mathcal{P}_{t^+|t}$.

We denote the *true* TFP series as $\{A_\tau\}_{\tau \geq t}$. As in Ilut and Schneider (2014), the stochastic process governing the TFP is unknown to the agents, who face ambiguity regarding its future realizations.

We now lay out the informational structure of the economy. For any t and any $P \in \mathcal{P}_{t^+|t}$, we define the *perceived* TFP as the random variable $A^P_{t^+|t}$ defined on a probability space $(\Omega^P, \mathscr{F}^P, P)$. For a given P, the σ-algebra $\mathscr{F}^P$ conveys the information available at time t. Let $\mathcal{A}_t$ be the set of all points that have non-zero probability under at least one $P \in \mathcal{P}_{t^+|t}$,

$$\mathcal{A}_t = \bigcup_{P \in \mathcal{P}_{t^+|t}} \operatorname{supp}(P) \subseteq \mathbb{R},$$

and $\underline{A}_t \in \mathbb{R}$ be

$$\underline{A}_t = \inf \mathcal{A}_t.$$

We assume that, in any period, all the possible outcomes are bounded from below by $\underline{A}_t$,

$$P\left(A^P_{t^+|t} \geq \underline{A}_t\right) = 1,$$

for all $P \in \mathcal{P}_{t^+|t}$, with $\underline{A}_t > 0$, and that the Dirac-delta distribution centered in $\underline{A}_t$ is an element of the set of perceived distributions,[2]

$$\mathcal{D}\left(\underline{A}_t\right) \in \mathcal{P}_{t^+|t}.$$

The former assumption is the continuous-time adaptation of Ilut and Schneider (2014), denoting the "*indicator of the quality of intangible information available at date t*

[1] This assumption can be seen as a limiting case of De Long et al. (1990).

[2] Since our notation is the same as Itskhoki and Moll (2019), we denote the Dirac-delta distribution with $\mathcal{D}$ to avoid confusion with the entrepreneurs' discount rate, δ, which will be defined below.

about TFP at $t+1$".[3] The latter simplifies the analysis by restricting the focus on $\mathcal{D}(\underline{A}_t)$ within $\mathcal{P}_{t^+|t}$. Specifically, for the Dirac-delta distribution, the upper and lower bounds of the expected TFP coincide at $\underline{A}_t$, whereas for all other elements in $\mathcal{P}_{t^+|t}$, the upper bound must be strictly higher than $\underline{A}_t$.[4] Consequently, ambiguity-averse agents will exclusively focus on $\mathcal{D}(\underline{A}_t)$, meaning the Dirac-delta distribution represents the worst-case scenario.

Workers and entrepreneurs behave as if they are optimizing based on a worst-case distribution, which is different from actual data-generating process.

For the sake of exposition, we start by setting each element of the sequence $\{A_\tau\}_{\tau\geq t}$ equal to A, and each perceived TFP lower bound $\{\underline{A}_\tau\}_{\tau\geq t}$ constant at $\underline{A}$, where $A > \underline{A}$. In Section 4, the lower bound is treated as time-varying to analyze the business cycle.

Workers The representative worker's preferences are given by

$$\min_{P\in\mathcal{P}_{t^+|t}} \mathbb{E}_P \int_t^\infty e^{-\rho(\tau-t)} u(c_\tau, \ell_\tau)\,\mathrm{d}\tau, \tag{1}$$

where ρ is the workers' discount rate, and

$$u(c,\ell) = \ln c - \psi\frac{\ell^{1+\phi}}{1+\phi},$$

with $\psi, \phi > 0$. The worker is subject to the flow budget constraint

$$c + \dot{b} = rb + w\ell. \tag{2}$$

The resulting optimality conditions are

$$-\frac{u_\ell}{u_c} = \psi c\ell^\phi = w, \tag{3}$$

$$\frac{\dot{u}_c}{u_c} = \frac{\dot{c}}{c} = \rho - r, \tag{4}$$

along with the proper transversality condition.

Entrepreneurs Assuming a logarithmic instantaneous utility function, an entrepreneur maximizes

$$\min_{P\in\mathcal{P}_{t^+|t}} \mathbb{E}_P \int_t^\infty e^{-\delta(\tau-t)} \ln c^e_\tau\,\mathrm{d}\tau, \tag{5}$$

[3] See Ilut and Schneider (2014, p. 2375). Similarly, Carli (2024) examines ambiguity in an environmental two-sector economy.

[4] This ensures that, by convexity, for all $P\in\mathcal{P}_{t^+|t}\backslash\mathcal{D}(\underline{A}_t)$ we have $\underline{A}_t < \mathbb{E}_P\left[A^P_{t^+|t}\right]$.

where δ is the entrepreneurs' discount rate, subject to the flow budget constraint

$$c^e + \dot{a} = \pi(a, z) + ra. \tag{6}$$

Given continuous time and logarithmic preferences, the optimal individual consumption for each entrepreneur is $c^e = \delta a$. Equation (6) characterizes the dynamics of entrepreneurial wealth, leading to a time-dependent distribution of assets, denoted as $\mathcal{G}_t(a)$.

We assume that idiosyncratic productivity follows a Pareto distribution with a shape parameter $\eta > 1$.[5] Productivity shocks are i.i.d. over time. The distribution function $\mathcal{F}(z)$ is supported on the interval $[1, \infty)$, and its density function is given by $\eta z^{-\eta-1}$.

Financial markets are incomplete, in that entrepreneurs face a collateral constraint given by

$$k \leq \lambda a, \tag{7}$$

with $\lambda \geq 1$.[6]

Upon making production decisions A is known, and the entrepreneur's production function can be expressed as

$$A(zk)^{\alpha} n^{1-\alpha}, \tag{8}$$

with $\alpha \in (0, 1)$, while profits are

$$A(zk)^{\alpha} n^{1-\alpha} - wn - rk.$$

Profit maximization gives

$$k(a, z) = \lambda a \, \mathbb{I}\left(z \geq \underline{z}\right), \tag{9}$$

$$n(a, z) = \left[\frac{1-\alpha}{w} A\right]^{\frac{1}{\alpha}} zk(a, z), \tag{10}$$

$$\pi(a, z) = \alpha \left[\frac{1-\alpha}{w}\right]^{\frac{1-\alpha}{\alpha}} A^{\frac{1}{\alpha}} zk(a, z) - rk(a, z), \tag{11}$$

[5] The shape parameter η determines the tail heaviness of the distribution. A higher η results in a thinner tail, indicating that extreme values are less common (i.e., the distribution of idiosyncratic productivity exhibits lower heterogeneity). As highlighted above, the motivation for this assumption is empirical, since the right tail of the wealth distribution in many countries is Pareto-shaped (Jones and Kim, 2018). Additionally, assuming a Pareto distribution simplifies the derivation of aggregate properties of our model economy.

[6] This assumption is equivalent to imposing a debt limit on entrepreneurs. An entrepreneur can leverage debt, d, backed by a fraction ϑ of the acquired capital, k, which acts as collateral (i.e., $d \leq \vartheta k$). ϑ can be interpreted as the loan-to-value ratio (LTV). Entrepreneur's net worth is then given by $a \equiv k - d$, and the debt constraint can be reformulated as (7), where $\lambda \equiv (1 - \vartheta)^{-1}$. Moreover, as $\lim_{\vartheta \to 1^-} \lambda = \infty$ it signals the efficient functioning of capital markets.

together with the appropriate transversality condition, where $\underline{z}$ is the minimum level of idiosyncratic productivity to enter the market. Let $\zeta \equiv \frac{r}{\alpha}\left[\frac{w}{1-\alpha}\right]^{\frac{1-\alpha}{\alpha}} A^{-\frac{1}{\alpha}}$,[7] then

$$\underline{z} = \max\left\{\zeta, 1\right\}. \tag{12}$$

Condition (12) states that the cutoff level of z is the maximum between the zero-profit productivity, ζ, and the lower bound of the distribution $\mathcal{F}(z)$, 1. For the sake of simplicity, we restrict our attention to the parameter space (and initial conditions) such that $\underline{z} = \zeta > 1$.

