REAP Rede de Economia Aplicada

Demand expectations and the timing of stimulus policies

Bernardo Guimaraes
Caio Machado

Julho, 2013
Working Paper 58

DEMAND EXPECTATIONS AND
THE TIMING OF STIMULUS POLICIES

Bernardo Guimaraes
Caio Machado

Bernardo Guimarães
Escola de Economia de São Paulo
Fundação Getúlio Vargas (EESP/FGV)
Rua Itapeva, nº 474
01332-000 - São Paulo, SP – Brasil
bernardo.guimaraes@fgv.br

Caio Machado
Escola de Economia de São Paulo
Fundação Getúlio Vargas (EESP/FGV)
Rua Itapeva, nº 474
01332-000 - São Paulo, SP – Brasil
caiohm@gmail.com

Demand expectations and the timing of stimulus policies[*]

Bernardo Guimaraes[†]   Caio Machado[‡]

July 2013

Abstract

This paper proposes a simple macroeconomic model with staggered investment decisions. The expected return from investing depends on demand expectations, which are pinned down by fundamentals and history. Owing to an aggregate demand externality, investment subsidies can improve welfare in this economy. The model can be used to address questions concerning the timing of stimulus policies: should the government spend more on preventing the economy from falling into a recession or on rescuing the economy when productivity picks up? Results show the government should strike a balance between both objectives.

JEL CLASSIFICATION: D84, E32, E62

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

1 Introduction

The recent recession is often explained by a combination of fundamental shocks and pessimistic expectations. The narrative goes like this: a negative shock from the financial sector spread to the rest of the economy and led firms to reduce investment. Low levels of economic activity have since persisted because, owing to low demand expectations, firms have been reluctant to resume previous levels of investment. In turn, this reduced investment has contributed to low demand, justifying pessimistic expectations. Hence a dynamic coordination problem lies at the heart of the recession.

Following the large investment slump of 2008-2009, stimulus packages around the world have been proposed and implemented.[1] These fiscal packages can be seen as attempts to mitigate dynamic coordination failures. By providing incentives for investment, governments hope to boost demand expectations and drive the economy to a situation with higher expected and realized economic activity.

This paper develops a simple macroeconomic model that captures that dynamic coordination problem among producers. The model features monopolistic competition and staggered investment decisions. An investment is a payment of a fixed cost that increases production capacity. Producers of each variety receive investment opportunities according to an exogenous Poisson process. That is a simple way to capture the idea that capital can not adjust overnight, leaving an important role in the model for expected demand. When deciding whether to invest or not, a producer has to form expectations on others' future decisions. If producers with subsequent investment opportunities choose to take them, the demand for a given variety shifts to the right, leading to a larger price and larger profits. Hence investment decisions are strategic complements, as in Kiyotaki (1988).

Investment decisions depend not only on expected 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, but once we allow for shocks, that is not true anymore, as in Frankel & Pauzner (2000).

Demand expectations are pinned down by fundamentals and history. The equilibrium of the model is characterized by a cutoff strategy given by a threshold that depends on the exogenous productivity parameter (fundamental) and the mass of producers that are currently operating at full capacity, which results from their recent choices (history). For a given level of fundamentals, producers choose to invest if the mass of producers operating at full capacity is sufficiently high, since that positively affects both demand today and

---

[1]The objective of stimulating the economy has been translated into concrete policies in a number of different ways, such as: cuts in energy prices, tax cuts, subsidized loans and fiscal incentives to investment (either to the whole economy or to specific industries). Khatiwada (2009) provides a comprehensive review of those policies.

the actions of others tomorrow. Recessions are triggered by shocks on fundamentals, but expectations about others' actions play a key role.

Investment generates positive externalities because a producer does not take into account the effects of her investment decisions on others' profits. Thus investment subsidies can mitigate coordination failures. The model provides a tractable framework that can be used to analyze stimulus policies that aim at affecting "market confidence". Demand expectations are important to determine investment decisions. Beliefs about others' actions are pinned down in the model, so we can understand how stimulus measures affect producers' demand expectations and investment.

The dynamic framework is particularly suitable to answer questions regarding the timing of govenment interventions. In particular, consider a policy maker with a certain budget to be spent in fiscal stimulus. Is preventing a recession better than rescuing the economy after an investment slump has already occurred? Or should policy makers give up early on avoiding a recession and provide incentives for producers once fundamentals have improved? How should incentives vary with economic activity and fundamentals?

The question can be posed in the following way: the equilibrium threshold can be represented as a curve in a two-dimensional space, with productivity in the horizontal axis and the measure of agents operating at full capacity in the vertical axis. Agents choose to invest if the economy is at the right of the threshold. The threshold is negatively sloped, implying that when the mass of producers operating at full capacity is larger, a producer requires a smaller level of productivity to invest. The government intervention aims at shifting the equilibrium threshold so that more producers will choose to invest. Besides translating the threshold to the left, how should it try to rotate the threshold? Should it try to estimulate investment primarily when most producers are still operating at full capacity, despite relatively weak fundamentals, in order to avoid an investment slump? Or should the policy maker focus on subsidizing investment when few producers are operating at full capacity, but productivity is picking up?

Neither of those is the answer. The government should shift the threshold to the left, increasing the region where investment occurs, but not rotate it. Trying to avoid a recession when productivity is very low, or trying to rescue the economy when the mass of agents is very low are both very expensive. Notwithstanding the importance of the demand externality, the equilibrium threshold features a balance between changes in fundamentals and in economic activity that should not be significantly affected by the government intervention.

The threshold arising from this policy (roughly) coincides with the results of a policy that prescribes a constant subsidy in every state of the world.[2] The optimal policy estab-

---

[2]The conclusions are based on numerical results, so we cannot make statements about the exact differences (if any) between a threshold implemented by a stimulus policy prescribing constant subsidies and the policy that shifts the threshold to the left (without rotation).

lishes a maximum level of subsidies to investment that is independent on productivity, capacity utilization and economic activity. Intuitively, at the margin, the objective is to get the largest possible mass of producers to invest with a given amount of resources, so preventing the economy from falling into a recession or rescuing the economy once productivity picks up are equally important.

