Demand expectations and the timing of stimulus policies*

Bernardo Guimaraes† Caio Machado‡

December 2014

Abstract

This paper proposes a simple macroeconomic model with staggered investment decisions. The model captures the dynamic coordination problem arising from demand externalities and fixed costs of investment. In times of low economic activity, a firm faces low demand and hence has less incentives for investing, which reinforces firms' expectations of low demand. In the unique equilibrium of the model, demand expectations are pinned down by fundamentals and history. Owing to the beliefs that arise in equilibrium, there is no special reason for stimulus at times of low economic activity.

JEL Classification: D84, E32, E62

Keywords: Demand expectations, coordination, fiscal stimulus, timing frictions

*We thank the editor Francesco Caselli, three anonymous referees, Marios Angeletos, Braz Camargo, Felipe Iachan, Nobuhiro Kiyotaki, V. Filipe Martins-da-Rocha, Stephen Morris, Alison Oliveira, Ana Elisa Pereira, Mauro Rodrigues and Mark Wright as well as seminar participants at U Carlos III, LACEA-LAMES 2014 (Sao Paulo), NASM-ES 2014 (Minneapolis), Princeton, PUC-Rio, Sao Paulo School of Economics – FGV, U Sao Paulo, SBE 2013 (Iguaçu), the XVII Workshop of International Economics and Finance (San Jose) and U Zurich. Bernardo Guimaraes gratefully acknowledges financial support from CNPq. Caio Machado gratefully acknowledges financial support from CAPES and FAPESP.

†Sao Paulo School of Economics - FGV

‡Sao Paulo School of Economics - FGV

1 Introduction

Pessimistic expectations are often said to play a role in recessions. This idea is captured by models with demand externalities that generate strategic complementarities in production.[1] In times of low economic activity, a firm faces low demand and thus has low incentives for investment. In a dynamic setting, this feedback effect may trap the economy in a regime of low output. One important policy issue is then the optimal stimulus policy for an economy subject to this dynamic coordination problem. Is there any special reason for stimulus at times of low economic activity owing to this coordination channel?

This paper develops a model that features this dynamic coordination problem. Investment decisions are staggered, hence economic activity is a state variable. A firm's optimal investment decision depends on its expectations about others' actions in the future. Naturally, the optimal policy also depends on those expectations. The model generates a unique set of rationalizable beliefs about others' actions. The beliefs that arise in equilibrium exactly offset the dynamic coordination problem faced by firms, so there is no special reason for stimulus at times of low economic activity.

The model features monopolistic competition. Investment is a payment of a fixed cost that increases production capacity. Returns to investment depend on future demand, and hence on whether producers with subsequent investment opportunities choose to take them as well. Thus investment decisions are strategic complements. Producers of each variety receive investment opportunities according to a Poison clock, a simple way to capture production decisions that can not adjust overnight. This assumption implies that investment decisions are not synchronized, so a producer has to form expectations about others' future decisions when deciding about investment.

Return to investment depends not only on demand but also on productivity. If the increase in production resulting from investing is large enough, then investing is a dominant strategy. Likewise, if productivity is very low, investing is a dominated strategy. In an intermediate range, a producer's decision depends on his expectations about the actions of others. In a world with no shocks, that gives rise to multiple equilibria. There is a region of parameters where investing is the optimal decision if agents expect others to do so, but refraining from investment is the best choice in case of pessimistic beliefs.

The solution to the planner's problem in the environment with no shocks differs from the decentralized equilibrium in three ways: (i) there is no room for multiplicity of beliefs, as one would expect; (ii) the planner requires a lower productivity to invest because it internalizes the benefits to consumers from selling at a cheaper price (monopoly distortion); and (iii) the difference between the planner's solution and the decentralized equilibrium is

[1]Seminal papers in this literature are Cooper and John (1988), Kiyotaki (1988) and Murphy et al. (1989).

particularly large at times of low economic activity.

The third point is particularly important. Agents might get stuck in a situation where economic activity is low, hence there is low demand and firms prefer not to invest even though productivity would be high enough to encourage investment if demand were high. In this situation, firms would like others to invest, so that demand would increase and then they would be happy to invest as well. A dynamic coordination problem means nobody wants to be the first to invest and the economy is trapped in a situation with low economic activity. The planner would be particularly keen to invest in this situation.

Once we allow for shocks, a unique equilibrium arises in the model, as in Frankel and Pauzner (2000). There is no room for arbitrary beliefs, agents' expectations are pinned down by the model. The basic idea is that fully pessimistic beliefs are not rationalizable in a region where a small shock to productivity would make it dominant for all firms to invest. Likewise, fully optimistic beliefs are not rationalizable in a region where a small shock to productivity takes the economy to a region where investing is a dominated strategy. Agents know all others will reason like this and try to anticipate what others will do. This process yields a unique rationalizable set of strategies and a unique set of beliefs.

The main result of the paper is that owing to the beliefs that arise in equilibrium, the third difference between the planner's solution and the decentralized equilibrium vanishes. The maximum amount of investment subsidies the planner is willing to provide at times of high and low economic activity is exactly the same. There is no special reason for subsidies at times of low economic activity.

How are equilibrium beliefs? Consider an agent indifferent between investing or not in a state of low economic activity. She understands that if fundamentals get a bit worse, firms will still be refraining from investing but there will be no major change in the state of the economy. Conversely, a slight improvement in fundamentals will trigger a recovery because firms will choose to invest and that will push the economy to a situation where investing is profitable for everyone. Owing to larger demand, firms will then have more incentives to invest, so it will take a large negative productivity shock to offset the benefit from increasing demand and stop the recovery. The fundamental asymmetry is that bad news basically leave the economy parked in a region of inaction, while good news drive the economy to a different state. The recovery is thus just waiting for a small piece of good news. Hence, in a pivotal circumstance, optimistic beliefs make perfect sense at times of low economic activity. The same reasoning implies that in a pivotal circumstance with high economic activity, beliefs will be more pessimistic: the economy is close to an investment slump.

Optimistic beliefs thus arise at times of decent fundamentals and low economic activity and that is what prevents the economy from being trapped in a regime of low output.

Equilibrium beliefs solve the dynamic coordination problem. The only remaining difference between the planner's solution and the decentralized equilibrium is the externality due to market power (which is unrelated to economic activity with (standard) constant elasticity specifications for consumption and production).

The demand externalities that play a key role in this paper are in the seminal contributions by Blanchard and Kiyotaki (1987), Kiyotaki (1988) and Murphy et al. (1989). When others produce more, the demand for a particular variety shifts to the right, and its producer finds it optimal to increase production. In Kiyotaki (1988), multiple equilibria arise because of increasing returns to scale. The model in this paper would also give rise to multiple equilibria in the absence of shocks to fundamentals or timing frictions, owing to the assumption of a fixed cost that increases production capacity.

A branch of the literature takes expectations to be driven by some "sunspot" variable, or simply, in the words of Keynes, by "*animal spirits"*. Depending on agents' expectations, coordination failures might arise and an inefficient equilibrium might be played.[2] Despite generating interesting insights, this approach does not allow us to understand how policies affect expectations. In models with multiple equilibrium, government policies can only hope to eliminate the "bad equilibrium". Here, in constrast, policies affect agents' beliefs about others' actions.

This paper is closely related to the theoretical contributions in Frankel and Pauzner (2000) and Frankel and Burdzy (2005) that resolve indeterminacy in dynamic models. They study models with time-varying fundamentals and timing frictions similar to the ones employed in this paper, and prove there is a unique rationalizable equilibrium in their models.[3] The results in Frankel and Pauzner (2000) guarantee a unique equilibrium in our baseline model and we use the results of Frankel and Burdzy (2005) in one extension of the model.[4] This paper is also related to the global games literature, which has been used to study a wide variety of economic problems that exhibit strategic complementarities, but differently from that literature, there is no asymmetric information in this model.[5]

There has been a lot of research incorporating strategics complementarities and coordination issues in macroeconomics.[6] However, there has not been much work applying those

[2]See, e.g., Cooper and John (1988), Benhabib and Farmer (1994) and Farmer and Guo (1994).

[3]Models with time-varying fundamentals and timing frictions have been used to study other dynamic coordination problems. Frankel and Pauzner (2002) employ a similar structure in order to analyze the timing of neighborhood change. Guimaraes (2006) studies speculative attacks. Levin (2009) studies the persistence of group behavior in a collective reputation model. He and Xiong (2012) study debt runs.

[4]See also Burdzy et al. (2001).

[5]See the seminal papers by Carlsson and Van Damme (1993) and Morris and Shin (1998). For a detailed survey, see Morris and Shin (2003).

[6]Angeletos and La'O (2010) and Angeletos and La'O (2013) show in an environment with noisy and dispersed information how self-fulfilling fluctuations can emerge. Expectations also play a key role in the literature of news-driven business cycles (e.g., Beaudry and Portier (2006)), but here expectations about future productivity depend solely on the current state of the economy. In the models of Lorenzoni (2009) and

theoretical insights to understand the effects of stimulus packages on coordination. One important exception is Sákovics and Steiner (2012). They build a model to understand who matters in coordination problems: in a recession, who should benefit from government subsidies? The results point that the government should subsidize sectors that have a large externality on others but that are not much affected by others' actions. Differently from a large literature that deals with coordination failures and expectations in macroeconomics, our focus is not on noisy and heterogeneous information, fundamentals are common knowledge here, all the action comes from dynamic frictions. This makes our framework specially suitable to understand the dynamic interplay between economic activity, productivity and beliefs that arise in equilibrium.

The paper is organized as follows. Section 2 presents the model in its simplest form. Section 3 describes the decentralized equilibrium and Section 4 shows the solution to the planner's problem. Section 5 then explains the result, highlighting the role of beliefs. Section 6 deals with important extensions: it shows a more complete macro model that leads to the same results; considers a process for productivity with mean reversion and discusses implementation of the optimal policy. Section 7 presents some numerical results and Section 8 concludes.

2 Model

2.1 Environment

Time is continuous. A composite good is produced by a perfectly competitive representative firm. At time t, Y_t units of the composite good are obtained by combining a continuum of intermediate goods, indexed by $i \in [0, 1]$, using the technology:

$$Y_t = \left(\int_0^1 y_{it}^{(\theta-1)/\theta} di \right)^{\theta/(\theta-1)}, \tag{1}$$

where y_{it} is the amount of intermediate good i used in the production of the composite good at time t and $\theta > 1$ is the elasticity of substitution. The zero-profit condition implies

$$\int_0^1 \tilde{p}_{it} y_{it} di = P_t Y_t, \tag{2}$$

Eusepi and Preston (2011), it is noisy information about current variables that leads to excessive optimism or pessimism about the future. Nimark (2008) builds a model where pricing complementarities together with private information help to explain the inertial behavior of inflation due to the inertial response of expectations (see also Angeletos and La'O (2009)). Chamley (2013) presents a model with decentralized trade, credit constraints and multiple equilibria where pessimistic expectations lead to precautionary savings, which in turn lead to low production.

where P_t is the price of the composite good and $\tilde{p}_{it}$ is the price of good i at time t.

There is a measure-one continuum of agents who discount utility at rate ρ. An agent's instantaneous utility at time t is given by $U_t = C_t$, where C_t is her instantaneous consumption of the composite good. The assumption of linear utility implies that policies will be concerned with inefficiencies in production but will not aim at providing insurance to the household.

Agent $i \in [0, 1]$ produces intermediate good i. Since y_{it} is the quantity produced by agent i at time t, her budget constraint is given by

$$P_t C_t \leq \tilde{p}_{it} y_{it} \equiv w_i P_t.$$

Prices are flexible and each price $\tilde{p}_{it}$ is optimally set by agent i at every time. Since goods are non storable, supply must equal demand at any time t.

The assumptions on technology aim at modelling staggered investment decisions in a simple and tractable way. There are 2 production regimes, a *High*-capacity regime and a *Low*-capacity regime. Agents get a chance to switch regimes according to a Poisson process with arrival rate α.[7] Once an individual is picked up, he chooses a regime and will be locked in this regime until he is selected again. Choosing the *Low* regime is costless. Choosing the *High* regime costs ψ units of the composite good.[8]

An agent in the *Low* regime can produce up y_{Lt} units at zero marginal cost at every time t, and an agent in the *High* regime can produce up to y_{Ht} units at zero marginal cost, with $y_{Ht} = A_t x_H$ and $y_{Lt} = A_t x_L$, where $x_H > x_L$ are constants and A_t is a time-varying productivity parameter.[9]

The *High* regime can be interpreted as the use of frontier technology, while the regime *Low* would correspond to a less productive technology. The cost ψ can be thought of as the cost difference between each technology and the difference $y_{Ht} - y_{Lt}$ as the resulting gain in productivity. Agents are locked in a regime until the next opportunity arises. In one interpretation, the equipament will break after some (random) time and the firm will then decide again between a more or less productive technology. Alternatively, that might capture attention frictions.[10]

[7] Real world investments require a lot of planning and take time to become publicly known, so investments from different firms are not syncronized. The Poisson process generates staggered investment decisions in a simple way. As an implication, investment decisions depend on expectations about others' actions in the near future. For further evidence on non-convex adjustment costs that lead to infrequent investment, see Hall (2000) and Cooper and Haltiwanger (2006).