Aggregation The evolution of $\mathcal{G}_t(a)$ over time does not pose a challenge, as (6) can be expressed as a linear function of a. Consequently, it becomes feasible to monitor the evolution of aggregate entrepreneurial wealth.

Let x represent the aggregate level of entrepreneurs' assets, specifically as $x = \int a \, \mathrm{d}i$, then $\int c^e \, \mathrm{d}i = \delta x$.

The aggregate levels of capital and labor demand are given by

$$\kappa = \lambda \underline{z}^{-\eta} x, \tag{13}$$

$$\ell = \lambda \frac{\eta}{\eta - 1}\left[\frac{1-\alpha}{w} A\right]^{\frac{1}{\alpha}} \underline{z}^{1-\eta} x. \tag{14}$$

Factor shares are given by

$$w\ell = (1-\alpha) y, \tag{15}$$

$$r\kappa = \frac{\alpha(\eta - 1)}{\eta} y, \tag{16}$$

with aggregate profits equal to $\frac{\alpha}{\eta} y$. The aggregate production is

$$y = \Theta A x^{\frac{\alpha}{\eta}} \kappa^{\frac{\alpha(\eta-1)}{\eta}} \ell^{1-\alpha}, \tag{17}$$

with $\Theta \equiv \left(\frac{\eta}{\eta-1}\right)^{\alpha} \lambda^{\frac{\alpha}{\eta}}$, while the law of motion of the entrepreneurs' wealth is

$$\dot{x} = \frac{\alpha}{\eta} y + (r - \delta) x. \tag{18}$$

Finally, the feasibility constraint is given by

$$\dot{\kappa} = y - \delta x - c. \tag{19}$$

[7] Derived by setting (11) to zero and solving for z.

3 Equilibrium under ambiguity

Given the initial aggregate levels for entrepreneurial wealth x_0 and workers' savings b_0, actual TFP A, and the worst-case distribution $\mathcal{D}(\underline{A})$, a competitive equilibrium consists of a sequence of aggregate variables $\{x, b, \kappa, y, c, \ell, \underline{z}\}_{\tau\geq t}$, wage $\{w\}_{\tau\geq t}$, and interest rate $\{r\}_{\tau\geq t}$ such that (i) workers maximize (1) subject to (2), taking prices as given; (ii) entrepreneurs maximize (5) subject to (6), taking prices as given; (iii) market clearing conditions hold; (iv) entrepreneurial wealth evolves according to (18); (v) aggregate wealth evolves according to (19). A steady state equilibrium is a competitive equilibrium in which real variables are time-invariant.

Ambiguous steady state In this model, the ambiguity relates to the future realization of the aggregate TFP. Consider now the following thought experiment. Based on the agents' expectation that the *true* level of TFP is $\underline{A}$, we can compute a worst-case steady state equilibrium, which we term *ambiguous*. Since $A > \underline{A}$, the economy never actually reaches this allocation. However, due to pessimistic expectations, agents believe the opposite and make their decisions accordingly. Let $\widehat{s}$ denote the steady state value for any variable s. *If* the true realization of the TFP were $\underline{A}$, we would have

$$\widehat{r} = \rho, \tag{20}$$

$$\widehat{c} = \left(1 - \frac{\delta}{\delta-\rho}\frac{\alpha}{\eta}\right)\widehat{y}, \tag{21}$$

$$\widehat{x} = \frac{1}{\delta-\rho}\frac{\alpha}{\eta}\widehat{y}, \tag{22}$$

$$\widehat{\kappa} = \frac{\alpha}{\rho}\frac{\eta-1}{\eta}\widehat{y}, \tag{23}$$

$$\widehat{b} = \left(\alpha - \frac{\delta}{\delta-\rho}\frac{\alpha}{\eta}\right)\frac{\widehat{y}}{\rho}, \tag{24}$$

$$\widehat{\ell} = \left[\frac{1-\alpha}{\psi}\left(1 - \frac{\delta}{\delta-\rho}\frac{\alpha}{\eta}\right)^{-1}\right]^{\frac{1}{1+\phi}}, \tag{25}$$

$$\widehat{w} = (1-\alpha)\frac{\widehat{y}}{\widehat{\ell}}, \tag{26}$$

$$\widehat{y} = \Theta\underline{A}\widehat{x}^{\frac{\alpha}{\eta}}\widehat{\kappa}^{\frac{\alpha(\eta-1)}{\eta}}\widehat{\ell}^{1-\alpha}. \tag{27}$$

Notice that (21) gives the necessary condition $(\delta-\rho)\eta > \delta\alpha$. Moreover, to have $\widehat{b} > 0$, from (24), we must have $(\delta-\rho)\eta > \delta$. Lastly, combining (22) and (23) with (13), we get a necessary condition for $\widehat{\underline{z}} > 1$, that is, $\rho\lambda > (\delta-\rho)(\eta-1)$.

The system of equations (20)-(27) represents the steady state of an economy conditional on $\mathcal{D}(\underline{A})$, with $A > \underline{A}$ (the lowest lower bound among all distributions within

$\mathcal{P}_{t^+|t}$). This instance, however, never occurs, since $A > \underline{A}$ by assumption. This implies that, under ambiguity, the economy never reaches such a long-run equilibrium. Nonetheless, as we will show, these conditions are essential to characterize the behavior of our economy. Indeed, agents systematically underestimate the actual value of the TFP, and the equilibrium dynamics arises from the interplay between *ex ante* pessimism and actual relizations. It then turns out that the ambiguous steady state, which is attained under the worst belief, is the point around which to linearize the behavior of the economy, since agents act *as if* it converges there in the long run (Ilut and Schneider, 2014).

It is convenient to express all endogenous variables in (4), (18), and (19) as functions of c, x, and κ. Equating (3) and (15) at the expected steady state, ambiguous aggregate production can be rewritten as

$$\widehat{y} = \Omega \underline{A}^{\frac{\omega}{\alpha}} \widehat{x}^{\frac{\omega}{\eta}} \widehat{\kappa}^{\frac{\omega(\eta-1)}{\eta}} \widehat{c}^{1-\frac{\omega}{\alpha}}, \tag{28}$$

where $\omega \equiv \alpha\frac{1+\phi}{\alpha+\phi}$ and $\Omega \equiv \left[\frac{1-\alpha}{\psi}\right]^{\frac{1-\alpha}{\alpha+\phi}} \Theta^{\frac{\omega}{\alpha}}$. Similarly, $\widehat{r}$ is obtained via (16).

The dynamics of the economy, linearized around the ambiguous steady state, is described by the following system

$$\begin{bmatrix} \dot{c} \\ \dot{x} \\ \dot{\kappa} \end{bmatrix} = \mathbf{J} \begin{bmatrix} c - \widehat{c} \\ x - \widehat{x} \\ \kappa - \widehat{\kappa} \end{bmatrix}, \tag{29}$$

where the $\mathbf{J}$ is the Jacobian matrix and is computed in the Appendix.

To ensure stability, $\mathbf{J}$ must have one eigenvalue with a positive real part, associated with the forward-looking variable, c, and two eigenvalues with negative real parts, associated with the predetermined variables, x and κ. We restrict our analysis to the parameter space that ensures that these conditions are satisfied. Therefore, from the stable branch of (29), consumption becomes[8]

$$c = \widehat{c} - \frac{Q_x(x - \widehat{x}) + Q_\kappa(\kappa - \widehat{\kappa})}{Q_c}. \tag{30}$$

Once c is pinned down, given that x and κ are predetermined, and $\widehat{c}$, $\widehat{x}$, and $\widehat{\kappa}$ are known, the model is solved. Notice that (30) depends on worst-case expectation $\underline{A}$ at the steady state, through (28). Moreover, since $\widehat{c}$ is increasing in $\underline{A}$, a decrease in ambiguity (i.e., a higher $\underline{A}$) implies a higher workers' consumption in the ambiguous steady state.

