DISCUSSION PAPER SERIES

No. 9977

DEMAND EXPECTATIONS AND THE TIMING OF STIMULUS POLICIES

Bernardo Guimarães and Caio Machado

INTERNATIONAL MACROECONOMICS

Centre for Economic Policy Research

www.cepr.org

Available online at: **www.cepr.org/pubs/dps/DP9977.php**

ISSN 0265-8003

DEMAND EXPECTATIONS AND THE TIMING OF STIMULUS POLICIES

Bernardo Guimarães, Fundacao Getulio Vargas and CEPR
Caio Machado, Fundacao Getulio Vargas

Discussion Paper No. 9977
May 2014

Centre for Economic Policy Research
77 Bastwick Street, London EC1V 3PZ, UK
Tel: (44 20) 7183 8801, Fax: (44 20) 7183 8820
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 9977

May 2014

ABSTRACT

Demand expectations and the timing of stimulus policies*

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 and E62
Keywords: coordination, demand expectations, fiscal stimulus and timing frictions

Bernardo Guimarães
São Paulo School of Economics
Fundacao Getulio Vargas
Rua Itapeva, 474
Sao Paulo
BRAZIL

Email: bernardo.guimaraes@fgv.br

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=160881

Caio Machado
São Paulo School of Economics
Fundacao Getulio Vargas
Rua Itapeva, 474
Sao Paulo
BRAZIL

Email: caiohm@gmail.com

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=178452

*We thank Braz Camargo, Francesco Caselli, Felipe Iachan, V. Filipe Martins-da-Rocha, Alison Oliveira, Mauro Rodrigues and Mark Wright as well as seminar participants at LACEA 2013 (Mexico City), Sao Paulo School of Economics – FGV, U Sao Paulo, SBE Meeting 2013 (Igua¸cu) and the XVII Workshop of International Economics and Finance (San Jose). Bernardo Guimaraes gratefully acknowledges financial support from CNPq. Caio Machado gratefully acknowledges financial support from CAPES and FAPESP.

Submitted 02 May 2014

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.

The dynamic coordination problem among firms has implications to the optimal timing of stimulus policies. 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 for investment vary with economic activity and fundamentals? This paper develops a simple macroeconomic model that captures that dynamic coordination problem among producers and can be used to answer these questions.

The model features monopolistic competition and staggered investment decisions. 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, as in Kiyotaki (1988). Producers of each variety receive investment opportunities according to a Poison clock. 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 about others' future decisions.

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

[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.

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

In order to study optimal policy in this model, we characterize the central planner's problem and obtain analytical results that relate the planner's choices and the decentralized equilibrium. We then calibrate the model and solve it numerically for cases we cannot get analytical results. We are thus able to answer questions about the timing of stimulus policies in a model where expectations about others' actions ('producers' confidence') are pinned down in the model and 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. Stimulus policies have a direct effect on agents' incentives for investment but also an indirect effect through beliefs about others' actions. But what is the optimal timing of stimulus policies?

The question can be posed in the following way: the equilibrium threshold can be represented as a curve in a two-dimensional space, with (log) 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 producers will require lower productivity or demand to invest. Besides translating the threshold to the left, how should it try to rotate the threshold? Should it try to stimulate investment primarily when most producers are still operating at full capacity, despite relatively weak productivity, 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 harder to avoid a recession when productivity is very low, or putting more emphasis on rescuing the economy when the mass of agents investing is very low are both inefficient. 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 affected by the government intervention.

The externalities generated by investment are proportional to the private gains from investing and related to the markup charged by firms. That helps understanding why the equilibrium threshold should not be rotated. The slope of the threshold reflects the relative effects of expected demand and productivity perceived by a producer. Since the externalities from investment are proportional to a producer's gains, there is no reason to affect this slope.

Another implication of the relation between the externalities and the private gains from investment is that the distortion in producers' decisions can be corrected by a constant subsidy. That is also true if the planner faces costs to monitor investment. In case the planner has limited resources to subsidize investment, an analogous result arises: the optimal policy establishes a maximum level of subsidies to investment that is independent of productivity, capacity utilization and economic activity.