[8] Investment is thus a binary decision. As shown in Gourio and Kashyap (2007), the extensive margin accounts for most of the variation in aggregate investment, so a binary choice set can capture much of the action in investment.

[9] The assumption of zero marginal cost is relaxed in Section 6.1.

[10] In another possible interpretation, ψ could be the cost of hiring a worker that cannot be fired until his contract expires. In that case, the fixed cost would not be paid at once, but that makes no difference in the

Investment requires agents to acquire a stock ψ of composite goods, which cannot be funded by their instantaneous income, so we assume agents can trade assets and borrow to invest. Owing to the assumption of linear utility, any asset with present value equal to ψ is worth ψ in equilibrium. For example, an agent might issue an asset that pays $(\rho + \alpha)\psi dt$ at every interval dt until the investment depreciates ($\rho\psi dt$ would be the interest payment and $\alpha\psi dt$ can be seen as an amortization payment since debt is reduced from ψ to 0 with probability αdt). Since agents are risk neutral, other types of assets would deliver the same results.

Let $a_t = \log(A_t)$ vary on time according to

$$da_t = \sigma dZ_t, \tag{3}$$

where $\sigma > 0$ and Z_t is a standard Brownian motion.

2.2 The agent's problem

The composite-good firm chooses its demand for each intermediate good taking prices are given. Using (1) and (2) and defining $p_{it} \equiv \tilde{p}_{it}/P_t$, we get

$$p_{it} = y_{it}^{-1/\theta} Y_t^{1/\theta}, \tag{4}$$

for $i \in [0, 1]$. Since marginal cost is zero and marginal revenue is always positive, an agent in the *Low* regime will produce y_{Lt}, and an agent in the *High* regime will produce y_{Ht}. Thus at any time t, there will be two prices in the economy, p_{Ht} and p_{Lt} (associated with production levels y_{Ht} and y_{Lt}, respectively). Hence the instantaneous income available to individuals in each regime is given by

$$w_{Ht} = p_{Ht} y_{Ht} = y_{Ht}^{\frac{\theta-1}{\theta}} Y_t^{\frac{1}{\theta}} \tag{5}$$

and

$$w_{Lt} = p_{Lt} y_{Lt} = y_{Lt}^{\frac{\theta-1}{\theta}} Y_t^{\frac{1}{\theta}}. \tag{6}$$

Moreover, using (1),

$$Y_t = \left(h_t y_{Ht}^{\frac{\theta-1}{\theta}} + (1 - h_t) y_{Lt}^{\frac{\theta-1}{\theta}} \right)^{\frac{\theta}{\theta-1}}, \tag{7}$$

where h_t is the measure of agents locked in the *High* regime.

Combining (5), (6) and (7), we get the instantaneous income of individuals locked in each regime. Let $\pi(h_t, a_t)$ be the difference between instantaneous income of agents locked in the *High* regime and agents locked in the *Low* regime when the economy is at (h_t, a_t).

model.

Then, using $y_{Lt} = e^{a_t} x_L$ and $y_{Ht} = e^{a_t} x_H$,

$$\pi(h_t, a_t) = e^{a_t} \left(h_t x_H^{\frac{\theta-1}{\theta}} + (1-h_t) x_L^{\frac{\theta-1}{\theta}} \right)^{\frac{1}{\theta-1}} \left(x_H^{\frac{\theta-1}{\theta}} - x_L^{\frac{\theta-1}{\theta}} \right). \tag{8}$$

Function π is increasing in both a_t and h_t. The effect of a_t captures the supply side incentives to invest: a larger a_t means a higher productivity differential between agents who had invested and those who had not. The effect of h_t captures the demand side incentives to invest: a larger h_t means a higher demand for a given variety. The equilibrium price of a good depends on how large y_{it}/Y_t is, so a producer benefits from others producing y_{Ht} regardless of how much she is producing. Nevertheless, since $\theta > 1$, an agent producing more reaps more benefits from a higher demand.

One key implication of (8) is that there are strategic complementarities: the higher the production level of others, the higher the incentives for a given agent to increase her production level.

A strategy is as a map $s(h_t, a_t) \mapsto \{Low, High\}$. An agent at time $t = \tau$ that has to decide whether to invest will do so if

$$\int_\tau^\infty e^{-(\rho+\alpha)(t-\tau)} E_\tau[\pi(h_t, a_t)] dt \geq \psi. \tag{9}$$

In words, investing pays off if the discounted expected additional profits of choosing the *High* regime are larger than the fixed cost ψ. Future profits $\pi(h_t, a_t)$ are discounted by the sum of the discount rate and depreciation rate $(\rho + \alpha)$.[11]

Investment decisions depend on expected profits. Producers will decide to invest not only if productivity is high, but also if they are confident they will be able to sell their varieties at a good price. Hence investment decisions crucially depend on demand expectations, which in turn are determined by expectations about the path of a_t and h_t.

3 Equilibrium

3.1 Benchmark case: no shocks

Consider the case where the fundamental a does not vary over time, $\sigma = 0$. Proposition 1 characterizes conditions under which we have multiple equilibria in this case.

Proposition 1 (No Shocks)**.** *Suppose $\sigma = 0$ and $a = \mu$. There are strictly decreasing functions $a^L : [0,1] \mapsto \Re$ and $a^H : [0,1] \mapsto \Re$ with $a^L(h) < a^H(h)$ for all $h \in [0,1]$ such that*

[11] As a tie breaking convention, an agent chooses *High* whenever she is indifferent between regimes *High* and *Low*.

1. *If $a < a^L(h_0)$ there is a unique equilibrium, agents always choose the Low regime;*

2. *If $a > a^H(h_0)$ there is an unique equilibrium, agents always choose the High regime;*

3. *If $a^L(h_0) < a < a^H(h_0)$ there are multiple equilibria, that is, both strategies High and Low can be long-run outcomes.*

Proof. See Appendix B. □

Figure 1 illustrates the result of Proposition 1. If the productivity differential is sufficiently high, agents will invest as soon as they get a chance and the economy will move to a state where $h = 1$ (and there it will rest). If the productivity differential is sufficiently low, the gains from investing are offset by the fixed cost, so not investing is a dominant strategy. In an intermediate area, there are no dominant strategies, the optimal investment decision depends on expectations about what others will do and there are multiple equilibria.

Figure 1: Equilibria without shocks

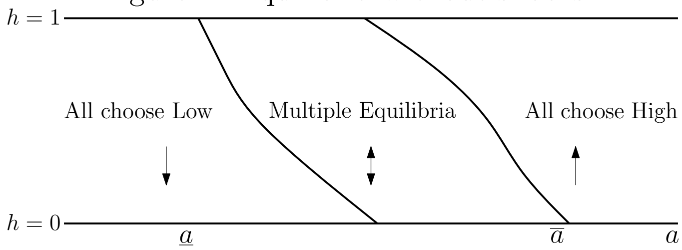

The boundaries of the region with multiple equilibria are obtained using the expression for the payoff from investing in (9). The set of states where agents are indifferent between investing or not assuming everyone will choose *High* from then on is the curve that separates the region of multiplicity from the region where all agents choose the low regime. The other boundary is calculated in the same way, assuming everyone will choose *Low* in the future.

Cycles are possible in this economy, but their existence depends on exogenous changes in beliefs. Demand expectations are not pinned down by the parameters that characterize the economy and its current state. Small subsidies to investment in the multiplicity region have no effects on beliefs.

3.2 The case with shocks

We now turn to the general case where productivity varies over time, $\sigma > 0$. We say that an agent is playing according to a threshold $a^* : [0, 1] \mapsto \Re$ if she chooses *High* whenever $a_t > a^*(h_t)$ and *Low* whenever $a_t < a^*(h_t)$. Function a^* is an equilibrium if the strategy profile where every player plays according to a^* is an equilibrium.

The model can be seen as a particular case of Frankel and Pauzner (2000). Hence we can apply Theorem 1 in their paper to show there is a unique rationalizable equilibrium where agents play according to a decreasing threshold $a^*(h)$.[12]

Proposition 2 (Frankel and Pauzner, 2000)**.** *Suppose $\sigma > 0$. There is a unique rationalizable equilibrium in the model. Agents invest if and only if $a > a^*(h)$, where a^* is a decreasing function.*

The proof shows that a unique equilibrium survives iterative elimination of strictly dominated strategies. Intuitively, consider a situation where productivity is relatively low, so a firm is only willing to invest if the probability the following firms will also invest is very high. In Figure 1, that would correspond to a point in the multiplicity region but close to its left boundary. In a world with shocks, the economy might cross to the region where investing is a dominated strategy. That imposes a cap on the probability that others will invest in the near future – the belief that all of them will invest is not rationalizable. In consequence, some dominated strategies are eliminated, which imposed further limits on beliefs agents can hold. Iterating on this process leads to a unique equilibrium.

The equilibrium is characterized by a threshold. A larger h implies that agents are willing to invest for lower values of a, as in Figure 2. Beliefs about others' investment decisions are pinned down by fundamentals (a) and history (h). Shocks to a_t and movements in h_t might affect expectations about others' actions.

Figure 2: Equilibrium with shocks

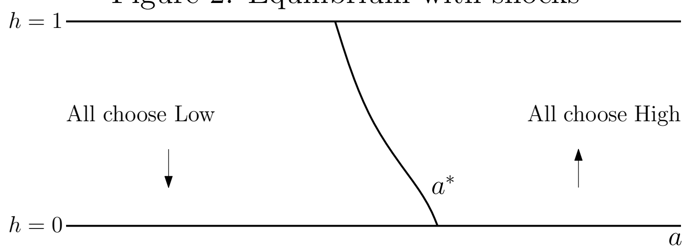

Let $V(a, h, \tilde{a})$ be the utility gain from choosing *High* obtained by an agent in state (a, h) that believes others will play according to threshold $\tilde{a}$. Then

$$V(a, h, \tilde{a}) = \int_0^\infty e^{-(\rho+\alpha)t} E[\pi(h_t, a_t)|a, h, \tilde{a}]dt - \psi, \tag{10}$$

where $E[\pi(h_t, a_t)|a, h, \tilde{a}]$ denotes the expectation of $\pi(h_t, a_t)$ of an agent in state (a, h) that believes others will play according to a threshold $\tilde{a}$. An agent choosing when $a = a^*(h)$ and

[12] We have shown instantaneous payoffs are increasing in a and h and an argument similar to the proof of Proposition 1 can be used to show that investing is a dominant action for high enough a and a dominated action for low enough a. We will show the existence of dominance regions for a more general process for a_t in Proposition 5, so the proof of Proposition 2 is omitted.

believing all others will play according to the cutoff a^* is indifferent between *High* and *Low*, which means that $V(a^*(h), h, a^*) = 0$, for every h.

4 The planner's problem

Proposition 2 shows there is a unique rationalizable equilibrium in the model. Although agents face a dynamic coordination problem, a unique set of rationalizable beliefs emerges and, from the point of view of an individual firm, pins down the optimal decision. It is then natural to ask about the beliefs that arise in equilibrium and, in particular, about the inefficiencies that might exist in the model.

One particularly important question is about whether inefficiencies are more pronounced at times of low economic activity. For instance, suppose the current value of h is low and the economy is at the left but close to the equilibrium threshold. A firm chooses not to invest and the economy might be stuck in a regime with low economic activity for a while. Is that situation particularly inefficient? Would a social planner be particularly interested in stimulating investment when h is low?

The planner maximizes expected welfare, given by:

$$E_\tau(W) = E_\tau \int_\tau^\infty e^{-\rho(t-\tau)} \left(Y(h_t, a_t) - \alpha\psi I(t)\right) dt \tag{11}$$

where $Y(h, a)$ is given by (7) and $I(t) \in [0, 1]$ is the decision of the planner about investing at time t.

The path of a is exogenously given and the path of h depends on future decisions of the planner, which is taken as given by the planner at a certain point in time. The planner chooses investment $I(\tau)$ at every point in time, which affects h in the following way: investing dI today raises h by αdI, but that increase depreciates at rate α. Hence

$$\frac{dh_t}{dI(\tau)} = \alpha e^{-\alpha(t-\tau)}$$

The first order condition of $E_\tau(W)$ with respect to investment $I(\tau)$ at a given time τ implies that the planner is indifferent between any level of investment if:

$$\int_\tau^\infty e^{-\rho(t-\tau)} E_\tau \left(\frac{\partial Y(h_t, a_t)}{\partial h} \alpha e^{-\alpha(t-\tau)} \right) dt - \alpha\psi = 0$$

Since

$$\frac{\partial Y(h_t, a_t)}{\partial h} = e^{a_t} \frac{\theta}{\theta - 1} \left(h_t x_H^{\frac{\theta-1}{\theta}} + (1 - h_t) x_L^{\frac{\theta-1}{\theta}} \right)^{\frac{1}{\theta-1}} \left(x_H^{\frac{\theta-1}{\theta}} - x_L^{\frac{\theta-1}{\theta}} \right)$$

we get that the planner chooses to invest at time τ ($I(\tau) = 1$) if:

$$\int_{\tau}^{\infty} e^{-(\rho+\alpha)(t-\tau)} E_{\tau} \left[\frac{\theta}{\theta-1} \pi(h_t, a_t) \right] dt \geq \psi \tag{12}$$

where $\pi(h_t, a_t)$ is given by (8). That is a necessary condition for optimality. In principle, it is difficult to characterize the planner's solution because expectations about the path of (a, h) have to be taken into account in the solution for the optimal decision but the path of h will be optimally chosen by the planner.