The demand externalities that play a key role in this paper are in the seminal contributions by Blanchard & Kiyotaki (1987) and Kiyotaki (1988). 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.[3] Despite generating interesting insights, this approach does not allow us to understand how policies directly affect expectations.[4]

This paper is closely related to the theoretical contributions in Frankel & Pauzner (2000) and Frankel & 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.[5] The uniqueness result in Frankel & Pauzner (2000) requires very small mean reversion, but Frankel & Burdzy (2005) generalize some of the results in Frankel & Pauzner (2000) for more general stochastic processes.[6] We use some of their theoretical results to show that a threshold equilibrium exists, and that there is a unique rationalizable equilibrium for a sequence of models that converges to our model. 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.[7]

There has not been much work applying those theoretical insights to understand the effects of stimulus packages on coordination. One important exception is Sákovics & Steiner (2012). They build a model to understand who matters in coordination problems:

---

[3]See, e.g., Cooper & John (1988).

[4]Government policies can surely affect the region where multiple equilibrium exist, but when the economy is inside that region, models with multiple equilibria cannot predict which one will be played.

[5]Models with time-varying fundamentals and timing frictions have been used to study other dynamic coordination problems. Frankel & 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 & Xiong (2012) study debt runs.

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

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

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.

Coordination and strategic complementarities also play a key role in the models of Angeletos & La'O (2010) and Angeletos & La'O (2013). They show in an environment with noisy and dispersed information how self-fulfilling fluctuations can emerge. The reasons for strategic complementarities in those models and in this one are similar. However, 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 study the dynamic interplay between economic activity and productivity. Expectations also play a key role in the literature of news-driven business cycles (e.g., Beaudry & Portier (2006)), but here expectations about future productivity depend solely on the current state of the economy. In the models of Lorenzoni (2009) and Eusepi & Preston (2011), it is noisy information about current variables that leads to excessive optimism or pessimism about the future.

The paper is organized as follows. Section 2 presents the model. Section 3 presents the policies analyzed in the paper, Section 4 describes and analyzes the results and section 5 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

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

where P_t is the price of the composite good and 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. Agent $i \in [0,1]$ produces intermediate good i. Her instantaneous utility at time t is given by $U_t = C_t$, where C_t is her instantaneous consumption of the composite good. Since y_{it} is

the quantity produced by agent i at time t, her budget constraint is given by

$$P_t C_t \leq p_{it} y_{it} \equiv w_i.$$

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

Investment is a sunk cost that reduces marginal cost of production. The assumptions on technology aim at modelling staggered investment decisions in a simple and tractable way. There are 2 states, *High* and *Low*. An agent in state *Low* can produce up y_{Lt} units at zero marginal cost at every time t, and an agent in state *High* 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. Agents get a chance to switch states according to a Poisson process with arrival rate α. Once an individual is picked up, he chooses a state and will be locked in this state until he is selected again. Choosing state *Low* is costless. Choosing state *High* implies a one-off cost ψ in units of the composite good (ψ is a stock).

Choosing state *High* can be interpreted as an investment decision. The cost ψ can be thought of as the cost of a machine and the difference $y_{Ht} - y_{Lt}$ as the resulting gain in productivity. This machine will become obsolete after some (random) time (so α also plays the role of a depreciation rate). Moreover, agents are locked in a state until the next investment opportunity arises, which captures the idea that firms cannot change their capital level overnight.[8] 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.

Investment requires agents to acquire a stock 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, there are other types of assets that would deliver the same results. The assumption of linear utility implies that consumption smoothing plays no role in the model, all results come from investment decisions.

---

[8]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 model.

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

$$da_t = \eta(\mu - a_t)dt + \sigma dZ_t, \tag{3}$$

where $\eta \geq 0$, $\sigma > 0$ and Z_t is a standard Brownian motion. The parameter η determines how fast a_t returns to its mean, given by μ.

2.2 Equilibrium

The composite-good firm chooses its demand for each intermediate good taking prices are given. Using (1) and (2), we get

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

for $i \in [0,1]$, and the price of the composite good is given by:

$$P_t \equiv \left(\int_0^1 p_{it}^{1-\theta} di \right)^{1/(1-\theta)}.$$

Since marginal cost is zero and marginal revenue is always positive, an agent in state *Low* will produce y_{Lt}, and an agent in state *High* 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 state is given by

$$w_{Ht} = p_{Ht} y_{Ht} = y_{Ht}^{\frac{\theta-1}{\theta}} Y_t^{\frac{1}{\theta}} P_t \tag{4}$$

and

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

Moreover, using (1),

$$Y_t = \left(h_t y_{Ht}^{\frac{\theta-1}{\theta}} + (1-h_t) y_{Lt}^{\frac{\theta-1}{\theta}} \right)^{\frac{\theta}{\theta-1}}, \tag{6}$$

where h_t is the measure of agents locked in state H.

The indirect utility over consumption goods for an agent with income equal to w is given by w/P_t. Combining (4), (5) and (6), we get the instantaneous utility of individuals locked in each state:

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

and

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

Let $\pi(h_t, a_t)$ be the difference between instantaneous utility of agents locked in state *High*

and agents locked in state *Low* when the economy is at (h_t, a_t). Then

$$\pi(h_t, a_t) = e^{a_t} \left(h_t x_H^{\frac{\theta-1}{\theta}} + (1 - h_t)x_L^{\frac{\theta-1}{\theta}} \right)^{\frac{1}{\theta-1}} \left(x_H^{\frac{\theta-1}{\theta}} - x_L^{\frac{\theta-1}{\theta}} \right). \quad (7)$$

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 (7) 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

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

In words, investing pays off if the discounted expected additional profits of choosing state *High* 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 $(\rho + \alpha)$.[9]

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.

2.3 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*

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

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

---

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

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 A. □

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. If the fundamental a is in an intermediate area, there are no dominant strategies: the optimal investment decision depends on expectations about what others will do. 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.

Figure 1: Equilibria without shocks
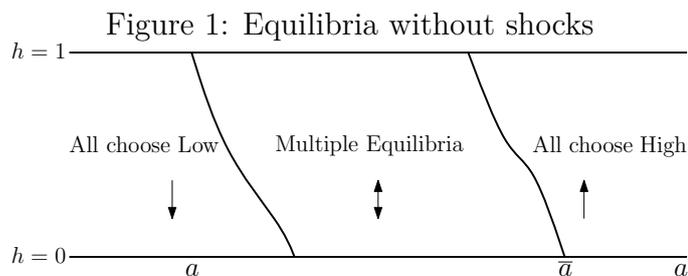
2.4 General case

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.