The tension between long-term pessimistic expectations and $A > \underline{A}$ produces two countervailing effects. In the short run, it leads to a reduction in workers' consumption as

[8] Q_c, Q_x, and Q_κ are the entries of the left-eigenvector associated to the eigenvalue with positive real part. See the Appendix for the derivation.

agents increase their savings to prepare for future expected downturns – i.e., ambiguity induces higher savings due to precautionary motive like in Aiyagari (1994). Conversely, in the long run, the economy converges to a steady state where the stream of workers' excessive savings allows them to attain a consumption which is higher than that they would achieve in a situation without uncertainty about the realization of TFP. To illustrate this point, we compare in Figure 1 the equilibrium dynamics of our economy (blue line), in which the actual TFP stands at A while the perceived future TFP is $\underline{A}$, with that of an economy with the same fundamentals but without ambiguity (black line), as in Itskhoki and Moll (2019).[9]

Along the convergence towards the steady state, compared to the case with no uncertainty, precautionary motive *initially* leads to a higher labor supply, as workers are willing to work more and accumulate more bonds – i.e., workers consume less (1e) to bolster savings (1b), offering higher labor hours at lower wages, as shown in Figure 1g and Figure 1h. Subsequently, workers' savings fuel aggregate wealth (1c) and drive higher wages, thus contracting labor demand (1h) and reducing entrepreneurial wealth (1a), which in turn lead to a lower interest rate (1d) and a lower productivity cutoff (1f).[10] Accordingly, aggregate production goes up (1i). Lastly, regarding consumption (1e), we notice that once the level of savings becomes sufficiently high, it surpasses the level attained in an economy without uncertainty.[11] Notice that the over-accumulation of assets, compared to the case of no ambiguity, also endures in the long run, characterized by a persistent mismatch between expectations and realizations.

We show how the interplay between ambiguous expectations and realizations can be captured by constructing vector fields induced by ambiguity. In each period, ambiguity generates two distinct vector fields: one *ex ante*, based on current forward-looking ambiguity, and one *ex post*, reflecting the actual realization of the TFP. The dynamics of workers' consumption implied by the vector fields is shown in Figure 6 (in the Appendix), while Figure 2 represents their evolution over time. In each figure, the orange line denotes the *ex ante* vector field as $\dot{c}(A, \underline{A})$, while the light blue line denotes the *ex post* vector field as $\dot{c}(A, A)$. In both cases, the first entry, A, specifies the current, known, TFP, while the second entry represents the expected ambiguous TFP, $\underline{A}$, and the realized TFP, A, respectively.

Pessimistic expectations systematically underestimate the actual realization of aggregate TFP, thus inducing an acceleration in consumption. The growth rate of

[9]We use the parameters from the calibration in Section 4. However, in this section, we set the variances to zero.

[10]The fact that $r \neq \rho$ at the steady state is similar to Moll et al. (2022). However, in their case, the presence of dissipation shocks, which leave agents with zero assets and only their labor income, causes $r > \rho$.

[11]In Carbonari and Maurici (2023), where workers are *hand-to-mouth*, the effect of ambiguity on workers' consumption is the opposite, resulting in a lower steady-state level compared to the case without uncertainty.

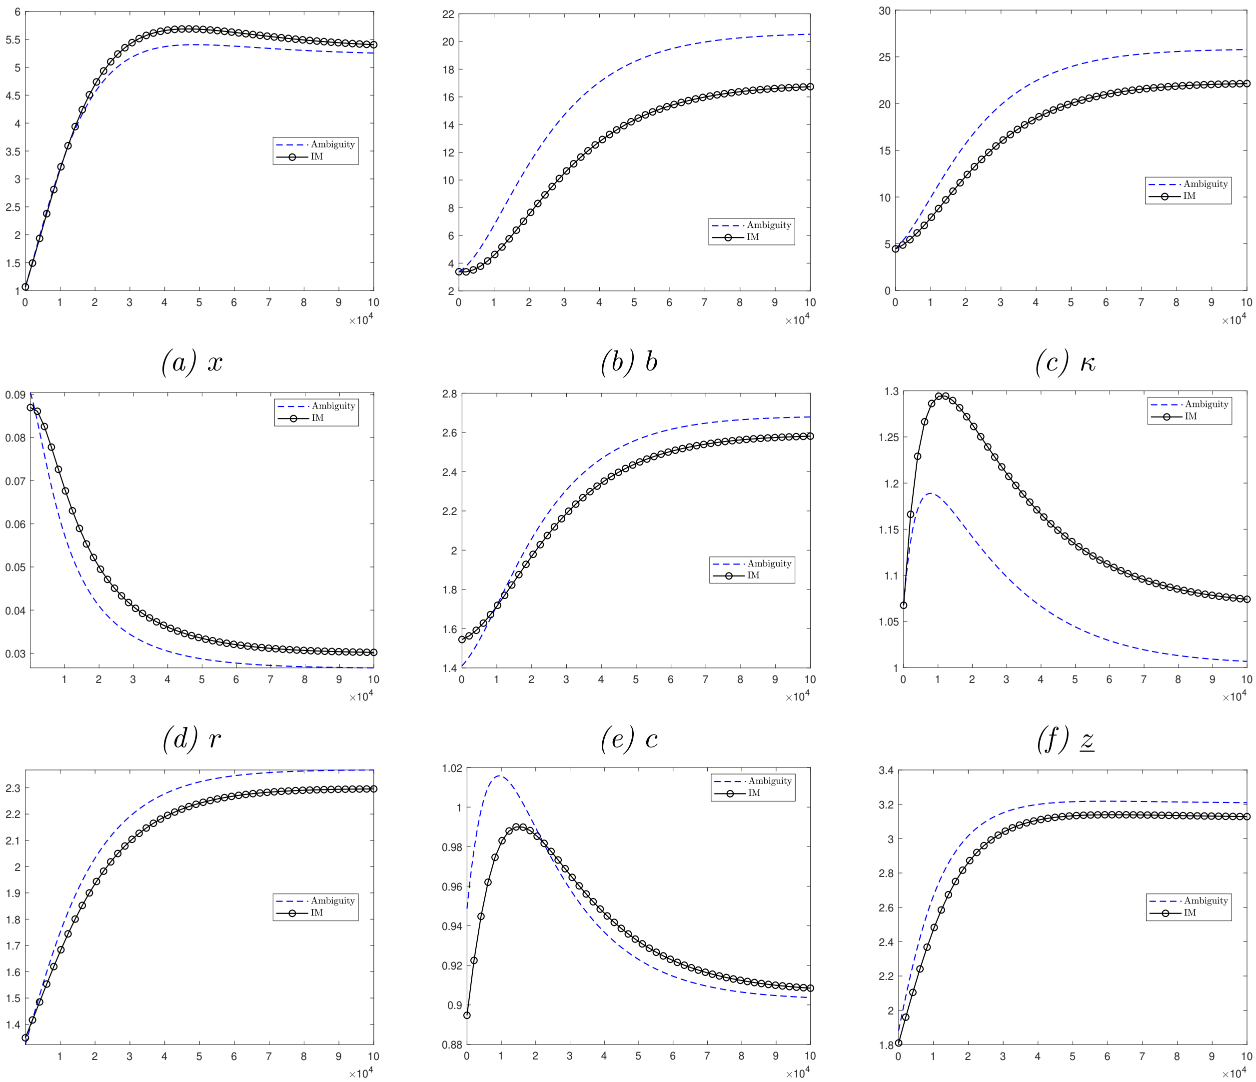

(a) x *(b)* b *(c)* κ

(d) r *(e)* c *(f)* $\underline{z}$

(g) w *(h)* ℓ *(i)* y

Figure 1: Dynamics of the ambiguous economy (blue) and the Itskhoki and Moll (2019) economy (black).

actual consumption eventually reaches a peak and then starts to decline towards zero. At the steady state, the orange line lies above the blue one, indicating that if the worst-case scenario were to actually occur, workers would instantly increase their consumption. This would happen because an economy with ambiguity accumulates more savings than one without. Therefore, if the worst-case scenario materializes, equation (24) would become valid at the actual steady state, and workers would liquidate the excess savings, resulting in a consumption spike. Now, since we consider an economy where the worst-case scenario *never* realizes, but the economy still reaches a steady state,