The demand externalities that play a key role in this paper are in the seminal contributions by Blanchard and 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.[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 uniqueness result in Frankel and Pauzner (2000) requires very small mean reversion,

[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.

but Frankel and Burdzy (2005) generalize some of the results in Frankel and Pauzner (2000) for more general stochastic processes.[4] 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.[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 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 part of the 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 study the dynamic interplay between economic activity and productivity.

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.

[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 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.

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.

The assumptions on technology aim at modelling staggered investment decisions in a simple and tractable way. Investment is a binary decision, a payment of a sunk cost that reduces marginal cost of production. 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. There are 2 production regimes, a *High*-capacity regime and a *Low*-capacity regime. 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. Agents get a chance to switch regimes according to a Poisson process with arrival rate α. 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 implies a one-off cost ψ in units of the composite good (ψ is a stock).

Choosing the *High* regime is 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 regime until the next investment opportunity arises, which captures the idea that firms cannot change their capital level overnight.[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.

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.[8]

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 The agent's problem

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:

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

[7] 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.

[8] The assumption of linear utility implies that consumption smoothing plays no role in the model, all results come from investment decisions.

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

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

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 the *High* regime.

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

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

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 the *High* regime and agents locked in the *Low* regime when the economy is at (h_t, a_t). Then, using $y_{Lt} = e^{a_t} x_L$ and $y_{Ht} = e^{a_t} x_H$,

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

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

$$\int_{\tau}^{\infty} e^{-(\rho+\alpha)(t-\tau)} E_{\tau}[\pi(h_t, a_t)]dt \geq \psi. \tag{8}$$

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 $(\rho + \alpha)$.[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 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 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. Small subsidies to investment in the multiplicity region have no effects on beliefs.

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

Figure 1: Equilibria without shocks

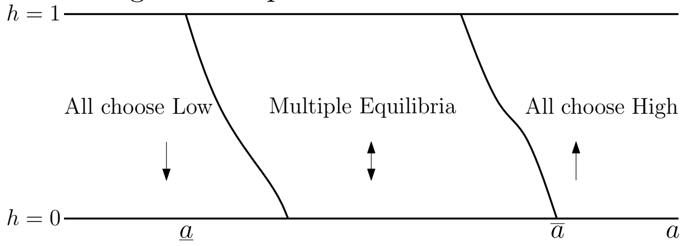

2.4 The case without mean reversion

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.

We start with the special case where $\eta = 0$, i.e., the fundamental process has no mean reversion. The model is a particular case of Frankel and Pauzner (2000) and we can apply Theorem 1 in their paper to show there is a unique rationalizable equilibrium where agents play according to a decreasing threshold $a^*(h)$. Figure 2 shows an example of equilibrium in the model.

Figure 2: Equilibrium with shocks

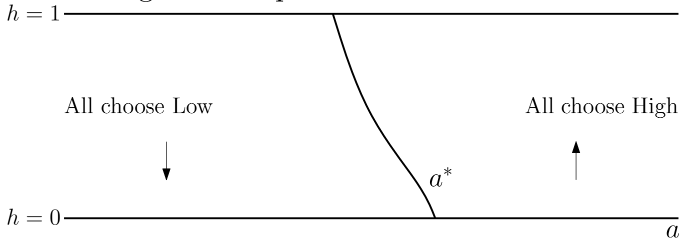

2.5 General case

We now turn to the general case where productivity varies over time, and there is mean reversion $\eta > 0$.

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 and 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, as in Figure 2. 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.

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

$$V(a,h,\tilde{a}) = \int_0^\infty e^{-(\rho+\alpha)t} E[\pi(h_t,a_t)|a,h,\tilde{a}]dt - \psi, \tag{9}$$

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.

2.5.1 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 and Burdzy (2005), we need to make two changes in the model. First, the diffusion process for a_t is given by (3), but the mean-reversion parameter η_t varies over time so that

$$\eta_t = \begin{cases} \eta & \text{if } t < T \\ 0 & \text{otherwise} \end{cases}, \tag{10}$$

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

$$\hat{\pi}(h,a) = \begin{cases} \pi(h,a) & \text{if } a < M \\ \pi(h,M) & \text{otherwise} \end{cases}, \tag{11}$$

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 3 (Uniqueness, Frankel and Burdzy (2005))**.** *Suppose $\sigma > 0$, the mean reversion parameter η_t is given by (10) and the relative payoff of investing is given by (11). 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 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 (9) 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.