Proposition 2 (Existence). *Suppose $\sigma > 0$. There exists a strictly decreasing function a^* such that a^* is an equilibrium.*

Proof. See Appendix A. □

Proposition 2 builds on Frankel & Pauzner (2000) to show that a threshold equilibrium always exists. The threshold function a^* is decreasing in h, so a larger h implies that agents are willing to invest for lower values of a. In a threshold equilibrium, beliefs about others' investment decisions are pinned down by fundamentals (a) and history (h). Demand expectations fluctuate because shocks to a_t and movements in h_t might trigger changes on expectations about others' actions. Figure 2 shows an example of equilibrium in the model.

Figure 2: Equilibrium with shocks
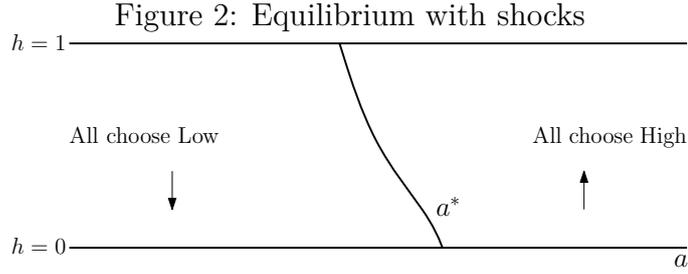
Let $V(a, h, \tilde{a})$ be the utility gain from choosing *High* obtained by an agent in state (a, h) that believes others will play according to threshold $\tilde{a}$. Then

$$V(a, h, \tilde{a}) = \int_0^\infty e^{-(\rho + \alpha)t} E[\pi(h_t, a_t)|a, h, \tilde{a}]dt - \psi, \tag{8}$$

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 $\tilde{a}$. An agent choosing when $a = a^*(h)$ and 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.

2.4.1 On equilibrium uniqueness

We do not have a strong uniqueness result. However, we can show that our model can be seen as a limiting case of a sequence of models that have a unique rationalizable equilibrium.

In order to apply the results of Frankel & Burdzy (2005), we need to make two changes in the model. First, the diffusion process for a_t is given by equation 3, but the mean-reversion parameter η_t varies over time so that

$$\eta_t = \begin{cases} \eta & \text{if } t < T \\ 0 & \text{otherwise} \end{cases}, \tag{9}$$

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

$$\hat{\pi}(h, a) = \begin{cases} \pi(h, a) & \text{if } a < M \\ \pi(h, M) & \text{otherwise} \end{cases}, \tag{10}$$

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 & Burdzy (2005), we can prove there is a unique equilibrium in this model.

Proposition 3 (Uniqueness, Frankel & Burdzy (2005)). *Suppose $\sigma > 0$, the mean rever-sion parameter η_t is given by (9) and the relative payoff of investing is given by (10). Then*

10

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 A. □

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 (8) 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 & Pauzner (2000) could be applied to show equilibrium uniqueness. However, in the presence of mean reversion, the last step in the proof of Frankel & 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 & 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 state 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 & 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.

3 Stimulus policies

The price of a particular variety depends positively on the quantity produced of other goods. For instance, if others are selling y_{Lt} units of their goods, a producer will face low demand and will only be able to sell y_{Ht} units at a low price. Consequently, one's profits are increasing on others' output. Since investing has a positive externality on other agents, we expect that without any intervention there will be underinvestment in this economy.

We now entertain the possibility of investment subsidies in this economy. A stimulus policy consists in a potentially state-dependent investment subsidy for agents that choose to invest. Formally, a policy specifies a subsidy $\varphi(h, a)$ that will be given in state (h, a) for those who pay the fixed cost ψ. We assume the policy is perfectly anticipated by all agents.

Suppose the government gives a constant subsidy, that is, $\varphi(h, a) = d > 0$ for all (h, a). It can be shown that the equilibrium threshold shifts to the left (from a^* to a_p^* in Figure 3). For any given h, agents require a smaller a to invest. By effectively reducing the fixed cost ψ, the policy maker can affect the equilibrium and lead to more investment, weakening coordination failures.

Figure 3: Constant subsidy
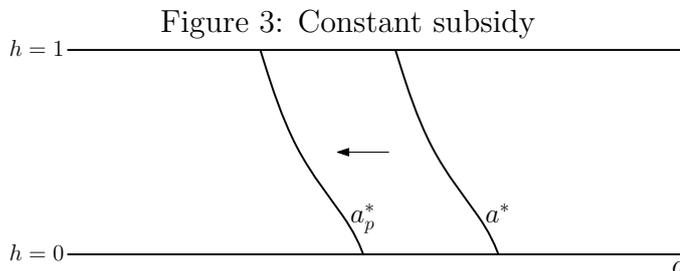
However, a constant subsidy would be quite expensive. An agent choosing when a is high enough would not require extra incentives to invest. Thus there would be cheaper ways to provide the same amount of incentives. In practice, subsidies are costly because they have to be financed with distortionary taxation. Hence it is important to know which stimulus policies can achieve a given result with a minimum amount of government spending.

We now focus on *minimal spending policies*, which are the cheapest way to subsidize producers that implements a certain threshold.[10]

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

---

[10]There would be cheaper ways to implement a threshold if policies were allowed to determine payments from producers that strictly prefer to invest or that are not investing. However, that would not be a stimulus policy. The objetive of this paper is to understand which policies minimize spending. It is also important to understand which policies minimize dead-weight losses from taxation, but that is beyond the scope of this paper.