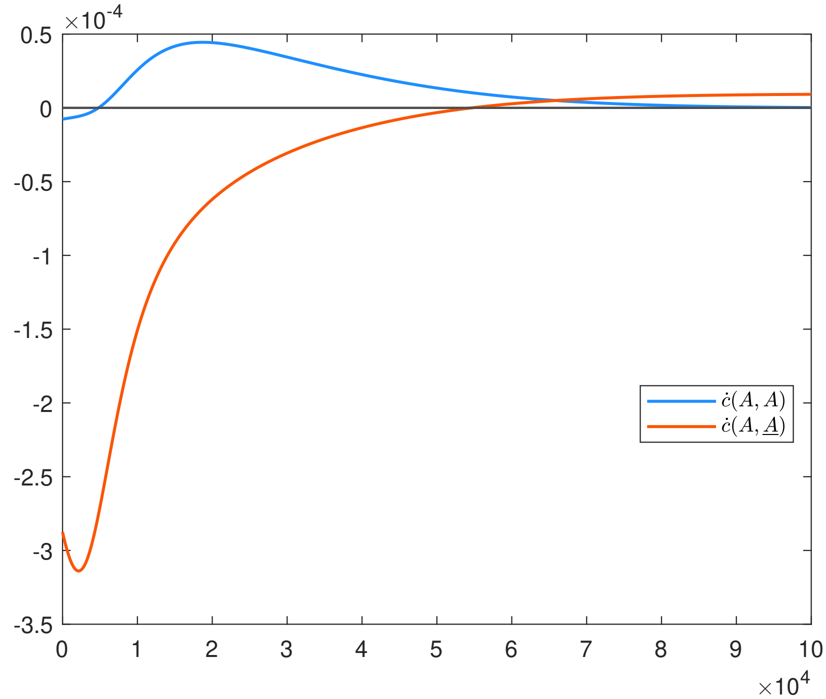

Figure 2: Dynamics of the vector fields.

which is not given by (20)-(27), the situation is different. This explains why, at the actual steady state, the *ex ante* vector field is always upward sloping.

Discussion An important feature of the model is that agents systematically underestimate the effective realization of the TFP. As in Carbonari and Maurici (2023), a key assumption is that a Dirac-delta distribution is an element of the set of perceived distributions. The Dirac-delta allows us to reduce a stochastic optimization problem, in which the distribution used by agents to weigh the future is a choice variable, into a deterministic worst-case optimization. This eliminates stochasticity because the distribution is fixed and consists of a single point.

The Dirac-delta distribution centered in the ambiguity lower bound eliminates any unusual features that might arise from the actual realization of the TFP. While this approach significantly simplifies the calculations, it may appear quite strong. However, relaxing this assumption would alter the magnitude of our results (i.e., make them less extreme) but not their overall nature. This is because, by construction, for any element of $\mathcal{P}_{t^+|t}$, the ambiguous TFP is expected to be *at most* the true TFP, which can be thought of as a Dirac-delta centered in A, $\mathcal{D}(A)$. Consequently, for any less-extreme minimizing distribution, the results can always be expressed as a convex combination of the perfect-foresight case and the lower bound Dirac-delta.[12] When we recover the ambiguity-implied cycles, we also take into account for the possibility that $\underline{A}_t = A_{t^+}$, which represents correct forecasting.

4 Business cycle

This section is divided into two parts. In the first part, we examine the impulse-response functions (IRFs) following a temporary ambiguity shock. By employing IRFs, we can trace the temporal progression of these effects, providing insights into how different macroeconomic variables respond and adjust over time following such a shock. In the second part, we calibrate our model to the U.S. economy using quarterly data from 1960Q1 to 2023Q3 and also analyze the properties of the implied cycles.

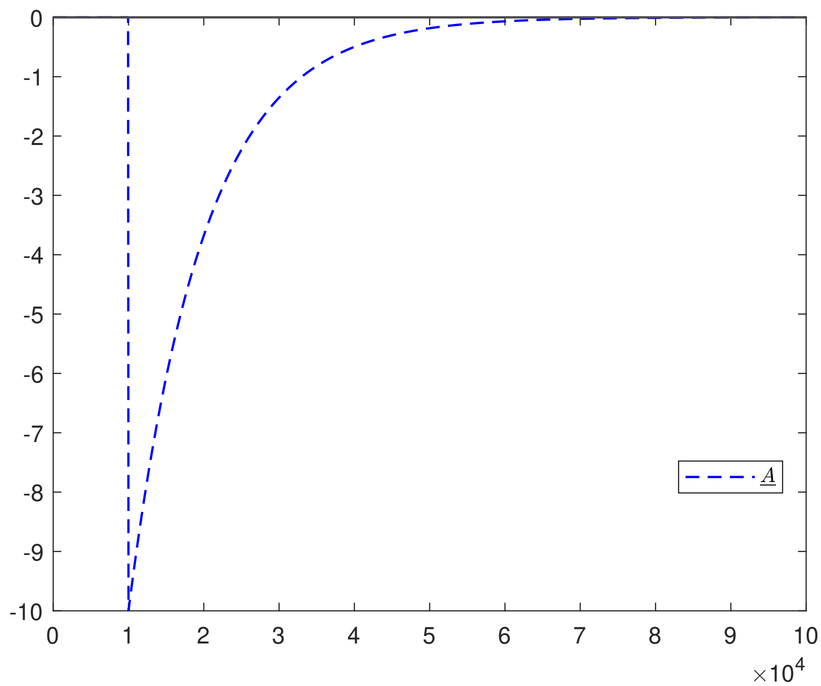

Figure 3: Ambiguity shock.

Impulse-Response Functions (IRFs) We now simulate the dynamics generated by an ambiguity shock. We assume that in the long run ambiguity disappears. Hence, the IRFs describe the short-run dynamics. We assume that the economy is at its non-ambiguous steady state, and, at time t, there is an ambiguity shock (-10%).[13] $\underline{A}$

[12] For a similar argument, see Battigalli et al. (2016).

[13] Note that a -10% shock, compared to a situation without ambiguity, corresponds to $0.9 \approx M = 0.93$, defined in the following paragraph.

falls and recovers over time according to

$$\dot{\underline{A}} = \rho_{\underline{A}} \left(A - \underline{A} \right). \tag{31}$$

To guarantee stability, we require $\rho_{\underline{A}} > 0$, which we set equal to 0.1.[14]