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. The natural questions that emerge are: what would a benevolent central planner choose, i.e., in which states would he invest? How would he implement this first-best?

Suppose we have already verified that the central planner's choice is to invest according to a threshold (we will show this later). Since we expect that agents underinvest in equilibrium, one should expect that the planner's threshold would be to the left of agent's threshold, meaning that the planner will invest in more states than agents do. But how would this threshold look like?

Figure 3: Policy alternatives

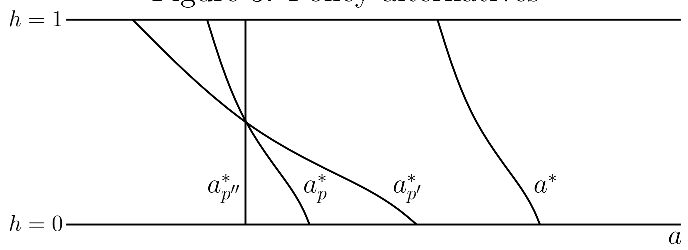

The thresholds a^*_p, $a^*_{p'}$ and $a^*_{p''}$ depicted in Figure 3 correspond to different policy objectives. A stimulus policy that implements the threshold a^*_p is not particularly concerned with either preventing the economy from falling into a recession or rescuing the economy when productivity picks up. In contrast, a policy that implements $a^*_{p'}$ prescribes 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. In the other extreme, a central planner could implement a vertical threshold ($a^*_{p''}$ in Figure 3), implying little effort in preventing recessions but a lot of effort to stimulate the economy as soon as possible when fundamentals pick up.

Each of the thresholds a^*_p, $a^*_{p'}$ and $a^*_{p''}$ correspond to different timings of stimulus policies. For example, consider the economy is in some state with $h = 1$ and to the right of both a^*_p and $a^*_{p'}$. Suppose productivity begins to fall. The central planner that chooses according to the threshold a^*_p will give up investing sooner than the central planner that implements $a^*_{p'}$. The latter will invest even when fundamentals are relatively lower. Now suppose the economy got into recession, a is to the left of both planner's thresholds and h is low. If productivity starts to increase, the social planner that implements a^*_p will start to invest earlier than the central planner that chooses according to $a^*_{p'}$.

3.1 The Central planner problem

We now characterize the central planner's problem. The planner maximizes expected welfare, given by:

$$E_\tau(W) = E_\tau \int_\tau^\infty e^{-\rho(t-\tau)} \left(Y(h,a) - \alpha\psi I(t)\right) dt \tag{12}$$

where $Y(h,a)$ is given by (6) 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

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

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

$$\int_\tau^\infty e^{-\rho(t-\tau)} E_\tau \left(\frac{dY(h,a)}{dh} \alpha e^{-\alpha(t-\tau)} \right) dt - \alpha\psi = 0$$

Since

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

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

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

where $\pi(h_t, a_t)$ is given by (7).

The expression for the planner decision is thus very similar to the agent's problem. The only difference is the term $\theta/(\theta-1)$ multiplying the benefit from investing.

In order to get some intuition, consider an expression for welfare that is given by the sum of individual agents' payoffs (which depend on others' actions as well): $W(h,a) = h u_H(h,a) + (1-h) u_L(h,a)$. We get that:

$$\frac{dW}{dh} = \left(u_H(h,a) - u_L(h,a)\right) + \left(h \frac{\partial u_H(h,a)}{\partial h} + (1-h) \frac{\partial u_L(h,a)}{\partial h} \right)$$

The agent considers only the first term in brackets, and not the externality on others' payoffs. However, in this case it turns out that:

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

Intuitively, the agent is taking into account the effect on its income but not the positive effect on others that come from her selling at a lower price. The lower the price elasticity, the larger the externality. Importantly, the externality is just a fraction of the agent's payoff from investing. In consequence, for a given return to investment, the externality is independent of h or a.