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

$$\varphi(h, a) = \begin{cases} \psi - \int_0^\infty e^{-(\rho+\alpha)t} E[\pi(h_t, a_t)|a, h, a_p^*]dt & \text{if } a_p^*(h) \leq a \leq \hat{a}(h) \\ 0 & \text{otherwise} \end{cases}. \quad (11)$$

Figure 4 shows three thresholds: a_p^* is the threshold the policy wishes to implement, a^* is the threshold without intervention and $\hat{a}$ is the best response of a player that believes others will play according to a_p^*. By definition, a^* is the best response to others playing according to a^*. 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 4). The 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^*

Figure 4: Example of minimal spending policy
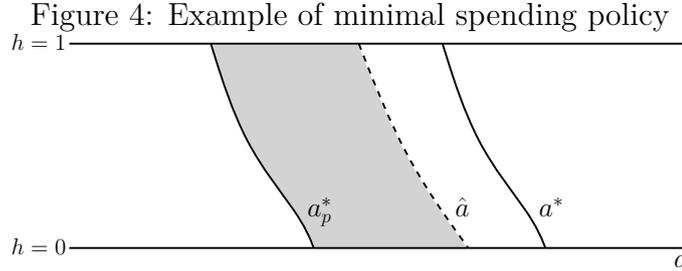
Notice that the equilibrium under the minimal spending policy is no longer unique. If agents believe others will play according to a^* their best response is to play according to a^*, and thus the policy has no effect at all.

Our aim is to compare the required government spending from policies that yield similar results in terms of utility gains. Consider the two thresholds a_p^* and $a_{p'}^*$ depicted in Figure 5 and assume both deliver the same lifetime utility for an agent born in a random state. The objective is to know which one is less costly.

The thresholds a_p^* and $a_{p'}^*$ correspond to different policy objectives. A stimulus policy that implements the threshold a_p^* is not particularly concerned with either preventing the economy from falling into a recession or rescuing the economy when productivity picks up. In contrast, a policy that implements $a_{p'}^*$ prescribes subsidies to investment when a is relatively low while h is still high, which might keep the economy away from the region where h falls down and avoid an investment slump.

Figure 5: Two different minimal spending policies
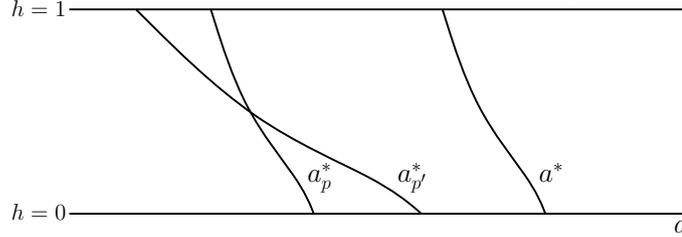
4 Numerical results

In order to solve the model numerically, we work with an approximation of the model presented in section 3. Now time is discrete and each period has length Δ, where Δ is a small number. Hence time $t \in \{0, \Delta, 2\Delta, 3\Delta, ...\}$. The stochastic process of a_t is given by

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

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 state. The instantaneous payoffs of being locked in each state 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 3.

4.1 Threshold Computation

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, pick an arbitrary threshold a_0^* and choose a finite grid for h in the interval $[0, 1]$. Then, for every point 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.[11]

---

[11]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. Both equilibrium thresholds and the one found using the first algorithm presented coincide, but these are more expensive in terms of computing time.

4.2 Policy Simulations

Once the equilibrium threshold a^* has been obtained, we consider some minimal spending policies parameterized in the following way:

$$a^*_{c,\xi}(h) = a^*(h) + \xi\vartheta(h) - c,$$

where $\vartheta(h) = 1 - 2h$, $c > 0$ and ξ is a real number. A policy is then a pair (ξ, c), where $c > 0$ implies the threshold is shifed to the left and $\xi \neq 0$ implies the threshold is rotated.

Initially, we set a grid for ξ that includes zero and a value for c. We then estimate the lifetime utility of a representative agent born in a random state (we simulate the economy and take out the first 50 years), given the stimulus policy $(0, c)$. We then do the same for different values of ξ and adjust the value of c for each policy so that the utility gain from all policies is (approximately) the same. At the end of this process, we have different stimulus policies that deliver the same utility but different slopes.

For each policy, we find the best response of an agent (the curve $\hat{a}$ in figure 4) given that others will follow the threshold prescribed by the policy. Then we compute the gain in utility from picking *High* for a set of points (in the gray area in Figure 4). Using interpolations, we can find the subsidy needed at each point $\varphi(h, a)$ to make the agent indifferent between investing or not. Finally, we simulate the economy several times and estimate the government spending under each policy by applying the formula given by (11).

4.3 Calibration

In the baseline calibration, 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.[12]

- The economy stays 30% of time at the left of the threshold, that is, agents are not investing 30% of the time, approximately.[13] Under the stimulus policy with $\xi = 0$, this number falls to 12%.

- 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.[14]

---

[12]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.

[13]When the economy is to the left of the treshold, 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.

[14]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 state *High* is equal to $(\rho + \alpha)\psi$. At time t, there are h_t agents in state *High*, so we subtract the cost of capital in the economy $h_t(\rho + \alpha)\psi$ from the total amount produced, given by (6).

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). In Appendix B we show that our results are robust to different specifications.

Table 1: Parameters

Parameter	Symbol	Value
Production state *High*	x_H	1.1
Production state *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

4.4 Results

Figure 6 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 7 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

Figure 6: Estimated threshold
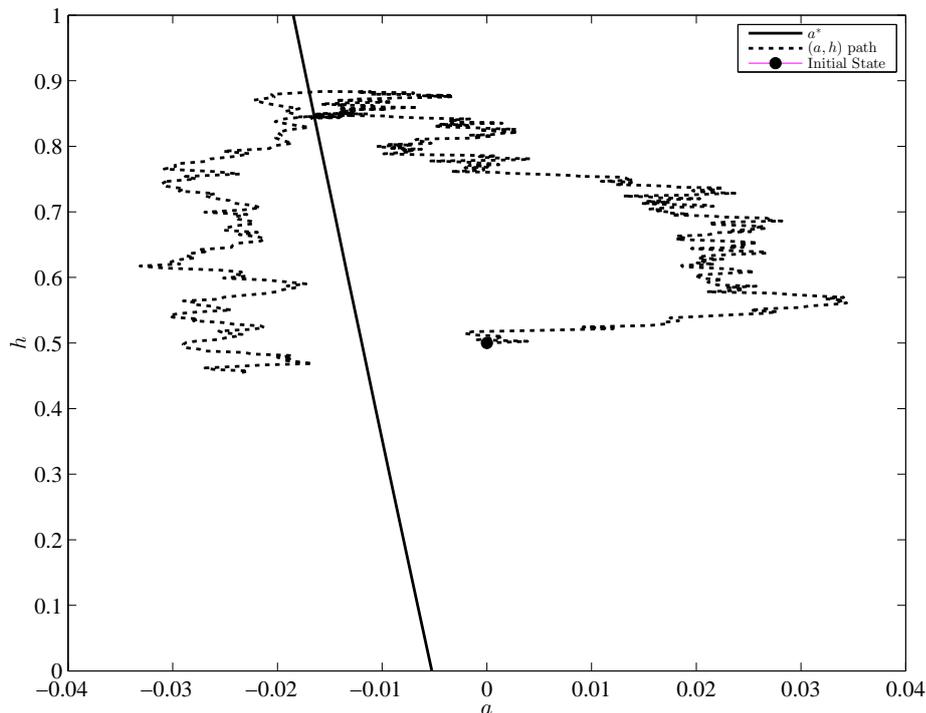
nothing about the exogenous movements in a but can increase the region where agents invest. Investment subsidies can bring output closer to the $h = 1$ curve.