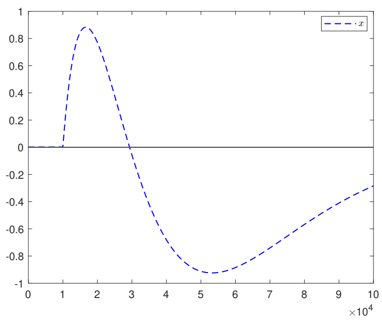

(a) x

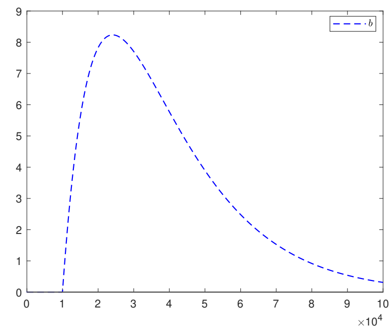

(b) b

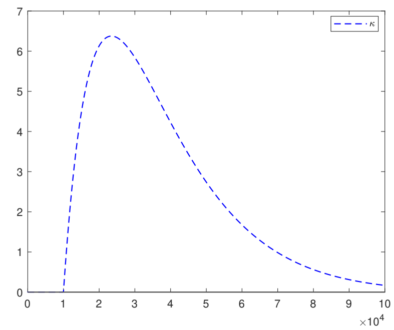

(c) κ

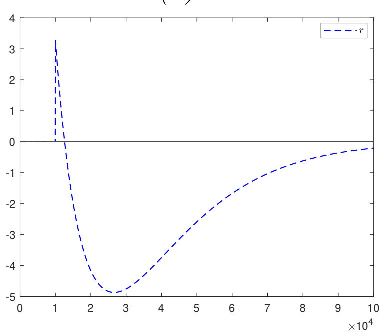

(d) r

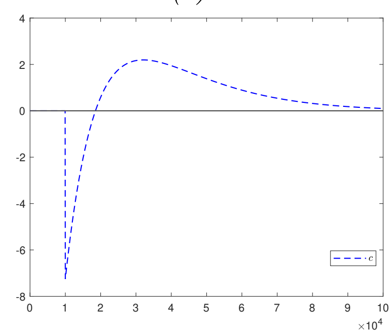

(e) c

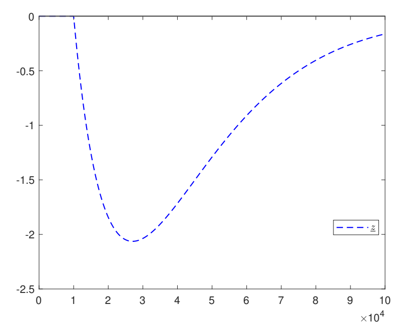

(f) $\underline{z}$

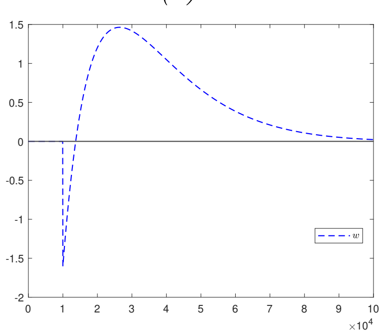

(g) w

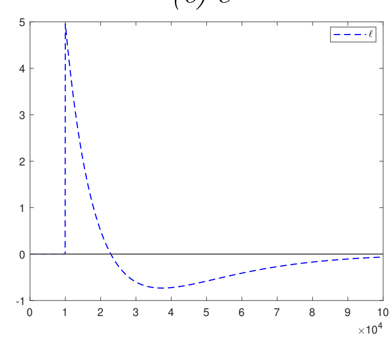

(h) ℓ

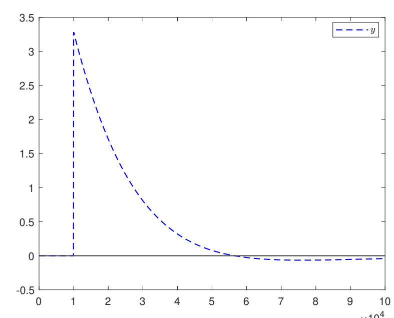

(i) y

Figure 4: Responses to an ambiguity shock.

The negative shock on $\underline{A}$ that hits the economy is shown in Figure 3 while Figure 4 reports the IRFs. Higher ambiguity causes an immediate drop in workers' consumption, accompanied by an increase in hours worked (4e and 4h) and asset holdings (4b). Excess

[14]Setting $\rho_{\underline{A}} = 0.1$ implies an exponential recovery, given by (31), equal to $e^{-\rho_{\underline{A}}} \approx 0.9$, which is standard in discrete-time models.

labor supply reduces wages (4g) while aggregate output sharply increases (4i), thereby boosting interest rate (4d) and entrepreneurs' profits and wealth (4a). Over time, the ambiguity shock vanishes and the expectation converges towards the true value of the TFP. After the initial increase driven by workers' savings (4b) and entrepreneurial wealth (4a), aggregate capital begins to decline (4c), mirroring the reduction in workers' savings. This leads to a reduction of the interest rate (4d) and the productivity cutoff (4f), both driven by the excess supply of capital during the transition dynamics. Workers' consumption increases due to higher levels of past savings and reduced necessity for precautionary savings (4e). Lastly, aggregate production, after the initial increase, rapidly returns to its long-run level, reflecting the pattern of hours worked.

Calibration and cycles In line with Altug et al. (2020), our parameter calibration aims to match two sets of cycle moments – correlations and standard deviations – of some filtered variables (i.e., total consumption, GDP, investments, and hours worked). For both simulated and actual data, we apply the HP-filter to their logarithms,[15] as in Hodrick and Prescott (1997), with a penalty parameter of 1600, consistent with standard practice for quarterly data.

Parameter	Value	Description	Source
ρ	0.03	workers' discount rate	standard value
δ	0.1	entrepreneurs' discount rate	standard value
α	0.33	capital share	standard value
λ	5	financial constraint	Gatt (2024)
ψ	1	marginal disutility of work	Itskhoki and Moll (2019)
φ	1.22	inverse Frisch elasticity	Chetty et al. (2011)
η	2.79	inverse skill-heterogeneity	calibrated
m	0.18	$\underline{A}_t/A_{t^+}$, lower bound	calibrated
M	0.93	$\underline{A}_t/A_{t^+}$, unconstrained mean	calibrated
σ	0.045	$\underline{A}_t/A_{t^+}$, unconstrained std. deviation	calibrated
σ_A	0.01	$\ln A_{t^+}$, noise std. deviation	calibrated

Table 1: Parameter values for calibration.

We model time-varying ambiguity, i.e., $\underline{A}_t = \underline{A}_t(A_{t^+}, m_t, M_t, \sigma_t)$, and allow for the

[15]Notice that in simulating the model, individual investments are periodically negative. To avoid taking the logarithm of a negative number we apply the following procedure. For each combination of parameters to calibrate, we evaluate the minimum of investments over time, $\underline{\dot{\kappa}}$, and define a new proxy for investments as $\dot{\kappa} + 1.01 \times |\underline{\dot{\kappa}}|$, which is strictly positive.

possibility that $\underline{A}_t = A_{t^+}$. When computing the cycles, we describe ambiguity as a rescaled truncated normal distribution, defined by

$$\underline{A}_t = A_{t^+} \times \begin{cases} m_t & \text{if} \quad \mathcal{N}(M_t, \sigma_t) \leq m_t, \\ \mathcal{N}(M_t, \sigma_t) & \text{if} \quad m_t < \mathcal{N}(M_t, \sigma_t) \leq 1, \\ 1 & \text{if} \quad \mathcal{N}(M_t, \sigma_t) > 1. \end{cases}$$

Then $m_t \in (0, 1]$ denotes the maximum level of ambiguity, while M_t and σ_t^2 are, respectively, the mean and variance parameters of a normal distribution.[16]

For calibration purposes, we assume the following process describing the logarithm of the TFP,

$$\ln A_{t^+} = \rho_A \ln A_t + \varepsilon_{t^+},$$

where $\varepsilon_{t^+} \sim \mathcal{N}(0, \sigma_A)$, while $\rho_A \in (-1, 1)$, which governs the speed of convergence, is set to 0.9.[17]

Table 1 presents the calibrated parameters.[18] We also compute the model-implied moments for the economy described by Itskhoki and Moll (2019), which serves as a reference case with no ambiguity.

	Consumption	*Production*	*Investment*	*Hours Worked*
Data, 1960Q1-2023Q3	1.13	1.52	6.40	0.93
Model with ambiguity	0.66	1.10	6.22	0.29
Model without ambiguity	0.67	1.00	8.93	0.15

Table 2: Std. deviations – Actual and simulated data.

Table 2 shows the implied standard deviations. As in the data, the volatility of consumption is less than that of aggregate production. However, both models predict that market consumption is too smooth, with standard deviations that are approximately 60% of the actual value. Regardless of the role of aggregate uncertainty,

[16] This characterization of ambiguity can easily accommodate exogenous learning processes, as in Benhabib et al. (2015).

[17] To recover the cycles implied by ambiguity, we employ discretization to transform the theoretical differential equations into difference equations. We set $\mathrm{d}\tau = 1/100$ to represent instantaneous time changes, implying that one quarter consists of 100 observations. Subsequently, we average every 100 simulated observations.

[18] We perform 120,000 iterations, discarding the initial 10% to get rid of the impact of initial conditions. This process is executed for each parameter combination, with each parameter drawn from a grid of 8 elements.

the standard deviation of output is roughly 70% of the observed value, while the volatility of overall investments exceeds that of output, reflecting a well-documented business cycle fact. Interestingly, ambiguity aversion appears to play a significant role in explaining investment volatility. Indeed, while our model perfectly reproduces this moment, the standard deviation for the model without uncertainty is approximately 40% higher than the observed value. Lastly, regarding the volatility of hours worked, although our model performs slightly better than the one with no uncertainty, the result is not satisfactory, with variability barely exceeding 30% of its observed value.