3.1.1 The timing of stimulus policies

The key implication of (13) for the timing of stimulus policies is in the next proposition:

Proposition 4. *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] When $\eta \to 0$, 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 (13) is satisfied. Multiplying both sides of (13) by $(\theta-1)/\theta$ yields the condition for an agent to invest in (9) 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 13 as

$$\int_\tau^\infty e^{-(\rho+\alpha)(t-\tau)}E_\tau\left[e^{\left(a_t+\log\left(\frac{\theta}{\theta-1}\right)\right)}g(h_t)\right]dt\geq\psi \tag{14}$$

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 (9) for the agents' decisions in the decentralized equilibrium (in the (a,h)-space). Moreover, if $\eta=0$, 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 4: Planner's problem

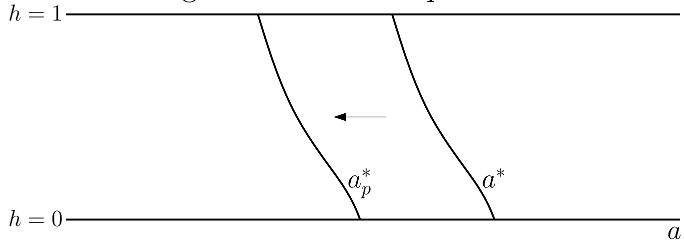

The solution to the planner's problem in (13) considers the benefits from investing by an individual producer multiplied by a constant larger than 1. Hence the only problem with the individual decision is that it requires a benefit from investing that is too high. A constant subsidy takes care of this problem.

As shown in (14), the expressions for the planner's problem in (13) is the same as the solution for the decentralized problem in (9) when a constant is added to the log of productivity. Without mean revertion in the process for a_t, that implies a translation of the equilibrium threshold, where that constant is subtracted from the productivity threshold. The slope of the equilibrium threshold a^* reflects the relative effects of expected demand and productivity on the expected profits from investing, and there is no reason for the planner's solution to affect this balance.

The argument does not work when the process for a_t exhibits mean reversion because in this case 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. Thus with mean reversion, the planner's decision is still given by (13), but Proposition 4 does not apply. However, we show in Section 4 that for a reasonable amount of mean reversion, the results are essentially unchanged.

The next proposition compares the return of investment and the gross product when the economy is about to enter (leave) an investment slump under the planner's threshold.

Proposition 5. *When $\eta \to 0$ we have that:*

1. $\frac{\partial Y(1,a_p^*(1))}{\partial h} > \frac{\partial Y(0,a_p^*(0))}{\partial h}$*: when the economy is at the planner's threshold, the instantaneous return of investing is higher when $h = 1$ than it is when $h = 0$.*

2. $Y(1, a_p^*(1)) > Y(0, a_p^*(0))$*: when the economy is at the planner's threshold, the gross product is higher when $h = 1$ than it is when $h = 0$.*

Proof. See Appendix A. □

Proposition 5 states that the planner requires a smaller instantaneous return to invest when the economy is at an investment slump and output is low. The intuition is that when economic activity is high, the planner knows it might fall and therefore is willing to invest only if the current return is really worth it. Conversely, when the economy is in recession, the planner expects activity to go up in the future (it can not go below zero), and therefore is willing to invest today even if the current return is not that high.

This idea seems to be in line with the Keynesian intuition that when economic activity is low, the government should be more willing to stimulate investment in order to help coordinate investment, even if the return to investment is not so large. However, as Proposition 4 shows, the government should not pay larger subsidies when economic activity is low, the planner's problem prescribes a constant subsidy. The conundrum is solved by noting that agents also internalize the possibility of an increase in h when the economy is at an investment slump. A corollary of Proposition 5 is that $\pi(1, a^*(1)) > \pi(0, a^*(0))$: agents also require a smaller instantaneous return to invest when h is low. The only thing they do not take into account is the positive effect of their investment on others' payoffs, but Proposition 4 shows that is proportional to the private return to investment, it does not vary with economic activity.

3.1.2 The planner's problem with a monitoring cost

One potential objection to this analysis is that it ignores the costs of subsidizing investment. When the economy is close to the planner's threshold a^*_P, the planner is close to indifferent between investing or not, so any cost to subsidize investment would make the planner choose no investment.