Besides amplifying the effects of negative shocks, the endogenous and staggered reaction of h also implies that low productivity periods have long-lasting negative effects. As shown in Figure 7, 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.

We now turn to the comparison of different policies. Figure 8 shows different stimulus policies corresponding to different values of ξ, that deliver the same welfare improvement. The thin dashed lines are the *loci* where an agent is indifferent between investing and not if she believes others will act according to the threshold implemented by the stimulus policy. The government pays subsidies in the area between the policy threshold and its respective thin dashed line.

Figure 8 helps understanding the different effects of each policy. Consider an economy in a recession, with $h = 0$ and $a = -0.045$, and a_t is moving up towards zero. The stimulus policy that implements the threshold with higher slope (the almost vertical line in Figure 8) kicks in as soon as fundamentals hits $a = -0.03$ and keeps paying subsidies for a long time. In contrast, the policy that implements the threshold with lower slope prescribes subsidies only when fundamentals are close to $a = -0.0075$. It is likely that as soon as subsidies start to be paid, they won't be needed anymore, since h will go up and

17

Figure 7: Output fluctuations
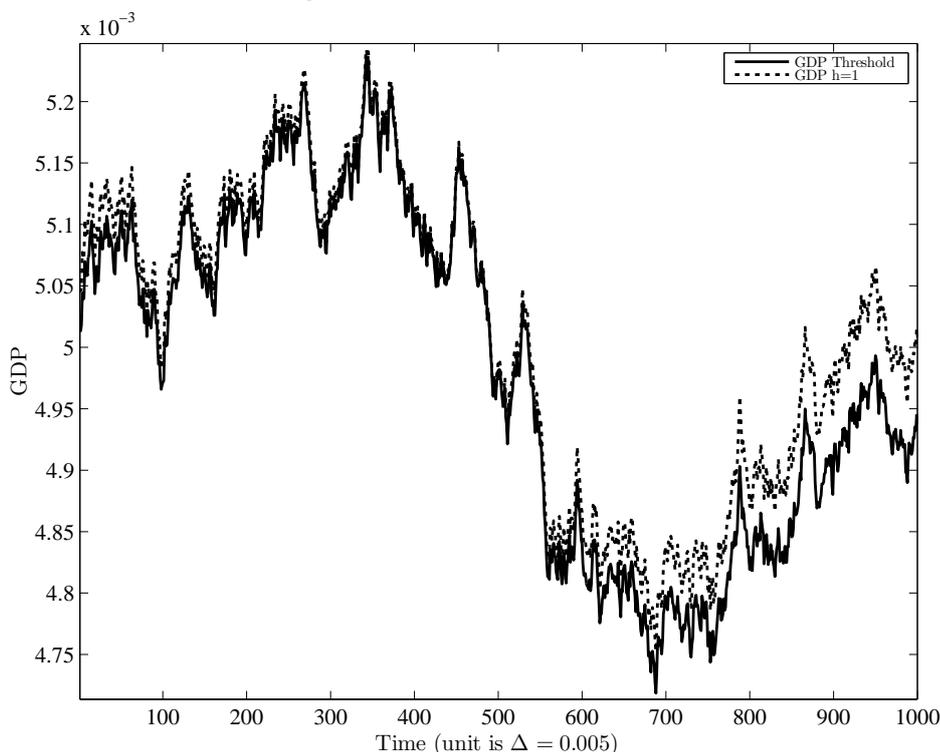
the economy will cross the $\hat{a}$ line. In short, while one policy prescribes large amounts of subsidies to take the economy out of a recession, the other policy does little about it.

The flipside of these policies can be seen when the economy is in good times but heading to a recession, say $a = 0$, $h = 1$ and the productivity parameter is moving down towards -0.045. Both policies prescribe inventives to invest as soon as a crosses their respective thin dashed lines (at about $a = -0.022$). The key difference is that now, if productivity keeps going down, the policy that implements the high-slope threshold soon gives up and the economy falls into a recession as soon as the value of a goes below -0.03. Subsidies will be given again whenever productivity gets past that point. In contrast, the low-slope stimulus policy prescribes a lot more subsidies to be spent in order to prevent the economy from falling into a recession.

This discussion highlights the trade-off involved in the choice of the timing of fiscal stimulus. The subsidies paid according to the low-slope policy to producers when h is high but a is low might prevent an investment slump. Anticipating that, demand expectations for a given variety will be larger, so producers will be more willing to invest – they will require less subsidies to choose *High*. However, the anticipation that a recession will last for a long time if negative shocks to a bring the economy to the no-investment region reduce incentives for investment. The choice of the timing of fiscal stimulus has to take into account that a subsidy for a producer at (a, h) affects not only her incentives to invest but also the incentives for other producers choosing before the economy might reach that