As for the remaining moments, Table 3 shows that our model outperforms the one without ambiguity, effectively mirroring the correlations between consumption, production, and investment. It also provides reasonable results for the joint movements of hours worked and consumption, as well as hours worked and production. Notably, regardless of the role of ambiguity, the model-implied correlations between hours worked and investment are too high compared to the data.

Data, 1960Q1-2023Q3	*Consumption*	*Production*	*Investment*	*Hours Worked*
Consumption	·	·	·	·
Production	80.79	·	·	·
Investment	56.44	88.05	·	·
Hours Worked	59.50	75.51	76.21	·
Model with ambiguity				
Consumption	·	·	·	·
Production	85.04	·	·	·
Investment	54.88	90.57	·	·
Hours Worked	43.02	84.07	98.93	·
Model without ambiguity				
Consumption	·	·	·	·
Production	99.95	·	·	·
Investment	99.05	99.41	·	·
Hours Worked	99.50	99.76	99.81	·

Table 3: Correlations – Actual and simulated data.

To identify the usual medium-term business cycle we apply a band-pass filter, as described in Baxter and King (1999), with frequency bands ranging from 1.5 to 8 periods. Figure 5 plots the ambiguity-implied cycles, with ambiguity read on the left axis and the other variables on the right axis, allowing for effective comparison despite different scales. Ambiguity serves as an amplification shock for both output and workers' savings, as shown in Figure 5a and Figure 5e, although it exhibits a low correlation with the latter. Conversely, ambiguity displays much greater volatility compared to the interest rate, as illustrated in Figure 5b. Finally, Figure 5c and Figure 5d show a one-to-one effect of ambiguity on wages and entrepreneurial wealth.

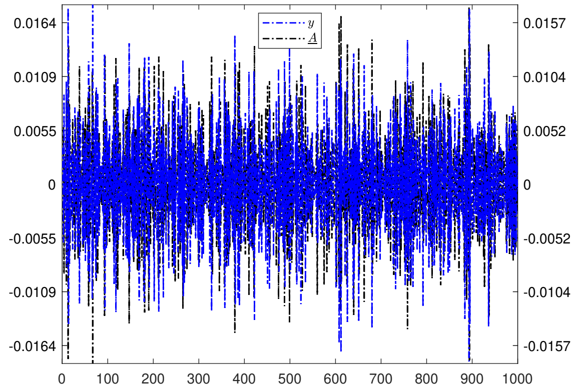

(a) Ambiguity and production.

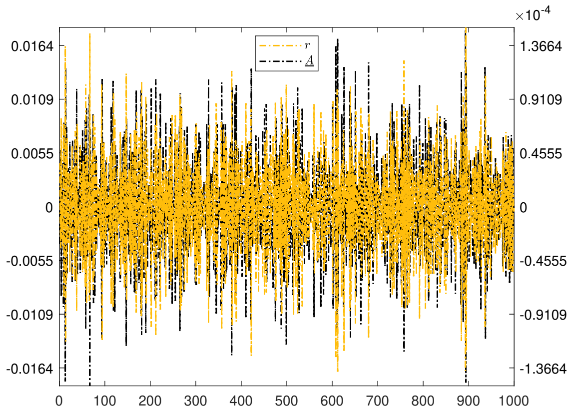

(b) Ambiguity and interest rate.

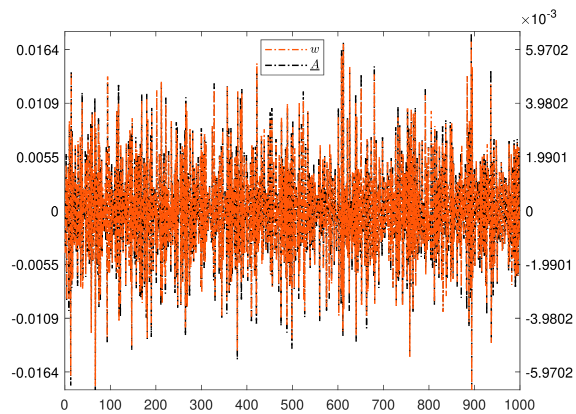

(c) Ambiguity and wages.

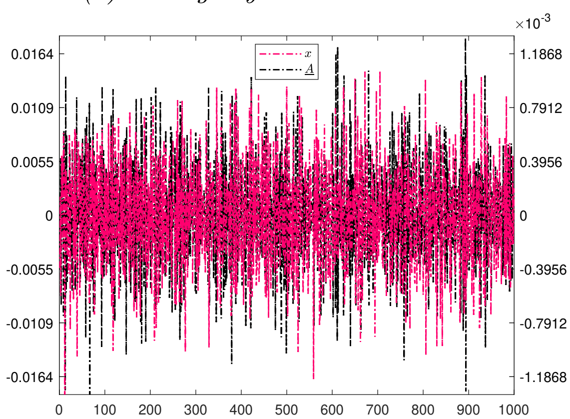

(d) Ambiguity and entrepreneurial wealth.

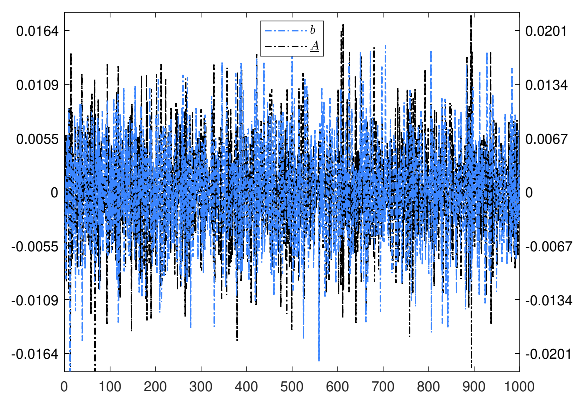

(e) Ambiguity and workers' savings.

Figure 5: Business cycle.

5 Conclusion

In this paper, we study a credit-constrained economy with heterogeneous entrepreneurs, as in Itskhoki and Moll (2019), where agents face Knightian uncertainty about the future realization of the aggregate TFP. We show that, similar to Ilut and Schneider (2014), incorporating worst-case scenarios into long-run expectations drives the equilibrium away from its non-ambiguous counterpart. Furthermore, we prove that the (local) stability properties of the ambiguous equilibrium are not affected by ambiguity.

Comparing our economy with one having the same fundamentals but without ambiguity, we observe that, due to precautionary motives, workers accumulate more assets. This offsets the reduction in entrepreneurial wealth, lowers interest rates, and reduces the productivity cutoff. Workers consume less in the early stages but end up consuming more at the ambiguity-implied steady state. Additionally, according to our calibration, we find an increase in aggregate production.

We examine the back-propagation mechanisms of ambiguity shocks. By characterizing short- and long-term dynamics, we observe how a temporary ambiguity shock impacts the accumulation processes of entrepreneurs and workers.

Employing discretization techniques and time-varying ambiguity modeling, we outline the cyclical patterns driven by ambiguity. To show the implications of ambiguity across key economic variables, we calibrate the model to match the U.S. economy. Overall, introducing ambiguity significantly improves the ability of the model to reproduce actual moments.

References

Yves Achdou, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. Income and wealth distribution in macroeconomics: A continuous-time approach. *The Review of Economic Studies*, 89(1):45–86, 2022.

S Rao Aiyagari. Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics*, 109(3):659–684, 1994.

Sumru Altug, Fabrice Collard, Cem Cakmaklı, Sujoy Mukerji, and Han Özsöylev. Ambiguous business cycles: A quantitative assessment. *Review of Economic Dynamics*, 38:220–237, 2020.

Abhijit V Banerjee and Benjamin Moll. Why does misallocation persist? *American Economic Journal: Macroeconomics*, 2(1):189–206, 2010.