Subsidizing investment requires that the planner monitors firms' investments. One simple way to capture that is to assume the planner faces a cost c for each unit of investment it subsidizes. Assume $c < \psi/(\theta - 1)$.[10] Then the planner chooses to maximize:

$$E_\tau(W) = E_\tau \int_\tau^\infty e^{-\rho(t-\tau)} \left(Y(h, a) - \alpha(\psi + c)I(t)\right) dt$$

which means the planner chooses to invest if:

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

which is basically what we got before, and an argument similar to the one in Proposition 4

[10] In case $c \geq \psi/(\theta - 1)$, monitoring costs are so large that subsidizing investment is never worth it.

shows that it leads to different translation of the threshold, but no rotation. The paralell shift of the equilibrium threhsold is robust to the inclusion of a monitoring cost.

3.2 Minimal spending policies

Proposition 4 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: suppose the government has a limited budget to spend with stimulus policies, cannot implement the planner's solution, but can commit to a given stimulus policy. Would it do anything different? For example, it could commit to subsidize investment when h is high but a is low, hoping agents would expect larger demand and invest more. Would it do so? Would it rotate the threshold?

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.

We now focus on *minimal spending policies*, which are the cheapest way to subsidize producers that implements a certain threshold.[11]

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

$$\varphi(h, a) = \begin{cases} \psi - \int_0^\infty e^{-(\rho+\alpha)t} E[\pi(h_t, a_t)|a, h, a_p^*]dt & \text{if } a_p^*(h) \leq a \leq \hat{a}(h) \\ 0 & \text{otherwise} \end{cases}. \tag{15}$$

Figure 5 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 5). The

[11] 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 objective 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.

Figure 5: Example of minimal spending policy

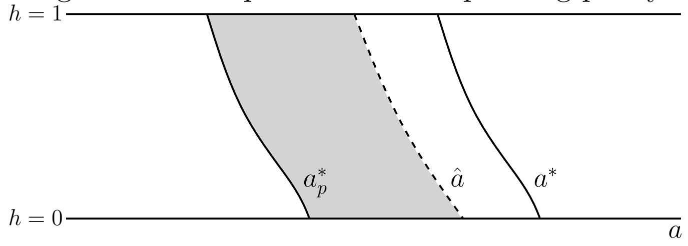

subsidy $\varphi(h, a)$ makes her indifferent between choosing *High* and *Low* given others will play according to a_p^*. Under those beliefs, playing according to a_p^* is a best response under this policy, so a_p^* is an equilibrium. Interestingly, no subsidies are needed in the area between $\hat{a}$ and a^*.[12]

For a given initial condition (h_τ, a_τ), the planner chooses a threshold $a_p^*(h)$ to maximize (12) subject to

$$\int_\tau^\infty e^{-\rho(t-\tau)} E_\tau \left[\varphi(h, a)\right] \leq C$$

where C is a constant and $\varphi(h, a)$ is given by (15).

In principle, the solution to this problem would depend on where the planner starts. In order to understand what a government would do in case of commitment, we consider the solution of a planner born in some random state, being the probability of being in a given state proportional to the frequency of that state (once the policy is implemented). That is similar in spirit to the idea of pre-commitment under a timeless perspective in Woodford (1999).

4 Numerical results

We now calibrate and solve the model numerically. That has three objectives: (i) to get a better understanding of the workings of the model; (ii) to check whether a reasonable amount of mean reversion affects the solution to the central planner problem; and (iii) to check whether commitment to minimal spending policies could lead to threshold rotation.

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

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

[12]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.

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.

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.[13]

4.2 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.[14]
- The economy stays 30% of time at the left of the threshold, that is, agents are not investing 30% of the time, approximately.[15]
- 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.[16]

[13] 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.

[14] 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.

[15] 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.

[16] 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 (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 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