18

Figure 8: Policies
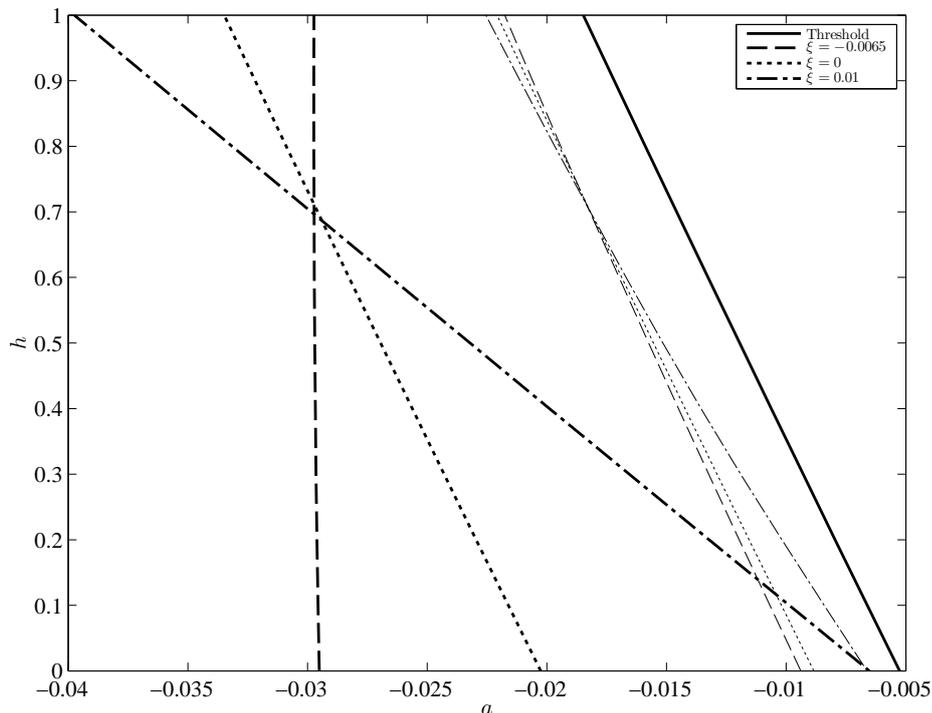
point.[15]

Figure 9 shows the amount of subsidies required to coax agents to invest in two situations: when the government is not intervening in the economy, and when it is implementing the stimulus policy with $\xi = 0$ in Figure 8. Under the stimulus policy, agents expect a larger demand for their goods. As a result, they require less subsidies to invest. The difference between both lines in Figure 9 corresponds to the gains from the increase in expected demand caused by the stimulus policy. In this model, policies that are expected to last and affect other agents are cheaper because staggered investment decisions are strategic complements.

Figure 10 shows our main result. It contrasts the 3 policies in figure 8 and 2 other intermediate policies. The figure suggests that the amount of subsidies required for a given utility level is convex in ξ with a minimum at $\xi = 0$. The cheapest policy is the one that shifts the threshold to the left without rotating it. In Appendix B, we show that the result is robust: the minimum spending policy under alternative parameters prescribes $\xi = 0$ in all specifications we tried.

The best policy does not change the shape of the threshold. Therefore, too much emphasis on preventing an investment slump is sub-optimal, but so is trying too hard to rescue the economy out of a recession. The government should not bias incentives one

---

[15]In the model, the timing of investment is exogenous. That assumption would have some important undesired effects if the stimulus policies analysed here provided incentives for producers to delay investment. However, that does not occur in the model, in equilibrium producers that receive subsidies are actually indifferent between investing or not. Larger subsidies only compensate for a lower productivity and lower expected demand.

Figure 9: One-off and Anticipated Subsidies
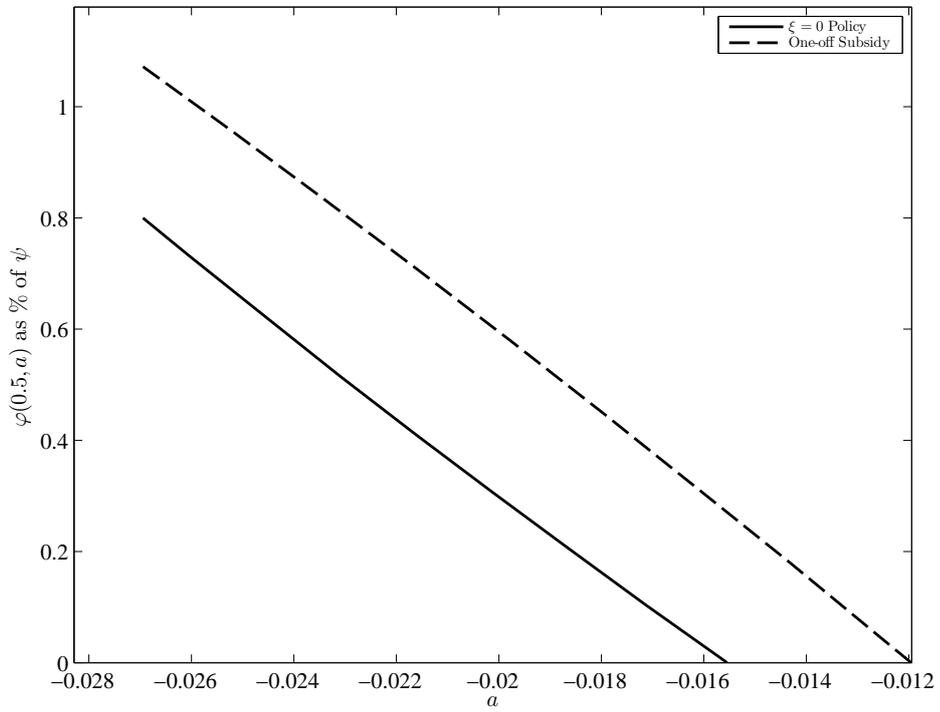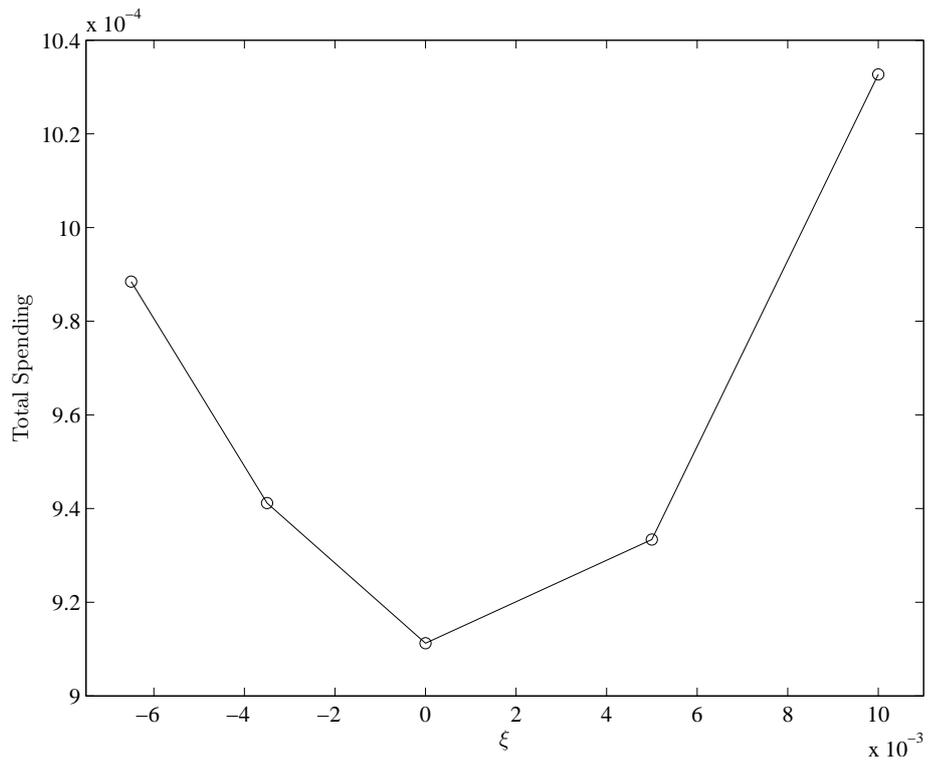
Subsidy with $h = 0.5$