However, mathematically, this problem is very similar to the agent's problem in the decentralized equilibrium. At every point in time, there is investment if (12) holds, taking into account that the path of h in the future will be determined by a similar choice. The only difference is that the planner and agents follow different decision rules, but even that difference is small: the expression for the planner's choice in (12) and the expression for a firm's decision in (9) differ only by the term $\theta/(\theta-1)$ multiplying the benefit from investing in (12). Hence we know the planner also chooses according to a threshold a_P^* such that (12) holds with equality at $a_P^*(h)$ for $h \in [0, 1]$.

The key implications for the optimal stimulus policies are in Proposition 3:

Proposition 3. *Optimal policy:*

1. *[Optimality of a constant subsidy] The planner's solution can be implemented by a constant subsidy of ψ/θ whenever an agent invests.*

2. *[Parallel shift of the threshold] The planner invests according to a threshold a_P^* such that for any $h \in [0, 1]$,*

$$a_P^*(h) = a^*(h) - \log\left(\frac{\theta}{\theta-1}\right)$$

where a^* is the threshold for the decentralized equilibrium.

Proof. First statement: The solution to the planner's problem prescribes investment if (and only if) the condition in (12) is satisfied. Multiplying both sides of (12) by $(\theta-1)/\theta$ yields the condition for an agent to invest in (10) in an economy where the cost for investing is $\psi - \psi/\theta$.

Second statement: Since $\pi(h_t, a_t)$ can be written as $e^{a_t} g(h_t)$, for some function $g(\cdot)$, we can rewrite condition 12 as

$$\int_{\tau}^{\infty} e^{-(\rho+\alpha)(t-\tau)} E_{\tau} \left[e^{\left(a_t + \log\left(\frac{\theta}{\theta-1}\right)\right)} g(h_t) \right] dt \geq \psi \tag{13}$$

Define $b_t = a_t + \log\left(\frac{\theta}{\theta-1}\right)$ and consider the planners' problem in the (b, h)-space. The expression for the planner's decisions is identical to the expression in (10) for the agents' decisions in the decentralized equilibrium (in the (a, h)-space). Moreover, the law of motion for b_t is exactly the same as the law of motion for a_t. Therefore, the solution for the problem must be the same as well.

We know there is a unique decentralized equilibrium given by a threshold a^*, hence $a^* = b^*$, which implies $a^*(h) = a^*_P(h) + \log\left(\frac{\theta}{\theta-1}\right)$ and yields the claim.

□

Figure 3: Planner's problem

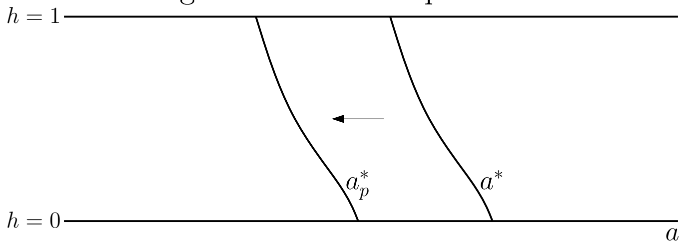

The solution to the planner's problem in (12) considers the benefits from investing by an individual producer multiplied by a constant larger than 1. Hence the only problem with the individual decision is that it requires a benefit from investing that is too high. A constant subsidy takes care of this problem.

As shown in (13), the expressions for the planner's problem in (12) is the same as the solution for the decentralized problem in (10) when a constant is added to the log of productivity. That implies the planner's threshold is a translation of the equilibrium threshold, where that constant is subtracted from the productivity threshold, as in Figure 3.[13]

The slope of the threshold affects the likelihood of a recession and its expected duration. If the threshold is close to a vertical line, h will start to fall when productivity is below some $a^\dagger$ but firms will resume investing whenever $a_t > a^\dagger$. A rotation of the threshold that reduces $a^*(1)$ but raises $a^*(0)$ implies there will be less ocasions where productivity will cross the threshold to the left of a^* when h is large, but when that happens, h is likely to go further down and it will take longer for h to increase again.

An implication of Proposition 3 is that a constant subsidy implements the planner's solution if and only if the planner is not willing to affect the slope of the equilibrium threshold. The proposition shows the planner is not concerned with the timing of stimulus

[13]This analysis ignores the costs of subsidizing investment but the result is robust to the inclusion of such costs. More specifically, assume the planner faces a monitoring cost c for each unit of investment it subsidizes, with $c < \psi/(\theta - 1)$. An argument similar to the one in Proposition 3 shows that the optimal policy can be implemented by a constant subsidy and leads to a different translation of the threshold, but no rotation. In Section 6.2 we assume that every unit of subsidy has a small welfare cost and obtain similar results.

policies: there is no special reason for subsidies at times of low (or high) economic activity (h); and there is no reason to affect the expected duration of recessions.

The planner's solution prescribes no extra stimulus for investment when h is low because equilibrium beliefs solve the dynamic coordination problem. The next section provides economic intuition for this result.

5 Understanding the result

Suppose the economy is at $h = 0$ and at the left but very close to the equilibrium threshold. Agents are stuck in a regime with low economic activity but would choose to invest if the economy switched to a high-h state. The planner could drive the economy to higher values of h but, as we've shown, this dynamic coordination problem does not provide a reason for stimulus at low h. Why?

The explanation can be divided in two parts: (i) the externality from investing is always a constant proportion of the individual return; and (ii) equilibrium beliefs compensate for the dynamic coordination problem. For the first part, note that welfare in this economy at time t can be written as $W(h,a) = hw_H(h,a) + (1-h)w_L(h,a)$. That yields:

$$\frac{\partial W}{\partial h} = (w_H(h,a) - w_L(h,a)) + \left(h\frac{\partial w_H(h,a)}{\partial h} + (1-h)\frac{\partial w_L(h,a)}{\partial h}\right)$$

An agent takes into account the effect of investment on her income but not the positive effect on others from selling at a lower price. However, owing to the constant demand elasticity,

$$\left(h\frac{\partial w_H(h,a)}{\partial h} + (1-h)\frac{\partial w_L(h,a)}{\partial h}\right) = \frac{1}{\theta-1}\left(w_H(h,a) - w_L(h,a)\right)$$

The externality is a constant proportion of the agent's payoff from investing and, consequently, independent of h or a for a given return to investment. Non-proportional externalities in payoffs may generate non-constant optimal subsidies but for reasons unrelated to the dynamic coordination problem. Sections 5.4 and 5.5 discuss how the model is affected when the externality depends on h.

For the second (and most important) part, we need to understand the beliefs that arise in equilibrium. It is instructive to look first at the case with no shocks. In this case, there are multiple equilibria, so beliefs outside the dominance regions are not determined by the model. We then move to a case with very small shocks, where a unique set of beliefs is pinned down by the model, and show how that affects the equilibrium.

5.1 Solutions for $\sigma = 0$ and $\sigma \to 0_+$

The uniqueness result and the expressions for the equilibrium and planner's thresholds hold for any $\sigma > 0$. In case $\sigma \to 0_+$, the expression for the threshold can be simplified. For any $h_0 \in [0, 1]$, suppose the economy is at the threshold, i.e., at $(a^*(h_0), h_0)$. The economy will soon move either up or down, but how exactly does it work and what are the probabilities the economy will go in each direction? This mathematical problem is studied by Burdzy et al. (1998) and their main result is that the economy will instantaneously move up in the direction of $(a^*(h_0), 1)$ with probability $1 - h_0$ and will move down in the direction of $(a^*(h_0), 0)$ with probability h_0. Using (10), $a^*_0 \equiv a^*(h_0)$ thus solves:

$$(1 - h_0) \int_0^\infty e^{-(\rho+\alpha)t}[\pi(h_t^\uparrow, a_0^*)]dt + h_0 \int_0^\infty e^{-(\rho+\alpha)t}[\pi(h_t^\downarrow, a_0^*)]dt = \psi \tag{14}$$

where $h_t^\uparrow = 1 - (1 - h_0)e^{-\alpha t}$, and $h_t^\downarrow = h_0 e^{-\alpha t}$. The solution to the planner's problem is similar. Using (12), we get that for any $h_0 \in [0, 1]$, the planner's threshold $a^*_{P0} \equiv a^*_P(h_0)$ solves:

$$(1 - h_0) \int_0^\infty e^{-(\rho+\alpha)t}[\pi(h_t^\uparrow, a_{P0}^*)]dt + h_0 \int_0^\infty e^{-(\rho+\alpha)t}[\pi(h_t^\downarrow, a_{P0}^*)]dt = \psi - \frac{\psi}{\theta} \tag{15}$$

For the planner, there is no difference between no shocks ($\sigma = 0$) or vanishing shocks ($\sigma \to 0_+$). Beliefs about the future are basically the same and the planner can effectively choose the path of h.[14]

However, the decentralized equilibrium in case $\sigma = 0$ can be very different from (14). In case of optimistic beliefs, the agent's threshold a^*_{opt} solves:

$$\int_0^\infty e^{-(\rho+\alpha)t}[\pi(h_t^\uparrow, a_{opt}^*(h_0))]dt = \psi \tag{16}$$

for every h_0, while in case of pessimistic beliefs, the agent's threshold a^*_{pes} is given by:

$$\int_0^\infty e^{-(\rho+\alpha)t}[\pi(h_t^\downarrow, a_{pes}^*(h_0))]dt = \psi \tag{17}$$

The irrelevance of vanishing shocks for the planner's solution highlights the point that

[14] The planner's problem with $\sigma = 0$ can be written in a different way: the planner chooses between always investing and never investing in the foreseeable future (any other option is dominated by one of these alternatives). Thus a^*_{P0} solves:

$$\int_0^\infty e^{-\rho t}\left(\left[h_t^\uparrow w_H(a_{P0}^*, h_t^\uparrow) + (1 - h_t^\uparrow) w_L(a_{P0}^*, h_t^\uparrow) - h_t^\uparrow \tilde{\psi}\right] - \left[h_t^\downarrow w_H(a_{P0}^*, h_t^\downarrow) + (1 - h_t^\downarrow) w_L(a_{P0}^*, h_t^\downarrow) - h_t^\downarrow \tilde{\psi}\right]\right) dt = 0$$

where $\tilde{\psi} = (\rho + \alpha)\psi$ is the user cost of capital. This expression seems very different from (15) but yields the same results.

very small fluctuations in a are not intrinsically important. Their effects on the decentralized equilibrium stem from the determination of beliefs in case $\sigma \to 0_+$.

5.2 The case with no shocks

Consider the case $\sigma = 0$. The equilibria of the model are depicted in Figure 4. The left threshold (good equilibrium) is the set of (a, h) where an agent is indifferent between investing or not assuming all others will invest given by (16). For the right threshold (bad equilibrium), the assumption is that no other agent will ever choose to invest, as in (17).

Figure 4: The case with no shocks

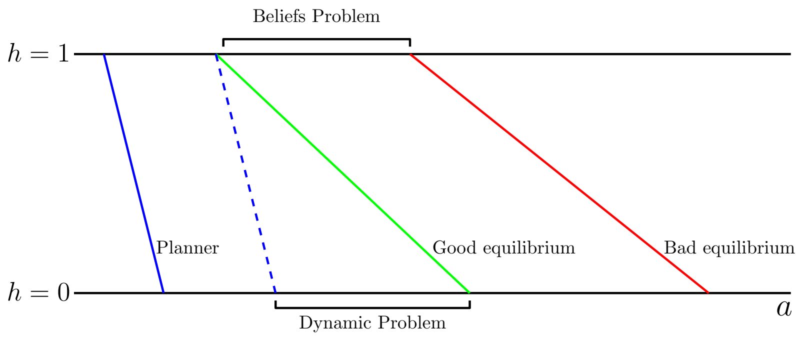

There are three differences between the planner's solution and the decentralized equilibrium: (i) the solution for the planner's problem is unique but there are multiple self-fulfilling equilibria; (ii) the planner takes into account the price externality from market power; and (iii) in the decentralized equilibrium, agents face a dynamic coordination problem: in a region of parameters, nobody wants to be the first to invest although investing is the socially optimal choice.

There are multiple equilibria in a region of parameters (the first difference between the planner's solution and the decentralized economy). However, regardless of whether we assume optimistic or pessimistic beliefs, the equilibrium threshold is further away from the planner's threshold for low values of h. The difference in slopes shows the planner would be willing to pay higher subsidies when h is low – as shown in Proposition 3, a constant subsidy would shift the threshold to the left without rotating it. That is the third difference between the planner's solution and the decentralized economy.

At the heart of this problem, there is a dynamic coordination failure. Agents effectively discount the future at rate $\alpha + \rho$ while the planner discounts the future at the much smaller rate ρ. In a region of parameters with low h and a just below the threshold, investment in the short run would drive the economy to the high regime and the planner takes that into account. Agents don't find it profitable to invest while demand is still low, but would be happy to sign a contract forcing everyone to take investment opportunities in the short run, as these losses for some would imply gains for all in the future.