4.3 Results

4.3.1 Equilibrium

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

Figure 6: Estimated threshold

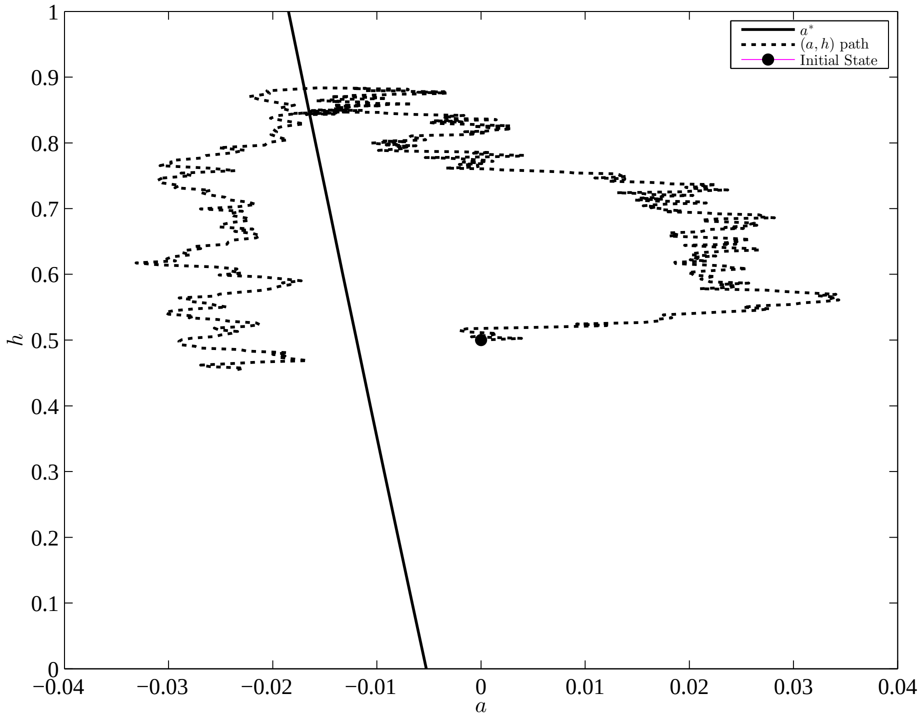

always equal to 1, because low values of the productivity parameter a lead to periods of low expected demand where agents choose not to invest. In this model, policies can do nothing about the exogenous movements in a but can increase the region where agents invest. Investment subsidies can bring output closer to the $h = 1$ curve.

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

4.3.2 The planner's problem

The result concerning the central planner's choice in Proposition 4 does not apply when the stochastic process of productivity is mean reverting. In what follows, we show that for reasonable values of the mean reversion parameter, the threshold implemented by the central planner is very close to a parallel shift of the equilibrium threshold.

Figure 8 shows the estimated planner's threshold together with a translation of the original threshold. As one can see in the figure, both thresholds practically coincide. The difference between the slope (measured here as the distance in the horizontal axis of a