Pierpaolo Battigalli, Simone Cerreia-Vioglio, Fabio Maccheroni, and Massimo Marinacci. A note on comparative ambiguity aversion and justifiability. *Econometrica*, 84(5): 1903–1916, 2016.

Marianne Baxter and Robert G King. Measuring business cycles: approximate band-pass filters for economic time series. *Review of Economics and Statistics*, 81(4): 575–593, 1999.

Jess Benhabib, Pengfei Wang, and Yi Wen. Sentiments and aggregate demand fluctuations. *Econometrica*, 83(2):549–585, 2015.

Truman F Bewley. Knightian decision theory. Part I. *Decisions in economics and finance*, 25:79–110, 2002.

Rhys Bidder and Ian Dew-Becker. Long-run risk is the worst-case scenario. *American Economic Review*, 106(9):2494–2527, 2016.

Francisco J Buera and Benjamin Moll. Aggregate implications of a credit crunch: The importance of heterogeneity. *American Economic Journal: Macroeconomics*, 7(3): 1–42, 2015.

Lorenzo Carbonari and Filippo Maurici. Firm heterogeneity, financial frictions and ambiguity. *Journal of Economic Dynamics and Control*, 155:104736, 2023.

Marco Carli. Green ambiguity. Working paper, available at SSRN, 2024.

Zengjing Chen and Larry Epstein. Ambiguity, risk, and asset returns in continuous time. *Econometrica*, 70(4):1403–1443, 2002.

Raj Chetty, Adam Guren, Day Manoli, and Andrea Weber. Are micro and macro labor supply elasticities consistent? a review of evidence on the intensive and extensive margins. *American Economic Review*, 101(3):471–75, 2011.

J Bradford De Long, Andrei Shleifer, Lawrence H Summers, and Robert J Waldmann. Noise trader risk in financial markets. *Journal of Political Economy*, 98(4):703–738, 1990.

Larry G Epstein and Shaolin Ji. Ambiguous volatility and asset pricing in continuous time. *The Review of Financial Studies*, 26(7):1740–1786, 2013.

Larry G Epstein and Martin Schneider. Recursive multiple-priors. *Journal of Economic Theory*, 113(1):1–31, 2003.

Larry G Epstein and Martin Schneider. Learning under ambiguity. *The Review of Economic Studies*, 74(4):1275–1303, 2007.

Xavier Gabaix, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. The dynamics of inequality. *Econometrica*, 84(6):2071–2111, 2016.

Patrick Gagliardini, Paolo Porchia, and Fabio Trojani. Ambiguity aversion and the term structure of interest rates. *The Review of Financial Studies*, 22(10):4157–4188, 2008.

William Gatt. Loan-to-value limits as a macroprudential policy tool: Developments in theory and practice. *Journal of Economic Surveys*, 38(1):232–267, 2024.

Itzhak Gilboa and David Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153, 1989.

Lars Peter Hansen and Jianjun Miao. Asset pricing under smooth ambiguity in continuous time. *Economic Theory*, 74(2):335–371, 2022.

Lars Peter Hansen and Thomas J Sargent. Robustness and ambiguity in continuous time. *Journal of Economic Theory*, 146(3):1195–1223, 2011.

Lars Peter Hansen and Thomas J Sargent. Structured ambiguity and model misspecification. *Journal of Economic Theory*, 199:105165, 2022.

Robert J Hodrick and Edward C Prescott. Postwar us business cycles: an empirical investigation. *Journal of Money, credit, and Banking*, pages 1–16, 1997.

Cosmin L Ilut and Martin Schneider. Ambiguous business cycles. *American Economic Review*, 104(8):2368–99, 2014.

Oleg Itskhoki and Benjamin Moll. Optimal development policies with financial frictions. *Econometrica*, 87(1):139–173, 2019.

Charles I Jones and Jihee Kim. A schumpeterian model of top income inequality. *Journal of Political Economy*, 126(5):1785–1826, 2018.

Peter Klibanoff, Massimo Marinacci, and Sujoy Mukerji. A smooth model of decision making under ambiguity. *Econometrica*, 73(6):1849–1892, 2005.

Peter Klibanoff, Massimo Marinacci, and Sujoy Mukerji. Recursive smooth ambiguity preferences. *Journal of Economic Theory*, 144(3):930–976, 2009.

Qian Lin and Frank Riedel. Optimal consumption and portfolio choice with ambiguous interest rates and volatility. *Economic Theory*, 71(3):1189–1202, 2021.

Claudio Michelacci and Luigi Paciello. Ambiguity aversion and heterogeneity in households' beliefs. *American Economic Journal: Macroeconomics*, 16(2):95–126, 2024.

Benjamin Moll, Lukasz Rachel, and Pascual Restrepo. Uneven growth: automation's impact on income and wealth inequality. *Econometrica*, 90(6):2645–2683, 2022.

Kristoffer P Nimark. Man-bites-dog business cycles. *American Economic Review*, 104 (8):2320–2367, 2014.

Appendix

Aggregation It is helpful to evaluate two integrals of the Pareto distribution,

$$\int_{z\geq \underline{z}} \mathrm{d}\mathcal{F} = \int_{z\geq \underline{z}} \eta z^{-\eta-1}\,\mathrm{d}z = \underline{z}^{-\eta},$$
$$\int_{z\geq \underline{z}} z\,\mathrm{d}\mathcal{F} = \int_{z\geq \underline{z}} \eta z^{-\eta}\,\mathrm{d}z = \frac{\eta}{\eta-1}\underline{z}^{1-\eta}.$$

Define $\mathcal{E} = \mathcal{G}\times\mathcal{F}$ as the joint distribution describing both entrepreneurs' assets and productivity. Aggregate demand for capital and labor, (13) and (14), are obtained by integrating (9) and (10),

$$\int\int_{z\geq \underline{z}} k(a,z)\,\mathrm{d}\mathcal{E} = \lambda x \int_{z\geq \underline{z}} \mathrm{d}\mathcal{F},$$
$$\int\int_{z\geq \underline{z}} n(a,z)\,\mathrm{d}\mathcal{E} = \lambda \left[\frac{(1-\alpha)}{w}A\right]^{\frac{1}{\alpha}} x \int_{z\geq \underline{z}} z\,\mathrm{d}\mathcal{F}.$$

Now, plug (10) into (8) to write individual production only in terms of k,

$$A\left[\frac{1-\alpha}{w}A\right]^{\frac{1-\alpha}{\alpha}} zk.$$

Aggregating,

$$y = A\left(\frac{\eta}{\eta-1}\right)^{\alpha} \underline{z}^{\alpha}\kappa^{\alpha}\ell^{1-\alpha},$$

which gives (17) once (13) is plugged in.

Linearizing It is useful to define some partial effects,

$$\frac{\partial y}{\partial c} = \left(1-\frac{\omega}{\alpha}\right)\frac{y}{c},$$
$$\frac{\partial y}{\partial x} = \frac{\omega}{\eta}\frac{y}{x},$$
$$\frac{\partial y}{\partial \kappa} = \frac{\omega(\eta-1)}{\eta}\frac{y}{\kappa}.$$