While the planner effectively decides where the economy goes, agents take beliefs as given. In consequence, the economy might be stuck in a recession trap (low h, not so low a). However, beliefs are exogenous in this reasoning. We now consider the case of a positive but very small σ to understand which beliefs arise in equilibrium.

5.3 The case with very small shocks

We now consider the case with $\sigma \to 0_+$. Differently from the previous case, beliefs now are uniquely determined by the model. The equilibrium is depicted in Figure 5.

Figure 5: The case with very small shocks

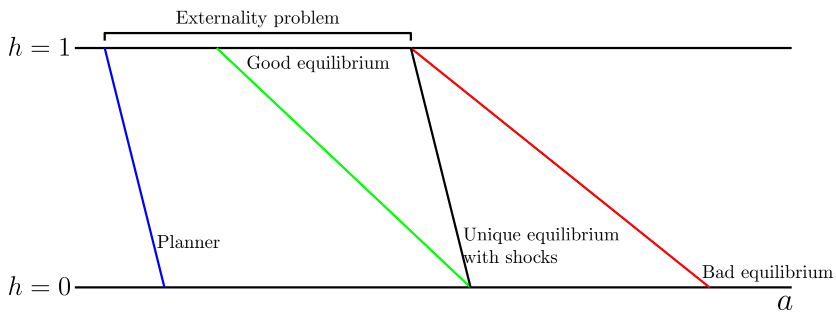

The result is completely different from the case with $\sigma = 0$. As in Frankel and Pauzner (2000), with very small but positive shocks, there is no role for arbitrary beliefs and no multiplicity of equilibria, so the first difference between the planner's problem and the decentralized equilibrium disappears. The surprising result from this paper is that the difference in slopes between the planner's and the agents' thresholds vanishes as well.

Since the productivity parameter moves very slowly, this change in the behavior of the economy must stem from the endogeneity of beliefs. As shown in Figure 5, the equilibrium threshold at $h = 0$ coincides with the 'good' equilibrium of the model with no shocks, hence agents are very optimistic at $h = 0$. The same reasoning also implies that agents are very pessimistic at $h = 1$.[15]

Intuitively, in a neighborhood of the equilibrium threshold, optimistic beliefs make perfect sense at $h = 0$, but no sense whatsoever at $h = 1$. Suppose the economy is exactly at the equilibrium threshold at $h = 0$. The economy would stay around there as long as $a < a^*(0)$, but any shock that moves a above $a^*(0)$ leads agents to invest and drives the

[15] An implication of Burdzy et al. (1998) is that in case $\sigma \to 0_+$, the equilibrium threshold connects the 'good-equilibrium' threshold (for the case $\sigma = 0$) at $h = 0$ and the 'bad-equilibrium' threshold at $h = 1$ as in Figure 5.

economy up in the picture. Since the slope of the threshold is negative, as soon as the economy is at $h > 0$, it is at the right of the threshold and hence moves up in the direction of $h = 1$. The fundamental asymmetry here is that a tiny negative shock basically leaves the economy where it is, while a tiny positive shock drives the economy up in the direction of $h = 1$. The same reasoning implies that in a neighborhood of the equilibrium threshold at $h = 1$, a regime switch is also expected and beliefs are pessimistic.

The intuition for equilibrium beliefs was explained for the case with $\sigma \to 0_+$ but Proposition 3 shows the result holds for any $\sigma > 0$. Graphically, for σ bounded away from 0, the planner's and agents' equilibrium thresholds would be parallel to each other as in Figure 5, but the agents' threshold would not touch the boundaries of the dominance regions.

The intuition for the case with σ bounded away from zero is very similar. When economic activity is low and agents are around the equilibrium threshold, they are optimistic: productivity is relatively good, the economy is parked in a region of inaction but is likely to leave that state soon. Again, the key asymmetry here is that a movement of a to the left does not significantly affect the state of the economy, but a movement of a to the right affects the mass of agents investing, raising demand in the economy and incentives for the following agents with investment opportunities to take them. The recovery is just waiting for a small piece of good news.

The only difference between the planner's solution and the decentralized equilibrium is the monopoly distortion in the investment decision, the externality shown in Figure 5. When beliefs are determined by the model, planner and agents solve a very similar problem. At every (a, h), investment is undertaken if its expected return pays off and the equilibrium (or planner's) threshold is a fixed point. "Pays off" means different things for agents and planner but the ratio is constant since the only difference is the externality from market power.

5.4 What if the externality is not a constant fraction of payoffs?

Different functional forms for preferences and technology could in principle lead to externalities varying with h. We now argue that for reasonable parameters (namely for $\rho << \alpha$), the effect of, say, larger externalities for low values of h would be mitigated in equilibrium.

Consider the case $\sigma \to 0_+$, assume the private gain from investing is $\pi = e^a g(h)$ and the planner's instantaneous gain is:

$$\Lambda(h, a) \equiv [1 + \kappa] \, e^a g(h) + \beta(h) \quad \text{with} \quad \int_0^1 \beta(h) dh = 0$$

The function $\beta(h)$ captures deviations of the planner's threshold from a parallel shift of the equilibrium threshold. Using Theorem 2 in Frankel and Pauzner (2000), the expression for the planner's threshold in (15) yields:

$$a_P(0) = \log\left(\frac{\alpha\psi}{\int_0^1 (1-h)^{\rho/\alpha} e^a (1+\kappa) g(h) dh + \int_0^1 (1-h)^{\rho/\alpha} \beta(h) dh}\right) \tag{18}$$

and

$$a_P(1) = \log\left(\frac{\alpha\psi}{\int_0^1 h^{\rho/\alpha} e^a (1+\kappa) g(h) dh + \int_0^1 h^{\rho/\alpha} \beta(h) dh}\right) \tag{19}$$

The second integral term in the denominator of both (18) and (19) corresponds to deviations from a parallel shift of the threshold – if that term is 0, deviations from the average externality do not show up in the equations. In general, the second integral terms in (18) and (19) will be small for two reasons: (i) $\beta(h)$ is the difference between the externality at a particular value of h and the mean externality (across h), while $e^a(1+\kappa)g(h)$ is the private gain plus the average externality, which would in general be much larger; and (ii) ρ/α tends to be small since ρ is the time discount rate (say 2% a year) and α is the frequency agents choose investment (say twice a year). Hence there will be little dispersion on weights attached to each $\beta(h)$ and the second integral will therefore be close to 0.

The second point is particularly related to our main contribution, so it is worth exploring it with an illustration. We now consider a linear example so that $\beta(h) = \zeta - 2\zeta h$, hence the externality is larger at $h = 0$ by ζ. In this case

$$\int_0^1 (1-h)^{\rho/\alpha} \beta(h) dh = \frac{\zeta\rho/\alpha}{(1+\rho/\alpha)(2+\rho/\alpha)} < \frac{\zeta}{2}\frac{\rho}{\alpha}$$

In words, the "extra" externality at $h = 0$ has to be divided by two and multiplied by ρ/α. This term is then added to the denominator of the expression for $a_{SP}(0)$ in (18) (and a similar term is subtracted from the expression for $a_{SP}(1)$). Since the time discount rate ρ is in general much smaller than the frequency of decisions α, only a small fraction of the "extra" externality at $h = 0$ is considered in the expression for $a_{SP}(0)$. Intuitively, when deciding on investment at the equilibrium threshold at $h = 0$, the planner takes into account that the economy will soon be moving up, so the "local" externality is not that important.

5.5 The case with vanishing frictions

In case of vanishing frictions ($\alpha \to \infty$), the economy moves very fast from $h = 0$ to $h = 1$ but agents' horizons become very short.[16] In the model with no shocks, agents take h into account in their decisions, regardless of whether we assume fully optimistic or fully pessimistic beliefs. In contrast, the planner does not take h into consideration. Moving the

[16]When taking the limit $\alpha \to \infty$, we fix the user cost of capital $(\rho + \alpha)\psi$, not the investment cost ψ.

economy to a different regime takes very little time and hence the transition is unimportant. The threshold from the planner's problem converges to a vertical line, as in Figure 6. This result holds for any $\sigma > 0$.

Figure 6: The case with vanishing frictions

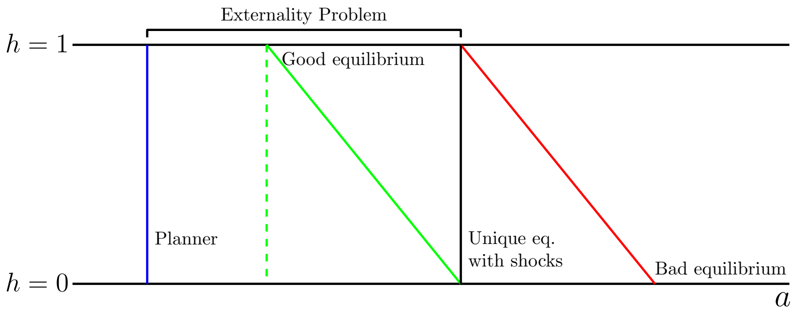

An implication of Proposition 3 is that in equilibrium agents also play according to a vertical threshold. History thus becomes irrelevant.[17] Interestingly, this result holds even in case of non-proportional externalities as discussed in Section 5.4. When $\alpha \to \infty$, the second integral in the denominators of (18) and (19) converge to zero, hence the differences in the externality from the High regime across h do not affect the planner's or agents' decisions. Hence for large α/ρ, the optimal policy prescribes an approximately constant subsidy even if externalities are larger when h is low.

6 Extensions

6.1 A macroeconomic model with labor

The underlying macroeconomic model presented in Section 2 is quite stylized. Firms cannot adjust their production in the intensive margin when economic conditions change. In this section we develop a standard macroeconomic model where firms still receive opportunities to increase their productivity according to a Poisson process, but can adjust the amount of labor used in production at each point in time. The qualitative results in Proposition 3 are

[17] Intuitively, for a large α, an agent at the equilibrium threshold and $h = 0$ knows the economy will move up with probability 1, while an agent at the equilibrium threshold and $h = 1$ is sure the economy will move down. They don't know their 'position in the queue', i.e., when they will have the next opportunity for revising their behavior: how much time will have elapsed and the value of h when they can choose again. On the one hand, the agent at $h = 0$ will experience lower values of h than the agent at $h = 1$ – the agent starting at $h = 0$ is likely to get an opportunity to revise behavior before h gets close to 1. On the other hand, for the agent at $h = 0$, the economy moves up very quickly at lower values of h, but slowly as h approaches 1, so the last firms to change their decision will spend relatively more time at high values of h. As it turns out, both effects exactly cancel each other, so the agents at $h = 0$ and $h = 1$ are indifferent between investing and not investing for the same value of a.

unchanged, i.e., the constant subsidy is still optimal and the planner chooses a threshold that is parallel to the agent's threshold.

There is a continuum of intermediate good firms indexed by $i \in [0,1]$ and a representative household who supplies labor and consume a final good. The final good is produced by a competitive firm whose production function is the Dixit-Stiglitz aggregator of intermediate goods y_{it}, given by equation (1) as before. The price of the final good is P_t. The household utility function is given by

$$\mathcal{U}(C_t, L_t) = C_t - \frac{1}{\gamma+1} L_t^{\gamma+1},$$

where C_t is the amount consumed of the final good, L_t is the total amount of labor supplied and $\gamma \geq 0$ parameterizes the Frisch elasticity of the labor supply. The household budget constraint is

$$P_t C_t \leq w_t L_t + \int_0^1 \Pi_{it} di,$$

where w_t is the wage and Π_{it} is the profit of firm i (the household owns all the firms).

As before, at a given point in time, intermediate good firms can be in either the *High* (H) or the *Low* (L) regime and opportunities to switch across regimes follow the same Poisson technology as in Section 2. Firms use labor to produce their differentiated goods. A firm i in regime $r \in \{L, H\}$ has production function

$$y_{it} = A_t X_r l_{it}^{\lambda},$$

where l_{it} denotes the amount of labor used in production, $\lambda \in (0,1)$ and $X_H > X_L$. Intermediate good firms maximize profit taking the demand schedule imposed by final goods firms as given. The stochastic process of $a_t = \log(A_t)$ is the same as in Section 2.

6.1.1 The agent's problem

We will start solving for the equilibrium in a given period t taking the proportion of firms in each regime as given. Since firms in a given regime and at a point in time make the same choices in equilibrium, we will abuse notation and index firms by the regime $r \in \{L, H\}$. We will also omit time subscripts in this subsection.