Figure 10: Government spending

way or the other. The best policy subsidizes producers according to the distance of the economy from the original threshold.

Interestingly, the threshold arising from this policy (roughly) coincides with the results of a policy that prescribes a constant subsidy in every state of the world. That has an interesting implication: the maximum level of subsidies should be the same when the economy is about to enter the no-investment region or when it is about to leave it. The government should not pay larger subsidies when a is falling in order to try to avoid an invesment slump or when productivity is relatively high despite low levels of investment.

Intuitively, the policy intervention aims at getting the largest possible mass of producers to invest, since a larger h will increase the returns to others' investments. Those close to the original threshold are nearly indifferent between choosing *Low* and *High*. It turns out that the distance from the original threshold can be seen as a proxy for returns to investment and, consequently, as an indicator of the level of subsidy required to coax a producer to invest. Hence subsidizing a producer far away from the initial threshold is costly, the same amount fo money would bring more investment from producers closer to the threshold. The fact that the economy will be often in some regions and rarely in others does not affect this reasoning, since the frequency the economy is at a given point affects how often subsidies have to be paid at that point but also how often that leads to investment. Nevertheless, it is important to stimulate investiment in a region that affects decisions of producers who would otherwise not invest, but all policies that shift the threshold to the left are doing so.

5   Concluding remarks

This paper proposes a tractable dynamic macroeconomic model with staggered investment decisions where demand expectations affect investment and might lead to coordination failures. The model generates a rich pattern of fluctuations in output and capacity utilization that can be illustrated in a simple diagram with 2 variables: productivity and measure of agents operating at full capacity. However, such simplicity comes at a price, in particular, producers are restricted to a binary set of actions, there are no other relevant state variables, and agents are risk-neutral. Future research might be able to extend this environment and relax some of those assumptions.

The model is consistent with policies that try to restore market confidence when the economy is at a recession, and was used to study the impact of different policies aiming at mitigating coordination failures. The equilibrium threshold for investment features a balance between economic activity and productivity. Stimulus policies should try to shift the threshold without affecting this balance. That means establishing a maximum level of subsidies (or tax cuts) to investment that is independent on productivity, capacity utilization and economic activity.

References

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Khatiwada, S. (2009), *Stimulus Packages to Counter Global Economic Crisis: A Review*, International Institute for Labour Studies.

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

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

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

A   Proofs

A.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 state 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

$$\overline{U}(h_0, a) = \int_0^\infty e^{-(\rho+\alpha)t}\pi(h_t^\uparrow, a)dt - \psi$$

$$= e^a \left(x_H^{\frac{\theta-1}{\theta}} - x_L^{\frac{\theta-1}{\theta}} \right) \int_0^\infty e^{-(\rho+\alpha)t} \left(h_t^\uparrow x_H^{\frac{\theta-1}{\theta}} + (1 - h_t^\uparrow)x_L^{\frac{\theta-1}{\theta}} \right)^{\frac{1}{\theta-1}} dt - \psi.$$

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. □

A.2 Proof of Proposition 2

In order to apply the existence arguments in Frankel & 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 & 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

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

And thus a_t conditional on a_0 is normally distributed with mean

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

and variance

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

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

$$E_0\left[e^{a_t}\right] = \exp\left\{\mu + e^{-\eta t}(a_0 - \mu) + \frac{1}{4}\frac{\sigma^2}{\eta}\left(1 - e^{-2\eta t}\right)\right\}. \tag{12}$$

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

$$\underline{W}(a_0) = \left(x_H^{\frac{\theta-1}{\theta}} - x_L^{\frac{\theta-1}{\theta}}\right)x_L^{\frac{1}{\theta}}\int_0^\infty e^{-(\rho+\alpha)t}E_0\left[e^{a_t}\right]dt - \psi$$

$$> \left(x_H^{\frac{\theta-1}{\theta}} - x_L^{\frac{\theta-1}{\theta}}\right)x_L^{\frac{1}{\theta}}\int_0^1 e^{-(\rho+\alpha)t}\inf\left\{E_0\left[e^{a_t}\right]\right\}_{t\in(0,1)}dt - \psi.$$

By (12), 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

$$\overline{W}(a_0) = \left(x_H^{\frac{\theta-1}{\theta}} - x_L^{\frac{\theta-1}{\theta}}\right)x_H^{\frac{1}{\theta}}\int_0^\infty e^{-(\rho+\alpha)t}E_0\left[e^{a_t}\right]dt - \psi$$

$$< \left(x_H^{\frac{\theta-1}{\theta}} - x_L^{\frac{\theta-1}{\theta}}\right)x_H^{\frac{1}{\theta}}\left(\int_0^Q e^{-(\rho+\alpha)t}\left(\sup\left\{E_0\left[e^{a_t}\right]\right\}_{t\in(0,Q)}\right)dt\right.$$

$$\left. + \int_Q^\infty e^{-(\rho+\alpha)t}\left(\mu + \frac{\sigma^2}{4\eta}\right)dt\right) - \psi.$$

By (12), we have that $\lim_{a_0\to-\infty}\sup\left\{E_0\left[e^{a_t}\right]\right\}_{t\in(0,1)} = 0$. 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 & 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.