Linearizing around the ambiguous steady state, i.e., computed at $A = \underline{A}$, yields

$$\left.\frac{\partial \dot{c}}{\partial c}\right|_{\underline{A}} = \left(1-\frac{\omega}{\alpha}\right)\widehat{r},$$
$$\left.\frac{\partial \dot{c}}{\partial x}\right|_{\underline{A}} = \frac{\omega}{\eta}\frac{\widehat{rc}}{\widehat{x}},$$
$$\left.\frac{\partial \dot{c}}{\partial \kappa}\right|_{\underline{A}} = \frac{\omega(\eta-1)-\eta}{\eta}\frac{\widehat{rc}}{\widehat{\kappa}},$$

$$\left.\frac{\partial \dot{x}}{\partial c}\right|_{\underline{A}} = \left[\frac{\alpha}{\eta} + \frac{\alpha(\eta-1)}{\eta}\frac{\widehat{x}}{\widehat{\kappa}}\right]\left(1-\frac{\omega}{\alpha}\right)\frac{\widehat{y}}{\widehat{c}},$$

$$\left.\frac{\partial \dot{x}}{\partial x}\right|_{\underline{A}} = \frac{\alpha\omega}{\eta^2}\left[1+(\eta-1)\frac{\widehat{x}}{\widehat{\kappa}}\right]\frac{\widehat{y}}{\widehat{x}} + \widehat{r} - \delta,$$

$$\left.\frac{\partial \dot{x}}{\partial \kappa}\right|_{\underline{A}} = \frac{\alpha\omega(\eta-1)}{\eta^2}\frac{\widehat{y}}{\widehat{\kappa}} + \frac{\omega(\eta-1)-\eta}{\eta}\widehat{r}\frac{\widehat{x}}{\widehat{\kappa}},$$

$$\left.\frac{\partial \dot{\kappa}}{\partial c}\right|_{\underline{A}} = \left(1-\frac{\omega}{\alpha}\right)\frac{\widehat{y}}{\widehat{c}} - 1,$$

$$\left.\frac{\partial \dot{\kappa}}{\partial x}\right|_{\underline{A}} = \frac{\omega}{\eta}\frac{\widehat{y}}{\widehat{x}} - \delta,$$

$$\left.\frac{\partial \dot{\kappa}}{\partial \kappa}\right|_{\underline{A}} = \frac{\omega(\eta-1)}{\eta}\frac{\widehat{y}}{\widehat{\kappa}}.$$

Using (20), (21), (22), and (23), we can write

$$\mathbf{J} = \begin{bmatrix} \left(1-\frac{\omega}{\alpha}\right)\rho & \frac{(\delta-\rho)\eta-\delta\alpha}{\alpha\eta}\omega\rho & \frac{\omega(\eta-1)-\eta}{\alpha\eta}\frac{(\delta-\rho)\eta-\delta\alpha}{(\delta-\rho)(\eta-1)}\rho^2 \\ -\frac{\omega-\alpha}{(\delta-\rho)\eta-\delta\alpha}\delta & \left(\frac{\omega}{\eta}-1\right)\delta+\rho & \frac{\rho\omega}{\eta} + \frac{\omega(\eta-1)-\eta}{\eta(\delta-\rho)(\eta-1)}\rho^2 \\ \left(1-\frac{\omega}{\alpha}\right)\frac{(\delta-\rho)\eta}{(\delta-\rho)\eta-\delta\alpha} - 1 & \frac{\omega}{\alpha}(\delta-\rho)-\delta & \frac{\omega}{\alpha}\rho \end{bmatrix}. \tag{32}$$

Since we focus on stable solutions, the nature of the endogenous variables implies that two of the eigenvalues of $\mathbf{J}$ must have negative real parts and one positive real part. Diagonalize $\mathbf{J}$ as $\mathbf{J} = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^{-1}$, where $\mathbf{Q}$ and $\mathbf{\Lambda}$ are, respectively, the matrix of eigenvectors and the eigenvalues of $\mathbf{J}$. Define the change of variables

$$\mathbf{g} = \mathbf{Q}^{-1}\begin{bmatrix} c-\widehat{c} \\ x-\widehat{x} \\ \kappa-\widehat{\kappa} \end{bmatrix} \rightarrow \dot{\mathbf{g}} = \mathbf{Q}^{-1}\begin{bmatrix} \dot{c} \\ \dot{x} \\ \dot{\kappa} \end{bmatrix} = \mathbf{\Lambda}\mathbf{g},$$

giving the solution

$$\mathbf{g}_t = e^{\mathbf{\Lambda}t}\mathbf{g}_0.$$

Stability is reached when $\lim_{t\to\infty}\mathbf{g}_t = \begin{bmatrix}0 & 0 & 0\end{bmatrix}^\top$. When the diagonal elements of $\mathbf{\Lambda}$ have negative real parts, the exponential pushes the equation to 0. Otherwise, call $\Lambda_U > 0$ the real part of the exploding eigenvalue in $\mathbf{\Lambda}$, and $\begin{bmatrix}Q_c & Q_x & Q_\kappa\end{bmatrix}$ the associated left-eigenvector. Again, stability implies

$$\lim_{t\to\infty} e^{\Lambda_U t}\begin{bmatrix}Q_c & Q_x & Q_\kappa\end{bmatrix}\begin{bmatrix} c-\widehat{c} \\ x-\widehat{x} \\ \kappa-\widehat{\kappa} \end{bmatrix} = 0.$$

Because $\lim_{t\to\infty} e^{\Lambda_U t} = \infty$, it must be that

$$\begin{bmatrix} Q_c & Q_x & Q_\kappa \end{bmatrix} \begin{bmatrix} c - \widehat{c} \\ x - \widehat{x} \\ \kappa - \widehat{\kappa} \end{bmatrix} = 0. \tag{33}$$

Rearranging (33) yields (30).

Notice that the transversality conditions hold, since, the one for workers can be written as

$$\lim_{\tau\to\infty} e^{-r_\tau(\tau-t)} b_\tau = 0,$$

where the sufficient conditions are $\lim_{\tau\to\infty} r_\tau > 0$ and $\lim_{\tau\to\infty} |b_\tau| < \infty$, both guaranteed by the stability implied by $\mathbf{J}$. The transversality condition for entrepreneurs follows the same reasoning.

Vector field The interplay between ambiguous expectations and realizations can be captured by constructing vector fields induced by ambiguity, see Figure 6. In each period, ambiguity generates two distinct vector fields: one *ex ante*, based on current forward-looking ambiguity, and one *ex post*, reflecting the actual realization of the TFP.

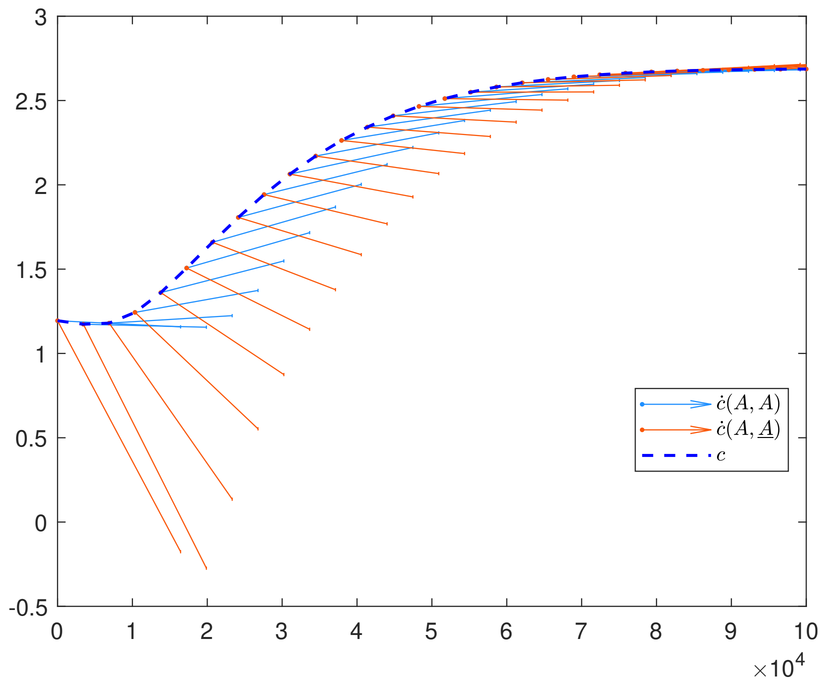

Figure 6: Consumption vector fields.

The empirical vector field (light blue line) shows three distinct phases: initially driving consumption downward to prioritize savings in anticipation of a potential negative TFP shock. As savings accumulate, individuals use their increased wealth to boost consumption. Finally, as the model stabilizes, the vector field flattens, indicating a stable consumption path.

In contrast, the theoretical ambiguous vector field (orange line) undergoes a steady rotation. Early expectations lead to reduced consumption, resulting in a counterclockwise rotation driven by excess savings and persistently low interest rates. At the steady state, the *ex ante* vector field slopes upwards, reflecting a more stable and positive consumption trajectory.