The demand schedule of the final good firm is the same as before and is given by (4). Therefore, we can write the profits of a firm in regime r as

$$\Pi_r = y_r^{\frac{\theta-1}{\theta}} Y^{\frac{1}{\theta}} - w \left(\frac{y_r}{AX_r} \right)^{\frac{1}{\lambda}}. \tag{20}$$

For a given wage w and final good output Y, the optimal choice of a firm in regime r is:

$$y_r = \left[\left(\frac{\theta}{\theta - 1} \right) \frac{1}{\lambda} \right]^{-\theta\lambda\zeta} (AX_r)^{\theta\zeta} Y^{\lambda\zeta} w^{-\theta\lambda\zeta}, \tag{21}$$

where $\zeta \equiv \frac{1}{\lambda+\theta(1-\lambda)}$. A firm producing y_r will demand $l_r = [y_r / (AX_r)]^{1/\lambda}$ units of labor. For a given wage w, the household will supply $w^{1/\gamma}$ units of labor. Thus, equilibrium in the labor market implies:

$$w = \left[h \left(\frac{y_H}{AX_H} \right)^{\frac{1}{\lambda}} + (1-h) \left(\frac{y_L}{AX_L} \right)^{\frac{1}{\lambda}} \right]^{\gamma}. \tag{22}$$

Finally, goods markets must clear:

$$Y = \left(h y_H^{\frac{\theta-1}{\theta}} + (1-h) y_L^{\frac{\theta-1}{\theta}} \right)^{\frac{\theta}{\theta-1}}. \tag{23}$$

Equations (21), (22) and (23) pin down the equilibrium of this economy for a given state (a, h). From these equations, we find equilibrium expressions for the final good production $Y(h, a)$, the production for firms in each regime $y_H(h, a)$ and $y_L(h, a)$, and the wage $w(h, a)$. These lead to expressions for hours worked $L(h, a)$ and profits $\Pi_H(h, a)$ and $\Pi_L(h, a)$.

The final good production in equilibrium is given by:

$$Y(h, a) = \left[\left(\frac{\theta}{\theta - 1} \right) \frac{1}{\lambda} \right]^{\frac{\lambda}{\lambda-\gamma-1}} \exp\left\{ \left(\frac{\gamma+1}{\gamma+1-\lambda} \right) a \right\} q(h)^{\frac{\lambda\gamma\delta(\theta-1)-1}{(\gamma+1)\delta-1}}, \tag{24}$$

where $q(h) \equiv \left[h (X_H)^{(\theta-1)\zeta} + (1-h) (X_L)^{(\theta-1)\zeta} \right]$ and $\delta \equiv \frac{1}{(1+\gamma)\theta-(\theta-1)\lambda}$. The difference in profits between firms in regime *High* and *Low*:

$$\tilde{\pi}(h, a) = \exp\left\{ \left(\frac{\gamma+1}{\gamma+1-\lambda} \right) a \right\} q(h)^{\frac{\lambda\gamma(\theta-1)-(\gamma+1)}{(\theta-1)(\lambda-\gamma-1)}} \left(X_H^{(\theta-1)\zeta} - X_L^{(\theta-1)\zeta} \right) \Phi, \tag{25}$$

where $\Phi \equiv \left[\left(\frac{\theta}{\theta-1} \right) \frac{1}{\lambda} \right]^{\frac{\lambda}{\lambda-\gamma-1}} - \left[\left(\frac{\theta}{\theta-1} \right) \frac{1}{\lambda} \right]^{\frac{\gamma+1}{\lambda-\gamma-1}}$.

It is no longer true that the gain of being in regime *High* above is always increasing in h. We need to impose the following restrictions on parameters:[18]

$$\lambda (\theta - 1) - 1 < 1/\gamma \tag{26}$$

This restriction is satisfied for standard parameters in the literature.[19] Notice that the Frisch

[18]If $\gamma = 0$ we cannot write this assumption this way, but the condition is always satisfied.

[19]For example, the restriction is satisfied with $\theta = 5$, $\lambda = 2/3$ and $\gamma = 1/3$. For a recent discussion about the Frisch elasticity $1/\gamma$, see Peterman (2014).

elasticity $1/\gamma$ cannot be too low relative to θ, which determines the intensity of demand externalities in the economy. As in the model of Section 2, when others produce more, an agent is more inclined to incur the fixed cost for the *High* regime, owing to the increase in price of his variety. However, in the model with labor, there is also an effect in the opposite direction: when others are in regime *High* and producing more, wages increase as well. If the Frisch elasticity is high, then the effect on wages is small and the overall incentives to be in regime *High* increase in h. Finally, λ also enters the expression since it affects the difference in labor used by firms in each regime.

Given that the assumption in (26) is satisfied, the equilibrium of the dynamic game is given as in Section 3, with $\pi(h, a)$ replaced by $\tilde{\pi}(h, a)$ from (25).

6.1.2 The planner's problem

The planner seeks to maximize the discounted sum of flow payoffs of the representative household, which is given by $\mathcal{U}^{SP} = Y - \frac{1}{\gamma+1} L^{\gamma+1}$. There are two different ways to consider the planner's problem: (i) the planner chooses whether firms invest and how much labor they hire (and hence how much they produce); or (ii) the planner will not interfere in firms' pricing decisions and in workers' labor supply, will only choose investment. Since we are concerned with policies to stimulate investment, we now focus on the second alternative, but the first one is presented in the appendix and the results are essentially the same as in Proposition 3.

The planner thus determines the choices of firms between the *High* and *Low* regimes, taking as given what firms and workers will choose at every point in time. The flow utility of the representative household can thus be rewritten as

$$\mathcal{U}^{SP}(h, a) = h\Pi_H(h, a) + (1 - h)\Pi_L(h, a) + \frac{\gamma}{\gamma + 1} w(h, a)^{\frac{\gamma+1}{\gamma}}. \tag{27}$$

where functions $Y(h, a)$, $L(h, a)$, $\Pi_H(h, a)$, $\Pi_L(h, a)$ and $w(h, a)$ refer to the decentralized equilibrium of this economy for a given h and a, as defined in the last subsection.

Repeating the steps leading to (12), we get that the solution to the planner's problem can be written as the solution to the decentralized equilibrium replacing the private gain from being in the *High* regime $\tilde{\pi}(h, a)$ by $\frac{\partial \mathcal{U}^{SP}}{\partial h}$. Taking partial derivatives of (27) with respect to h, using (22) and the equality between marginal revenue and marginal cost, and rearranging, we get:

$$\frac{\partial \mathcal{U}^{SP}(h, a)}{\partial h} = \tilde{\pi}(h, a) + \frac{1}{\theta} \frac{\partial Y(h, a)}{\partial h}.$$

Taking derivatives of (24) and rearranging leads to:

$$\frac{\partial \mathcal{U}^{SP}}{\partial h} = \left(\frac{1}{1-\chi}\right)\tilde{\pi}(h,a),$$

where $\chi = \frac{1}{\theta - \frac{\lambda}{\gamma+1}(\theta-1)}$.

Following the steps in Section 4, we can show that a constant subsidy $\chi\psi$ implements the planner's first best.[20] Moreover, the planner's threshold is parallel to the agent's threshold and given by

$$\tilde{a}_P^*(h) = \tilde{a}^*(h) - \left(\frac{\gamma+1-\lambda}{\gamma+1}\right)\log\left(\frac{1}{1-\chi}\right), \tag{28}$$

where $\tilde{a}^*(h)$ denotes the agent's threshold in the model of this section.

6.2 Implementation

Proposition 3 shows that a constant subsidy implements the first best. However, in a large set of states, much less generous subsidies would be enough to coax agents to invest. This leads to the following question: what if every unit of subsidy has a small welfare cost $\varepsilon \approx 0$? It is not difficult to show that the government will use minimal spending policies, as in Definition 1.

Definition 1. Let a^* be an equilibrium of the game and a_p^* a continuous function such that $a_p^*(h) < a^*(h)$, for every h. Let $\hat{a}$ be the boundary where an agent is indifferent between *High* and *Low* when others are playing according to a_p^*. The function $\varphi(h,a)$ is the *minimal spending policy* that implements a_p^* if

$$\varphi(h,a) = \begin{cases} \psi - \int_0^\infty e^{-(\rho+\alpha)t} E[\pi(h_t,a_t)|a,h,a_p^*]dt & \text{if } a_p^*(h) \le a \le \hat{a}(h) \\ 0 & \text{otherwise} \end{cases}. \tag{29}$$

Figure 7 shows 3 thresholds: a_p^* is the threshold implemented by the policy, $\hat{a}$ is the best response of a player that believes others will play according to a_p^* and a^* is the equilibrium threshold without intervention. By definition, a^* is the best response to others playing according to a^*. Now, the sheer change in beliefs affects agents' strategies: once they believe others will play according to a_p^*, they will be indifferent between *High* and *Low* at a threshold $\hat{a}$ such that $\hat{a}(h) < a^*(h)$ for all $h \in [0,1]$.

A government following a minimal spending policy is committed to give an investment subsidy to each agent in the region between a_p^* and $\hat{a}$ (the gray area in figure 7). The

[20]When the planner can affect not only investment but also labor and production at each t, the optimal investment subsidy is ψ/θ. Intuitively, when the only instrument available for the planner is an investment subsidy, the optimal subsidy is larger in order to partially offset the inefficiencies from low labor demand.

Figure 7: Example of minimal spending policy

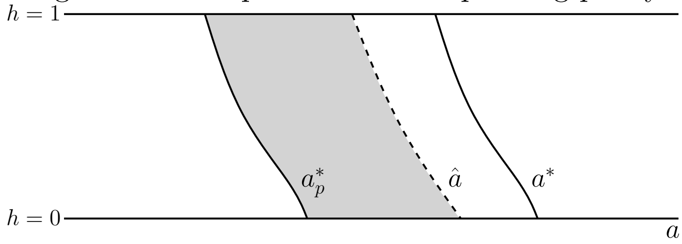

subsidy $\varphi(h, a)$ makes her indifferent between choosing *High* and *Low* given others will play according to a_p^*. Under those beliefs, playing according to a_p^* is a best response under this policy, so a_p^* is an equilibrium. Interestingly, no subsidies are needed in the area between $\hat{a}$ and a^*.[21]

Proposition 4 shows that minimum spending policies do not affect the main result of the paper.

Proposition 4. *In the model of Section 2 with minimal spending policies, the maximum optimal subsidy is ψ/θ for all $h \in [0, 1]$.*

Proof. See Appendix B. □

The result is intuitive. Under minimal spending policies, investment subsidies are equal to the minimum between how much the planner is willing to pay and how much is required to coax agents to invest. The result from Proposition 4 thus follows from the planner's willingness to subsidize being independent of h (Proposition 3).

6.3 Mean reversion

We now consider a_t follows a process with mean reversion. Let $a_t = \log(A_t)$ vary in time according to

$$da_t = \eta(\mu - a_t)dt + \sigma dZ_t \tag{30}$$

The parameter η determines how fast a_t returns to its mean, given by μ. Proposition 5 builds on Frankel and Pauzner (2000) to show that a threshold equilibrium always exists.

Proposition 5 (Existence)**.** *Suppose $\sigma > 0$. There exists a strictly decreasing function a^* such that a^* is an equilibrium.*

[21] The equilibrium under the minimal spending policy is no longer unique. If agents believe others will play according to a^* their best response is to play according to a^* as well, thus the policy has no effect at all. The amount of subsidies required to coax agents to invest depends on whether they expect others to respond to the stimulus policy. However, the government could implement this allocation through a contingent subsidy that would be essentially equivalent to a minimal spending policy: a large subsidy contingent on others not investing, and the subsidy prescribed by the minimal spending policy in case others invest as well.

Proof. See Appendix B. □

The first statement of Proposition 3 still holds when the process for a_t exhibits mean reversion, thus the planner's solution can be implemented by a constant subsidy. However, the second statement does not hold in this case because with mean reversion, translating the threshold is not isomorphic to re-labeling the a-axis. A translation of a^* also implies a different path of a_t and, consequently, a different balance between expected demand and productivity around that threshold. However, numerical results show that for a reasonable amount of mean reversion, the results are essentially unchanged.[22]

We do not have a strong uniqueness result in this case. However, the model can be seen as a limiting case of a sequence of models that have a unique rationalizable equilibrium. The following section discusses this point.

6.3.1 On equilibrium uniqueness

In order to apply the results of Frankel and Burdzy (2005), we need to make two changes in the model. First, the diffusion process for a_t is given by (30), but the mean-reversion parameter η_t varies over time so that

$$\eta_t = \begin{cases} \eta & \text{if } t < T \\ 0 & \text{otherwise} \end{cases}, \tag{31}$$

where T is a large number. Second, the difference between the instantaneous utility of agents locked in each regime is given by $\hat{\pi}$ instead of π, where

$$\hat{\pi}(h,a) = \begin{cases} \pi(h,a) & \text{if } a < M \\ \pi(h,M) & \text{otherwise} \end{cases}, \tag{32}$$

where M is a large number. One can verify that $\hat{\pi}(h,a)$ is Lipschitz in both a and h, and continuous. Using the results in Frankel and Burdzy (2005), we can prove there is a unique equilibrium in this model.

Proposition 6 (Uniqueness, Frankel and Burdzy (2005))**.** *Suppose $\sigma > 0$, the mean reversion parameter η_t is given by (31) and the relative payoff of investing is given by (32). Then there is a unique rationalizable equilibrium in the model. Agents follow cut-off strategies, and the cut-off can vary over time.*

Proof. See Appendix B. □

[22] See Guimaraes and Machado (2013).