$\square$

A.3 Proof of Proposition 3

Before we prove Proposition 3, it is useful to establish the following results.

Lemma 1. *Let $\hat{a}_t$ be the following latent variable*

$$\hat{a}_t = \begin{cases} a_t & \text{if } a_t < M \\ \tilde{a} & \text{otherwise} \end{cases},$$

where a_t is given by $da_t = \eta_t(\mu - a_t)dt + \sigma dZ_t$, with η_t given by (9). Then $\lim_{a_\tau \to \infty} E_\tau\left[e^{\hat{a}_t}\right] = e^M$ and $\lim_{a_\tau \to -\infty} E_\tau\left[e^{\hat{a}_t}\right] = 0$, for every $t > \tau$ and every $\tau \geq 0$. Moreover, e^M and zero are, respectively, an upper bound and a lower bound for $r_\tau(t) \equiv E_\tau\left[\hat{a}_t\right]$, for every $\tau \geq 0$.

Proof. First assume $\tau < T$. It follows from (A.2) that when $\tau < t < T$, $a_t | a_\tau$ has a normal distribution with mean and variance given by (A.2) and (A.2), respectively. In that case we have

$$E_\tau\left[e^{\hat{a}_t}\right] = \frac{1}{\Sigma_t\sqrt{2\pi}} \int_{-\infty}^{M} \exp\left\{a_t - \frac{1}{2}\left(\frac{a_t - \mu - e^{-\eta t}(a_\tau - \mu)}{\Sigma_t}\right)^2\right\} da_t$$
$$+ \left(1 - \Phi\left(\frac{M - \mu - e^{-\eta t}(a_\tau - \mu)}{\Sigma_t}\right)\right)e^M, \quad (13)$$

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

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,

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

where the last equality follows from equation A.2. Moreover,

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

We can show that $a_t | a_\tau$ follows a normal distribution, [16] and so,

$$E_\tau\left[e^{\hat{a}_t}\right] = \frac{1}{\sqrt{(\sigma^2(t - T) + \Sigma_T^2)\, 2\pi}} \int_{-\infty}^{M} \exp\left\{a_t - \frac{1}{2}\left(\frac{a_t - \mu - e^{-\eta T}(a_\tau - \mu)}{\sqrt{(\sigma^2(t - T) + \Sigma_T^2)}}\right)^2\right\} da_t$$
$$+ \left(1 - \Phi\left(\frac{M - \mu - e^{-\eta T}(a_\tau - \mu)}{\sqrt{(\sigma^2(t - T) + \Sigma_T^2)}}\right)\right)e^M. \quad (14)$$

Notice that both equations 13 and 14 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 equations 13 and 14 completes the proof for the case where $\tau < T$. The proof for the case where $\tau \geq T$ is very similar and

---

[16] We know that $a_T | a_\tau \sim N(\mu + e^{-\eta T}(a_\tau - \mu), \frac{\sigma^2}{2\eta}(1 - e^{-2\eta T}))$ 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.

therefore omitted. The last sentence of the Lemma comes directly from the inspection of equations 13 and 14 . □

Lemma 2. *Suppose $\sigma > 0$, the mean reversion parameter η_t is given by (9) and the relative payoff of investing is given by (10). 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 $\hat{\pi}(h_t, a_t) = \pi(h_t, \hat{a}_t)$, 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

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

But $\inf \left\{ E_0 \left[e^{\hat{a}_t} \right] \right\}_{t\in(0,1)}$ 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

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

But $\sup \left\{ E_0 \left[e^{\hat{a}_t} \right] \right\}_{t\in(0,Q)}$ 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 3. Since we have proved the existence of dominance regions, it follows from Theorems 1 and 4 in Frankel & Burdzy (2005) (since our model is a special case of their model).

□

B   Robustness

We run our policy exercise using other parameters. In order to get sufficiently different statistics from those obtained in the baseline calibration, we did not try to match the same

statistics as before. It led to very different statistics in most specifications. The only thing we imposed was that the economy experienced both crisis and recessions and that the policies did not completely shut down recessions. Table 2 reports the parameters chosen and the implied average time in recession with no intervention and under the intervention with $\xi = 0$. The parameter c for the parallel policy - which defines the distance between the original threshold and the parallel policy threshold - is not reported, since all the relevant information about this parameter is contained in the last two lines of Table 2. The values of ξ were chosen to contemplate the cases of an almost vertical threshold and a threshold were the government almost does not pay subsidies when fundamentals are picking up and the economy is leaving from a situation with $h = 0$, as in Figure 8. Figure 11 reports the results for government spending, showing that our results did not change qualitatively.

Table 2: Robustness check parameters

Parameter	Symbol	Specification					
		1	**2**	**3**	**4**	**5**	**6**
Production state *High*	x_H	1.1	1.1	1.1	1.05	2	1.1
Production state *Low*	x_L	1	1	1	1	1	1
Elasticity substitution	θ	6	6	6	6	6	6
Fixed cost of investing	ψ	0.0408	0.0806	0.0408	0.0413	0.795	0.0278
Mean of fundamental process	μ	0	0	0	0	0	0
Arrival rate of Poisson Process	α	2	1	2	1	1	3
Standard deviation of shocks	σ	0.03	0.1	1	0.03	0.3	0.03
Discount rate	ρ	0.03	0.03	0.03	0.01	0.03	0.03
Mean reversion intensity	η	0.7	0.7	2	1.5	0.7	0.7
Time interval lenght	Δ	0.005	0.005	0.005	0.005	0.005	0.005
Average time in recession (%)	-	26.15	43.20	33.39	38.50	37.70	55.95
Average time in recession parallel policy (%)	-	10.15	35.96	14.03	4.79	16.88	24.61

Figure 11: Government spending in each specification