As M and T approach infinity, this modified model converges to our model. For finite values of M and T, the environment is not stationary anymore: the equilibrium strategies might vary over time. Nevertheless, agents' behavior at time 0 is determined by a threshold that makes agents indifferent between *High* and *Low*. For large values of M and T, that threshold is arbitrarily close to the function $\tilde{a}$ that makes the expression in (10) equal to zero.

Why does the mean reversion need to die out eventually? In case of no mean reversion ($\eta = 0$), the iterative procedure in Frankel and Pauzner (2000) could be applied to show equilibrium uniqueness. However, in the presence of mean reversion, the last step in the proof of Frankel and Pauzner (2000) fails. Their proof relies on finding two boundaries, $a^1(h) < a^2(h)$ for every $h \in [0, 1]$, with the same shape, such that: (i) in any equilibrium that survives iterative elimination of strictly dominated strategies, agents play *Low* whenever the economy is to the left of $a^1(h)$ and *High* if the economy is to the right of $a^2(h)$; and (ii) there exists $\hat{h} \in [0, 1]$ such that an agent B at $(a^1(\hat{h}), \hat{h})$ and agent C at $(a^2(\hat{h}), \hat{h})$ are indifferent between *High* and *Low*. Since $a^1(h) < a^2(h)$, for every $h \in [0, 1]$, it cannot be the case that both are indifferent because both expect the same dynamics for h_t given any realization of the Brownian motion, but C expect larger values of a_t (because $a^2(\hat{h}) > a^1(\hat{h})$). However, this argument fails when the process for a_t exhibits mean reversion. In order to see this, consider the case where $a^1(\hat{h}) < \mu$ and $a^2(\hat{h}) > \mu$. Now, C expects a_t to fall, while B expects a_t to rise. Although C still expects larger values of a_t for any realization of the Brownian motion, B expects better relative dynamics for a_t, which can imply a more optimistic expectation about the dynamics of h_t.

Frankel and Burdzy (2005) overcome this problem by transforming the space and time of the stochastic process a_t, so the difference in instantaneous utility of agents locked in each regime can be written as a function of an i.i.d. process and time. Then, we can follow a procedure that is similar to Frankel and Pauzner (2000) *for every date* t in a transformed time-and-fundamental space. However, technical complications arise when the mean reversion lasts forever. For a given time t, in the transformed time-and-fundamental space, we may not be able to find a translation of a boundary such that every agent at *every date* $\tau > t$ chooses *Low* (or *High*), that is, the region where no action is dominant keeps expanding in time in the transformed fundamental space.

7 Numerical example

We now calibrate and solve the model numerically in order to get a better understanding of the workings of the model. In order to solve the model numerically, we work with an approximation of the model presented in Section 2. Now time is discrete and each period has length Δ, where Δ is a small number. Hence time $t \in \{0, \Delta, 2\Delta, 3\Delta, ...\}$. We also consider a stochastic process with mean reversion, so that

$$a_t = a_{t-1} + \eta(\mu - a_{t-1})\Delta + \sigma\sqrt{\Delta}\varepsilon_t,$$

where ε_t is an iid shock with a standard normal distribution. In the beginning of each period, after a_t is observed, $(1 - e^{-\alpha\Delta})$ individuals are randomly selected and get a chance to switch regime. The instantaneous payoffs of being locked in each regime are the same as before, but now agents discount utility by the factor $e^{-\rho\Delta t}$. When $\Delta \to 0$, this model converges to the model of Section 2.

Our algorithm aims at finding a threshold where agents are indifferent between actions *High* and *Low* if they believe others will play according to that threshold. The steps are basically the following: first, choose a finite grid for h in the interval $[0, 1]$ and pick an arbitrary threshold $a_0^*(h)$ for every h in the grid. Then, for each h in the grid, simulate n paths of a_t and h_t departing from $(a_0^*(h), h)$ assuming every agent will play according to a_0^*. Use those paths to estimate the gain in utility from picking *High* of an agent choosing at $(a_0^*(h), h)$. That yields an estimate of $V(a_0^*(h), h, a_0^*)$. If the gain in utility is close to zero in every point of the grid, stop. Otherwise, update a_0^* and repeat the simulation process that leads to the estimation of $V(a_0^*(h), h, a_0^*)$ until it converges.[23]

Parameters were chosen to satisfy the following criteria:

- The mean of output in peaks is about 4% higher than in troughs, which is roughly consistent with the data using the two-quarters definition of business cycles.[24]

- The economy stays 30% of time at the left of the threshold, that is, agents are not investing 30% of the time, approximately.[25]

- Once the economy goes to the left of the threshold, the mean time it stays there is 5 quarters. We consider that the economy went to the left of the threshold if it crossed it and remained there for at least 36,5 days.[26]

[23]Alternatively, we can assume every agent is choosing *Low*, find the threshold that determines the region where playing *High* is a dominant strategy (call it a_0^H), then assume all agents play according to a_0^H, find again the best response and keep iterating until it converges. We can also start by assuming all agents play *High*, find the region where playing *Low* is dominant and start the iterative process of eliminating dominated strategies from there. As suggested by the proof of equilibrium uniqueness, both equilibrium thresholds and the one found using the first algorithm presented coincide, but the former ones are more expensive in terms of computing time.

[24]According to the two-quarters definition of business cycles, a recession starts when output goes down for two consecutive quarters and ends when it increases for two consecutive quarters.

[25]When the economy is to the left of the threshold, no agent is investing. If that is interpreted as a recession, this calibration implies the economy is in recession 30% of the time. Owing to the lack of a positive trend in our productivity parameter, output is increasing roughly 50% of the time.

[26]That is because it is not reasonable to consider an economy is in a recession if unemployment fell for 3 consecutive days.

Output is computed net of depreciation, so the present value of output is equivalent to the present value of consumption (and thus utility) in the economy. The user cost of capital for an agent locked in the *High* regime is equal to $(\rho + \alpha)\psi$. At time t, there are h_t agents in the *High* regime, so we subtract the cost of capital in the economy $h_t(\rho + \alpha)\psi$ from the total amount produced, given by (7).

The parameters μ and x_L were normalized to zero and one, respectively. The chosen values of the parameters θ and ρ are standard in the literature, and α was made equal to 1, meaning that investment decisions are made once a year on average. All other parameters in the model were chosen to match the desired statistics. Table 1 shows the parameters (the time unit is years, when needed).

Table 1: Parameters

Parameter	**Symbol**	**Value**
Production regime *High*	x_H	1.1
Production regime *Low*	x_L	1
Elasticity substitution	θ	6
Fixed cost of investing	ψ	0.0806
Mean of fundamental process	μ	0
Arrival rate of Poisson Process	α	1
Standard deviation of shocks	σ	0.03
Discount rate	ρ	0.03
Mean reversion intensity	η	0.7
Time interval lenght	Δ	0.005

Figure 8 shows the equilibrium threshold and the path of the economy following a random realization of a_t. At the left of the threshold, agents do not invest, so h decreases; at the right of the threshold, agents invest, so h increases. A point (a, h) describes the current state of the economy and, together with the equilibrium threshold, determines agents' expectations about the future. In this example, the economy starts to the right of the threshold at $(0, 0.5)$, so h initially increases. About a year later, negative shocks to a bring the economy to the left of the equilibrium threshold and h starts to decrease. At that point, it is optimal for agents to choose *Low* because they expect others will do so.

Figure 9 shows output in the economy and what output would be in case $h = 1$. The variance of output in this economy is about 20% higher relative to the case where h is always equal to 1, because low values of the productivity parameter a lead to periods of low expected demand where agents choose not to invest. In this model, policies can do nothing about the exogenous movements in a but can increase the region where agents invest. Investment subsidies can bring output closer to the $h = 1$ curve.

Besides amplifying the effects of negative shocks, the endogenous and staggered reaction

Figure 8: Estimated threshold

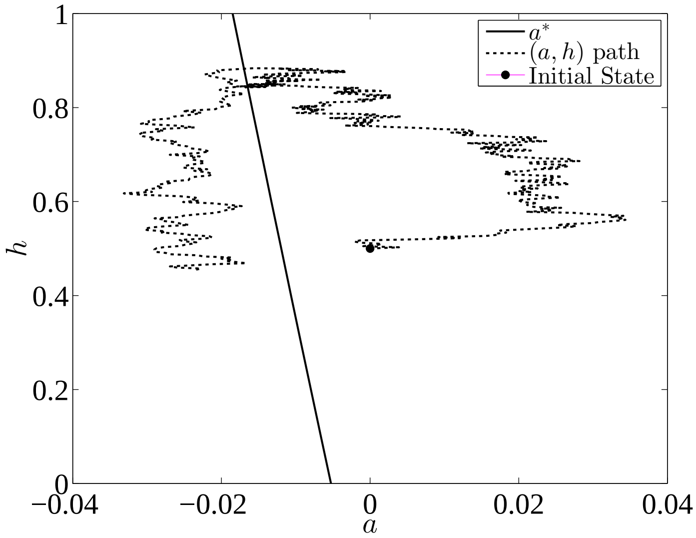

of h also implies that low productivity periods have long-lasting negative effects. As shown in Figure 9, output when the economy is coming back from a recession is lower than right before the recession for the same productivity parameter a. That occurs not only because staggered investment decisions mechanically add persistence to output, but also because agents require a higher productivity to invest when h is low.

8 Concluding remarks

This paper proposes a macroeconomic model that captures in a simple way the dynamic coordination problem arising from demand externalities and fixed investment costs. From a substantive point of view, the main result of the paper is the absence of a special reason for subsidies at times of low economic activity – a constant subsidy implements the planner's solution. From a methodological point of view, the paper highlights the importance of understanding beliefs that arise in equilibrium and their policy implications.

References

Angeletos, G.-M. and La'O, J. (2009). Incomplete information, higher-order beliefs and price inertia, *Journal of Monetary Economics* **56**: S19–S37.

Angeletos, G.-M. and La'O, J. (2010). Noisy business cycles, *NBER Macroeconomics Annual 2009*, Vol. 24, University of Chicago Press, pp. 319–378.

Angeletos, G.-M. and La'O, J. (2013). Sentiments, *Econometrica* **81**(2): 739–779.

Figure 9: Output fluctuations

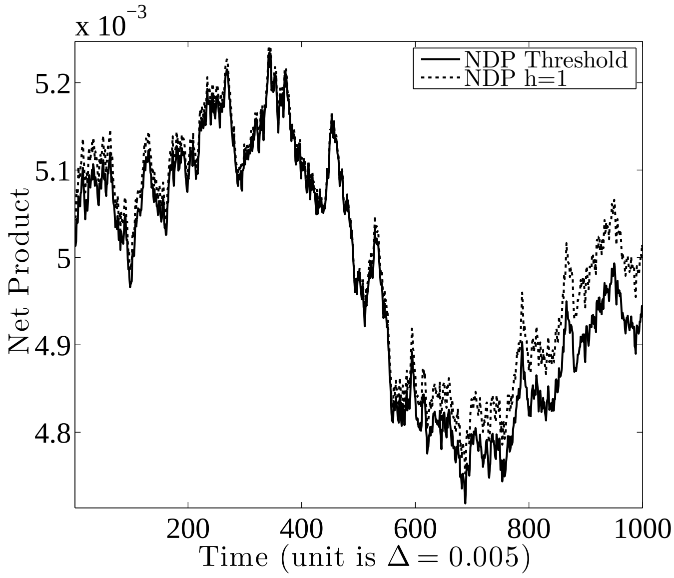

Arnold, B., Castillo, E. and Sarabia, J.-M. (1999). *Conditional Specification of Statistical Models*, Springer.

Beaudry, P. and Portier, F. (2006). Stock prices, news, and economic fluctuations, *The American Economic Review* **96**(4): 1293–1307.

Benhabib, J. and Farmer, R. E. (1994). Indeterminacy and increasing returns, *Journal of Economic Theory* **63**(1): 19–41.

Blanchard, O. and Kiyotaki, N. (1987). Monopolistic competition and the effects of aggregate demand, *American Economic Review* **77**(4): 647–66.

Burdzy, K., Frankel, D. and Pauzner, A. (1998). On the time and direction of stochastic bifurcation, *in* B. Szyszkowicz (ed.), *Asymptotic Methods in Probability and Statistics*, North-Holland, Amsterdam, pp. 483 – 500.

Burdzy, K., Frankel, D. and Pauzner, A. (2001). Fast equilibrium selection by rational players living in a changing world, *Econometrica* **69**(1): 163–189.

Carlsson, H. and Van Damme, E. (1993). Global games and equilibrium selection, *Econometrica* **61**(5): 989–1018.

Chamley, C. (2013). When demand creates its own supply: Savings traps, *Review of Economic Studies* **forthcoming**.

Cooper, R. and John, A. (1988). Coordinating coordination failures in keynesian models, *The Quarterly Journal of Economics* **103**(3): 441–463.

Cooper, R. W. and Haltiwanger, J. C. (2006). On the nature of capital adjustment costs, *The Review of Economic Studies* **73**(3): 611–633.

Eusepi, S. and Preston, B. (2011). Expectations, learning, and business cycle fluctuations, *American Economic Review* **101**(6): 2844–72.

Farmer, R. E. and Guo, J.-T. (1994). Real business cycles and the animal spirits hypothesis, *Journal of Economic Theory* **63**(1): 42–72.

Frankel, D. and Burdzy, K. (2005). Shocks and business cycles, *Advances in Theoretical Economics* **5**(1).

Frankel, D. and Pauzner, A. (2000). Resolving indeterminacy in dynamic settings: The role of shocks, *The Quarterly Journal of Economics* **115**(1): 285–304.

Frankel, D. and Pauzner, A. (2002). Expectations and the timing of neighborhood change, *Journal of Urban Economics* **51**(2): 295–314.

Gourio, F. and Kashyap, A. K. (2007). Investment spikes: New facts and a general equilibrium exploration, *Journal of Monetary Economics* **54**: 1–22.

Guimaraes, B. (2006). Dynamics of currency crises with asset market frictions, *Journal of International Economics* **68**(1): 141–158.

Guimaraes, B. and Machado, C. (2013). Demand expectations and the timing of stimulus policies, *CEPR Discussion Paper* (9977).

Hall, G. J. (2000). Non-convex costs and capital utilization: A study of production scheduling at automobile assembly plants, *Journal of Monetary Economics* **45**(3): 681–716.

He, Z. and Xiong, W. (2012). Dynamic debt runs, *Review of Financial Studies* **25**(6): 1799–1843.

Kiyotaki, N. (1988). Multiple expectational equilibria under monopolistic competition, *The Quarterly Journal of Economics* **103**(4): 695–713.

Levin, J. (2009). The dynamics of collective reputation, *The BE Journal of Theoretical Economics* **9**(1).

Lorenzoni, G. (2009). A theory of demand shocks, *American Economic Review* **99**(5): 2050–2084.

Morris, S. and Shin, H. S. (1998). Unique equilibrium in a model of self-fulfilling currency attacks, *American Economic Review* **88**(3): 587–597.

Morris, S. and Shin, H. S. (2003). Global games: Theory and applications, *in* M. Dewatripont, L. Hansen and S. Turnovsky (eds), *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress*, Vol. 1, Cambridge University Press.

Murphy, K. M., Shleifer, A. and Vishny, R. W. (1989). Industrialization and the big push, *The Journal of Political Economy* pp. 1003–1026.

Nimark, K. (2008). Dynamic pricing and imperfect common knowledge, *Journal of Monetary Economics* **55**(2): 365–382.

Peterman, W. B. (2014). Reconciling micro and macro estimates of the frisch labor supply elasticity.

Sákovics, J. and Steiner, J. (2012). Who matters in coordination problems?, *American Economic Review* **102**(7): 3439–3461.

A Optimal policy when the planner can choose labor

In this appendix we solve the planner's problem in the model of Section 6.1 assuming it can not only choose when a firm will invest but also how much each firm will produce at each t. In other words, we also allow the planner to correct the monopoly distortion at each time t, when the proportion of firms in regime $High$ is given. We can thus break the planner's problem into two objectives: (i) in a given period, taking the aggregate state (a_t, h_t) as given, the planner's wants to achieve the optimal production for each firm; (ii) taking as given that in each period the optimal production will be implemented, the planner wants to choose investment optimally.

It is well known that in models with CES preferences a subsidy of $1/\theta$ per dollar spent on wages will achieve the optimal outcome within a period – since the markup is constant, this subsidy will induce firms to set prices equal to the marginal cost. Thus all we need to do is to solve the model assuming this wage subsidy will be paid at each t and get the planner's optimal investment decision.

The steps are essentially the same as in Section 6.1, except that the profit of a firm, given by (20) before, is now given by

$$\Pi_r = y_r^{\frac{\theta-1}{\theta}} Y^{\frac{1}{\theta}} - \left(\frac{\theta-1}{\theta}\right) w \left(\frac{y_r}{AX_r}\right)^{\frac{1}{\lambda}}$$

Since the steps to find the expressions for the endogenous variables are the same as in Section 6.1, we will omit most derivations here. We get that the gain from being in regime

High is now given by

$$\tilde{\pi}_{LS}(h,a) = \exp\left\{\left(\frac{\gamma+1}{\gamma+1-\lambda}\right)a\right\} q(h)^{\frac{\lambda\gamma(\theta-1)-(\gamma+1)}{(\theta-1)(\lambda-\gamma-1)}} \left(X_H^{(\theta-1)\zeta} - X_L^{(\theta-1)\zeta}\right)\hat{\Phi}, \tag{33}$$

where $\hat{\Phi} \equiv \lambda^{\frac{\lambda}{1+\gamma-\lambda}} - \left(\frac{\theta-1}{\theta}\right)\lambda^{\frac{\gamma+1}{1+\gamma-\lambda}}$. Notice that condition (26) that guarantees complementarity remains unchanged. We can also find new expressions for the values in equilibrium of GDP $Y(h,a)$, wage $w(h,a)$, profits in each regime $\Pi_r(h,a)$, total labor supplied $L(h,a)$ and so forth.[27] As before, the planner flow payoff is:

$$\mathcal{U}^{SP}(h,a) = Y(h,a) - \frac{1}{\gamma+1}L(h,a)^{\gamma+1}, \tag{34}$$

where $Y(h,a)$ and $L(h,a)$ are now given by

$$Y(h,a) = \lambda^{\frac{\lambda}{1+\gamma-\lambda}} \exp\left\{\left(\frac{\gamma+1}{\gamma+1-\lambda}\right)a\right\} q(h)^{\frac{\lambda\gamma\delta(\theta-1)-1}{(\gamma+1)\delta-1}}$$

and

$$L(h,a)^{\gamma+1} = \lambda^{\frac{\gamma+1}{1+\gamma-\lambda}} \exp\left\{\left(\frac{\gamma+1}{\gamma+1-\lambda}\right)a\right\} q(h)^{\frac{\lambda\gamma\delta(\theta-1)-1}{(\gamma+1)\delta-1}}.$$

In the next paragraphs we present the main findings of this appendix.

Optimal combination of subsidies. Taking derivatives of (34) with respect to h we get that

$$\frac{\partial\mathcal{U}^{SP}(h,a)}{\partial h} = \left(\frac{\theta}{\theta-1}\right)\tilde{\pi}_{LS}(h,a).$$

This equation imply that our constant subsidy result still holds in this case. The planner combines the same (constant) investment subsidy ψ/θ of Section 4 with a labor subsidy that finances a share $1/\theta$ of the firm's cost of labor. Thus, the planner's pays a proportion $1/\theta$ of all costs incurred by a firm.

Parallel shift. We can now analyze how the planner's threshold look like in comparison to the agent's. From (33) one can verify that the planner's threshold can be written as

$$\tilde{a}_P^{**}(h) = \tilde{a}_{LS}^{*}(h) - \left(\frac{\gamma+1-\lambda}{\gamma+1}\right)\log\left(\frac{\theta}{\theta-1}\right), \tag{35}$$

[27] For simplicity, unless otherwise stated, we will use the same notation as in Section 6.1 to denote the functions that determine the endogenous variables in equilibrium.

where $\tilde{a}^*_{LS}(h)$ denotes the agent's threshold if the planner's only pays the labor subsidy but not the investment subsidy. Comparing the expressions in (33) and (25) one can verify that

$$\tilde{a}^*_{LS}(h) = \tilde{a}^*(h) - \left(\frac{\gamma + 1 - \lambda}{\gamma + 1}\right) \log\left(\frac{\hat{\Phi}}{\Phi}\right) \tag{36}$$

where $\tilde{a}^*(h)$ is the equilibrium threshold when the planner does not pay any subsidy (as in Section 6.1). Doing some algebra one can verify that $\hat{\Phi} > \Phi$ and therefore $\tilde{a}^*_{LS}(h)$ is to the left of $\tilde{a}^*(h)$. Therefore, the planner's threshold here is also parallel to the agent's threshold regardless of whether agents get no subsidy at all or get just the subsidy that corrects the monopoly distortions in the labor market within a period.

Some comparative statics. From (35) and (36), one can verify that the total shift to the left implemented by the planner's when he gives both the labor and the investment subsidy is

$$\tilde{a}^{**}_P(h) - \tilde{a}^*(h) = \left(\frac{\gamma + 1 - \lambda}{\gamma + 1}\right) \log\left(\frac{\hat{\Phi}}{\Phi}\frac{\theta}{\theta - 1}\right).$$

Remember that when it cannot subsidize labor (as in Section 6.1), the threshold shift is given by

$$\tilde{a}^*_P(h) - \tilde{a}^*(h) = \left(\frac{\gamma + 1 - \lambda}{\gamma + 1}\right) \log\left(\frac{1}{1 - \chi}\right),$$

where $\tilde{a}^*_P(h)$ is defined in (28). After some algebra, we can show that $\frac{\hat{\Phi}}{\Phi}\frac{\theta}{\theta-1} > \frac{1}{1-\chi}$, and thus $\tilde{a}^{**}_P(h) < \tilde{a}^*_P(h)$ for every h. Intuitively, in some states, it might be worth investing only if firms will not distort the choice of labor, in which case the gains from a higher productivity will be totally explored.

Next figure summarizes the relative position of each threshold, where $\tilde{a}^*_P(h)$ and $\tilde{a}^*_{LS}(h)$ could switch positions on the figure, depending on parameters.

Figure 10: Relative position of the thresholds

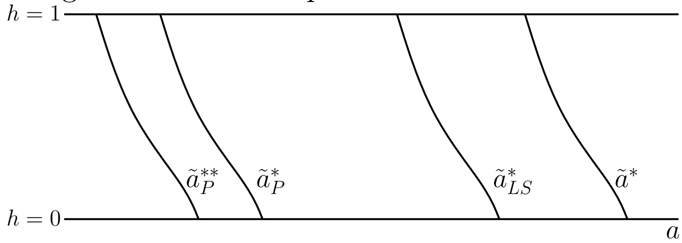

B Proofs

B.1 Proof of Proposition 1

Consider an agent deciding at time normalized to 0 who believes that every agent that will get an opportunity to change regime will choose *Low*. He assigns probability 1 that the path of h_t will be $h_t^{\downarrow} = h_0 e^{-\alpha t}$, which is independent of a. Thus, choosing *High* raises his payoff by

$$\begin{aligned}\underline{U}(h_0, a) &= \int_0^{\infty} e^{-(\rho+\alpha)t} \pi(h_t^{\downarrow}, a) dt - \psi \\ &= e^a \left(x_H^{\frac{\theta-1}{\theta}} - x_L^{\frac{\theta-1}{\theta}} \right) \int_0^{\infty} e^{-(\rho+\alpha)t} \left(h_t^{\downarrow} x_H^{\frac{\theta-1}{\theta}} + (1-h_t^{\downarrow}) x_L^{\frac{\theta-1}{\theta}} \right)^{\frac{1}{\theta-1}} dt - \psi.\end{aligned}$$

Therefore this agent will choose *High* iff $\underline{U}(h_0, a) \geq 0$. Now, $\underline{U}(h_0, a)$ is continuous and strictly increasing in a, $\lim_{a\to\infty} \underline{U}(h_0, a) = \infty$, and $\lim_{a\to-\infty} \underline{U}(h_0, a) = -\psi$. Thus for any h_0, there is $a = a_H(h_0)$ such that $\underline{U}(h_0, a) = 0$. Since $\underline{U}(h_0, a)$ is strictly increasing in a, for any $a' > a_H(h_0)$ we have $\underline{U}(h_0, a') > 0$ and thus choosing *High* is a strictly dominant strategy (any other belief about the path of h_t will raise the relative payoff of choosing *High*). Notice that $\underline{U}(h_0, a)$ is strictly increasing in both a and h_0 and thus $a_H(h_0)$ is strictly decreasing.

A similar argument proves that there exists a strictly decreasing threshold a^L such that if $a < a^L(h_0)$, *Low* is a dominant action. Consider an agent who believes others will choose *High* after him. He believes that the motion of h_t will be given by $h_t^{\uparrow} = 1 - (1-h_0)e^{-\alpha t}$, so choosing *High* instead of *Low* raises his payoff by

$$\begin{aligned}\overline{U}(h_0, a) &= \int_0^{\infty} e^{-(\rho+\alpha)t} \pi(h_t^{\uparrow}, a) dt - \psi \\ &= e^a \left(x_H^{\frac{\theta-1}{\theta}} - x_L^{\frac{\theta-1}{\theta}} \right) \int_0^{\infty} e^{-(\rho+\alpha)t} \left(h_t^{\uparrow} x_H^{\frac{\theta-1}{\theta}} + (1-h_t^{\uparrow}) x_L^{\frac{\theta-1}{\theta}} \right)^{\frac{1}{\theta-1}} dt - \psi.\end{aligned}$$

This agent will choose *Low* whenever $\overline{U}(h_0, a) < 0$ and, as in the previous case, we can show that there exists a strictly decreasing threshold a^L such that if $a < a^L(h_0)$, *Low* is a dominant action. Since for every h_0 and $t > 0$ we have $h_t^{\uparrow} > h_0 > h_t^{\downarrow}$, $\overline{U}(h_0, a) > \underline{U}(h_0, a)$. This implies $a^H(h_0) > a^L(h_0)$.

Take a pair (a, h_0) such that $a_L(h_0) < a < a_H(h_0)$. Since $a < a_H(h_0)$, if an agent believes that the path of h_t will be $h_t^{\downarrow}$, then $\underline{U}(h_0, a) < 0$ and thus his optimal strategy is to play *Low*. Therefore this belief is consistent and the strategy profile where every player plays *Low* is an Nash equilibrium. Likewise, since $a > a_L(h_0)$ the strategy profile where every player plays *High* is also a Nash equilibrium. Hence, there is multiplicity in this set. □

B.2 Proof of Proposition 4

For a given $h \in [0,1]$ the maximum amount of subsidies the planner have to pay to implement his threshold $a_p^*(h)$ is $\varphi(h, a_p^*(h))$, since agents payoffs are increasing in a. But from Proposition 3 we know that an agent will be indifferent between investing and not investing at $(h, a_p^*(h))$ if the cost is $\psi - \psi/\theta$ and he believes that the other's will play according to $a_p^*(h)$. Thus, $\varphi(h, a_p^*(h)) = \psi/\theta$, which is independent of h. □

B.3 Proof of Proposition 5

In order to apply the existence arguments in Frankel and Pauzner (2000), it suffices to show that playing *High* is a dominant choice for some large enough a and that *Low* is a dominant choice for some small enough a. This is so because i.i.d. shocks are needed just to show uniqueness, and Corollary 1 in Burdzy et al. (1998) guarantees that Lemma 1 in Frankel and Pauzner (2000), used in their proof, holds for our more general process for a_t.

Solving $da_t = \eta(\mu - a_t)dt + \sigma dZ_t$ we get that

$$a_t = a_0 e^{-\eta t} + \mu(1 - e^{-\eta t}) + \sigma \int_0^t e^{\eta(s-t)} dZ_s. \tag{37}$$

And thus a_t conditional on a_0 is normally distributed with mean

$$E_0\left[a_t\right] = \mu + e^{-\eta t}(a_0 - \mu). \tag{38}$$

and variance

$$Var_0\left[a_t\right] = \frac{\sigma^2}{2\eta}\left(1 - e^{-2\eta t}\right). \tag{39}$$

Therefore, e^{a_t} conditional on a_0 follows a log-normal distribution with mean

$$E_0\left[e^{a_t}\right] = \exp\left\{\mu + e^{-\eta t}(a_0 - \mu) + \frac{1}{4}\frac{\sigma^2}{\eta}\left(1 - e^{-2\eta t}\right)\right\}. \tag{40}$$

Consider an agent deciding at some point $(0, a_0)$ who believes that $h_t = 0$ for every $t \geq 0$. His utility gain from choosing *High* is

$$\begin{aligned} \underline{W}(a_0) &= \left(x_H^{\frac{\theta-1}{\theta}} - x_L^{\frac{\theta-1}{\theta}}\right) x_L^{\frac{1}{\theta}} \int_0^\infty e^{-(\rho+\alpha)t} E_0\left[e^{a_t}\right] dt - \psi \\ &> \left(x_H^{\frac{\theta-1}{\theta}} - x_L^{\frac{\theta-1}{\theta}}\right) x_L^{\frac{1}{\theta}} \int_0^1 e^{-(\rho+\alpha)t} \inf\left\{E_0\left[e^{a_t}\right]\right\}_{t\in(0,1)} dt - \psi. \end{aligned}$$

By (40), we have that $\lim_{a_0\to\infty} \inf\left\{E_0\left[e^{a_t}\right]\right\}_{t\in(0,1)} = \infty$. Thus, there exists some large enough a^{**} such that *High* is a strictly dominant action when $a > a^{**}$.

Now consider an agent deciding at some point $(1, a_0)$, with $a_0 < \mu$, who believes that $h_t = 1$, for every $t \geq 0$. His gain in utility of choosing *High* is given by

$$\begin{aligned}\overline{W}(a_0) &= \left(x_H^{\frac{\theta-1}{\theta}} - x_L^{\frac{\theta-1}{\theta}}\right) x_H^{\frac{1}{\theta}} \int_0^\infty e^{-(\rho+\alpha)t} E_0\left[e^{a_t}\right] dt - \psi \\ &< \left(x_H^{\frac{\theta-1}{\theta}} - x_L^{\frac{\theta-1}{\theta}}\right) x_H^{\frac{1}{\theta}} \left(\int_0^Q e^{-(\rho+\alpha)t} \left(\sup\left\{E_0\left[e^{a_t}\right]\right\}_{t\in(0,Q)}\right) dt \right. \\ &\qquad \left. + \int_Q^\infty e^{-(\rho+\alpha)t} \left(\mu + \frac{\sigma^2}{4\eta}\right) dt\right) - \psi.\end{aligned}$$

By (40), we have that $\lim_{a_0\to-\infty} \sup\left\{E_0\left[e^{a_t}\right]\right\}_{t\in(0,Q)} = 0$. For large enough Q, the integral term is small enough, so $\overline{W}(a_0) < 0$. Hence there exists some small enough a^{**} such that *Low* is a strictly dominant action when $a < a^{**}$.

We have shown the existence of dominant regions. Now, as in Frankel and Pauzner (2000) we can iteratively eliminate strictly dominated strategies. This process converges to a threshold a^* such that agents are indifferent between investing or not at $(a^*(h), h)$ for all $h \in [0, 1]$, if they believe the others will play according to a^*. Given a threshold a^*, notice that payoffs are increasing in a and h. Thus, playing according to a^* is an equilibrium.

□

B.4 Proof of Proposition 6

In order to prove Proposition 6, we need to show the existence of dominance regions, so that we can apply the results on Frankel and Burdzy (2005). First, Lemma 1 constructs a process u_t that allows us to replace payoffs $\hat{\pi}(\cdot)$ by the original payoffs $\pi(\cdot)$. Then, Lemma 2 proves the existence of dominance regions.

Lemma 1. Let u_t *be the following stochastic process*

$$u_t = \begin{cases} a_t & \text{if } a_t < M \\ M & \text{otherwise} \end{cases},$$

where a_t *is given by* $da_t = \eta_t(\mu - a_t)dt + \sigma dZ_t$*, with* η_t *given by (31). Then* $\lim_{a_\tau\to\infty} E_\tau\left[e^{u_t}\right] = e^M$ *and* $\lim_{a_\tau\to-\infty} E_\tau\left[e^{u_t}\right] = 0$*, for every* $t > \tau$ *and every* $\tau \geq 0$.

Proof. First assume $\tau < T$. It follows from (37) that when $\tau < t < T$, $a_t|a_\tau$ has a normal distribution with mean and variance given by $\mu + e^{-\eta(t-\tau)}(a_\tau - \mu)$ and $\frac{\sigma^2}{2\eta}\left(1 - e^{-2\eta(t-\tau)}\right)$, respectively. In that case we have

$$E_\tau\left[e^{u_t}\right] = \frac{1}{\Sigma_t\sqrt{2\pi}}\int_{-\infty}^{M}\exp\left\{a_t - \frac{1}{2}\left(\frac{a_t-\mu-e^{-\eta(t-\tau)}(a_\tau-\mu)}{\Sigma_t}\right)^2\right\}da_t$$
$$+\left(1-\Phi\left(\frac{M-\mu-e^{-\eta(t-\tau)}(a_\tau-\mu)}{\Sigma_t}\right)\right)e^M, \quad (41)$$

where Φ is the standard normal distribution and $\Sigma_t \equiv \sigma\sqrt{\frac{1}{2\eta}\left(1-e^{-2\eta(t-\tau)}\right)}$.

Now fix $t \geq T$. In that case, we have that $a_t|a_T$ follows a normal distribution with mean a_T and variance $\sigma^2(t-T)$. Therefore, by the law of iterated expectations,

$$E_\tau\left[a_t\right] = E_\tau\left[E_\tau\left[a_t|a_T\right]\right] = E_\tau\left[a_T\right] = \mu + e^{-\eta(T-\tau)}(a_\tau-\mu),$$

where the last equality follows from (38). Moreover, by the law of total variance,

$$Var_\tau\left[a_t\right] = E_\tau\left[Var_\tau\left[a_t|a_T\right]\right] + Var_\tau\left[E_\tau\left[a_t|a_T\right]\right] = \sigma^2(t-T)+\Sigma_T^2.$$

We can show that $a_t|a_\tau$ follows a normal distribution[28] and so,

$$E_\tau\left[e^{u_t}\right] = \frac{1}{\sqrt{\left(\sigma^2(t-T)+\Sigma_T^2\right)2\pi}}\int_{-\infty}^{M}\exp\left\{a_t - \frac{1}{2}\left(\frac{a_t-\mu-e^{-\eta(T-\tau)}(a_\tau-\mu)}{\sqrt{\left(\sigma^2(t-T)+\Sigma_T^2\right)}}\right)^2\right\}da_t$$
$$+\left(1-\Phi\left(\frac{M-\mu-e^{-\eta(T-\tau)}(a_\tau-\mu)}{\sqrt{\left(\sigma^2(t-T)+\Sigma_T^2\right)}}\right)\right)e^M. \quad (42)$$

Notice that both (41) and (42) are continuous on a_τ and that they coincide at $t=T$. Taking limits with $a_\tau \to \infty$ and $a_\tau \to -\infty$ of (41) and (42) completes the proof for the case where $\tau < T$. The proof for the case where $\tau \geq T$ is very similar and therefore omitted. □

Lemma 2. *Suppose $\sigma > 0$, the mean reversion parameter η_t is given by (31) and the relative payoff of investing is given by (32). Then, if M is sufficiently high, there are constants a' and a'', with $a' < a''$ such that if $a(h) > a''$ it is strictly dominant to play High and if $a(h) < a'$ it is strictly dominant to play Low.*

Proof. First, notice that we can write $\hat{\pi}(h_t, a_t) = \pi(h_t, u_t)$, for every t. Assume that an agent deciding at some period normalized to 0 has the belief that $h_t = 0$, for every $t \geq 0$. Thus, his payoff of investing is given by

[28] We know that $a_T|a_\tau \sim N(\mu + e^{-\eta(T-\tau)}(a_\tau-\mu), \frac{\sigma^2}{2\eta}(1-e^{-2\eta(T-\tau)}))$ and $a_t|a_T, a_\tau \sim N(a_T, (t-T)\sigma^2)$. Since $E_\tau\left[a_t|a_T\right]$ is linear on a_T and $Var_\tau\left[a_t|a_T\right]$ does not depend on a_T we guarantee bivariate normality of the vector (a_t, a_T) conditional on a_τ (see Arnold et al. (1999), p. 56) and therefore its marginal distributions are normal.

$$
\begin{aligned}
\underline{\mathcal{U}}(a_0) &\equiv x_L^{\frac{1}{\theta}} \left(x_H^{\frac{\theta-1}{\theta}} - x_L^{\frac{\theta-1}{\theta}} \right) \int_0^\infty e^{-(\rho+\alpha)t} E_0 \left[e^{u_t} \right] dt - \psi \\
&> x_L^{\frac{1}{\theta}} \left(x_H^{\frac{\theta-1}{\theta}} - x_L^{\frac{\theta-1}{\theta}} \right) \int_0^1 e^{-(\rho+\alpha)t} E_0 \left[e^{u_t} \right] dt - \psi \\
&> x_L^{\frac{1}{\theta}} \left(x_H^{\frac{\theta-1}{\theta}} - x_L^{\frac{\theta-1}{\theta}} \right) e^{-(\rho+\alpha)} \inf \left\{ E_0 \left[e^{u_t} \right] \right\}_{t\in(0,1)} - \psi.
\end{aligned}
$$

But $\inf \left\{ E_0 \left[e^{u_t} \right] \right\}_{t\in(0,1)}$ converges to e^M when a_0 goes to ∞. Thus as long as M is suficiently high, we can get an a'' such that $\underline{\mathcal{U}}(a'') > 0$.

Now consider an agent deciding at some period τ normalized to zero that believes that $h_t = 1$ for every $t \geq 0$. His gain in utility of investing is given by

$$
\begin{aligned}
\overline{\mathcal{U}}(a_0) &\equiv x_H^{\frac{1}{\theta}} \left(x_H^{\frac{\theta-1}{\theta}} - x_L^{\frac{\theta-1}{\theta}} \right) \int_0^\infty e^{-(\rho+\alpha)t} E_0 \left[e^{u_t} \right] dt - \psi \\
&< x_H^{\frac{1}{\theta}} \left(x_H^{\frac{\theta-1}{\theta}} - x_L^{\frac{\theta-1}{\theta}} \right) \left(\int_0^Q e^{-(\rho+\alpha)t} \sup \left\{ E_0 \left[e^{u_t} \right] \right\}_{t\in(0,Q)} dt + \int_Q^\infty e^{-(\rho+\alpha)t} e^M dt \right) - \psi.
\end{aligned}
$$

But $\sup \left\{ E_0 \left[e^{u_t} \right] \right\}_{t\in(0,Q)}$ goes to zero as a_0 goes to $-\infty$. For large enough Q, the second integral term is small enough, so for sufficiently small a', we get $\overline{\mathcal{U}}(a') < 0$. □

Proof of Proposition 6. We have proved the existence of dominance regions, hence our model is a special case of Frankel and Burdzy (2005) and uniqueness follows from Theorems 1 and 4 in that paper.

□