Figure 7: Output fluctuations

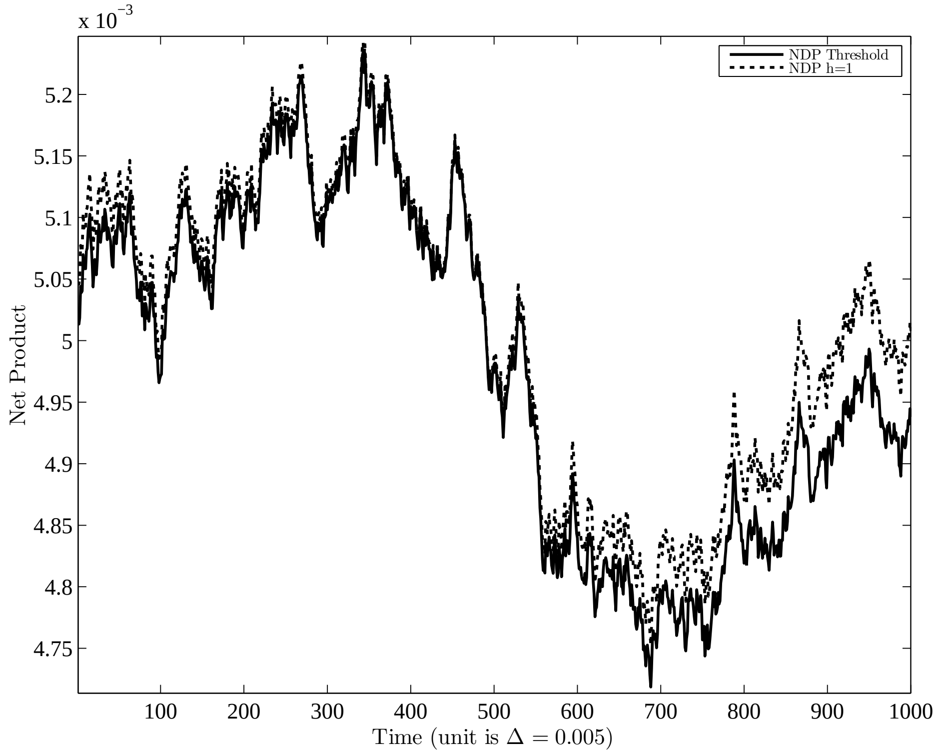

Figure 8: Planner's threshold with mean reversion

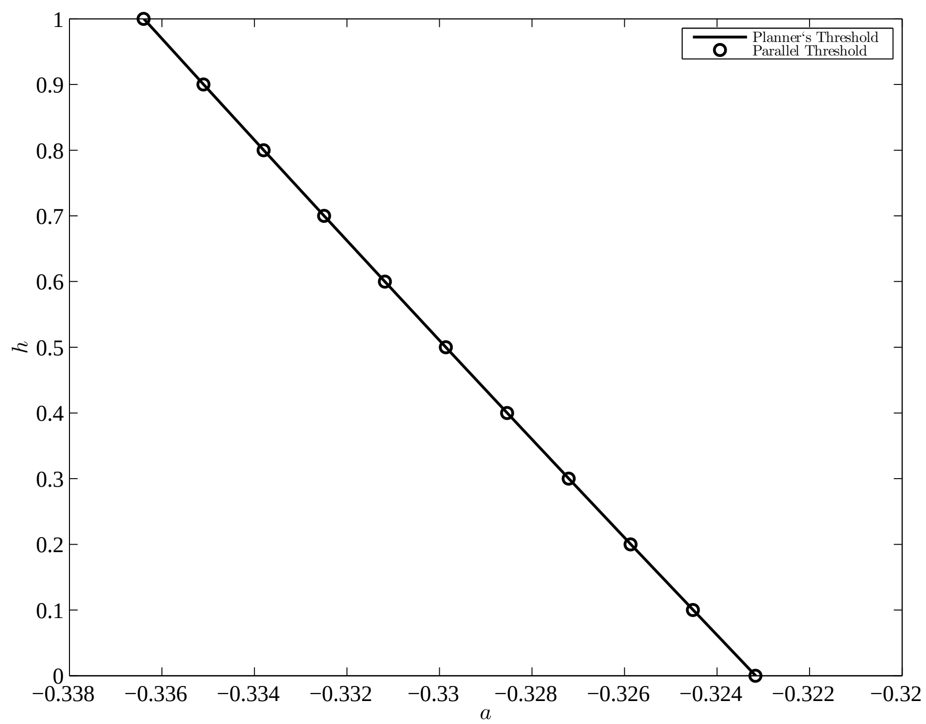

threshold extreme points, i.e., $|a(1) - a(0)|$) of the planner's threshold and the agent's threshold was no more than 0.1% of the slope of the agent's threshold. In appendix B we show that results are robust to alternative parametrizations: we did the same exercise under different specifications and the difference between the slope of both thresholds was on average around 2%.

Another interesting feature of the mean reverting process is that the shift in the threshold is larger than in the case without mean reversion. That is because the productivity parameter is expected to return to its mean in the near future, so the planner is more willing to invest for a given level of fundamentals.

4.3.3 Minimal spending policies

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.

For the sake of computational efficiency, instead of solving the problem proposed in Section 3.2, we approximate the solution of the dual of problem of minimizing the government spending subject to some utility level. Therefore, for some random initial state (h_τ, a_τ) the government chooses a threshold $a^*_p(h)$ to minimize

$$\int_\tau^\infty e^{-\rho(t-\tau)} E_\tau \left[\varphi(h, a)\right] dt$$

subject to

$$E_\tau[W] \geq \overline{U},$$

where $\overline{U}$ is a constant and $\varphi(h, a)$ is given by (15).

Initially, we set a grid for ξ that includes zero and a value for c.[17] 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 5) given

[17]We chose a value c such that under the parallel stimulus policy the economy stayed approximately 12% of the time to the left of threshold.

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 5). 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 (15).

We now turn to the comparison of different policies. Figure 9 shows different stimulus policies corresponding to different values of ξ, that deliver the same welfare improvement. The average duration of a recession decreases with the slope of the threshold. In the examples we tested, it varies from about 3 quarters in the policy represented by the threshold with a higher slope to almost 5 quarters in the policy corresponding to the lower slope.[18]

Figure 9: Policies

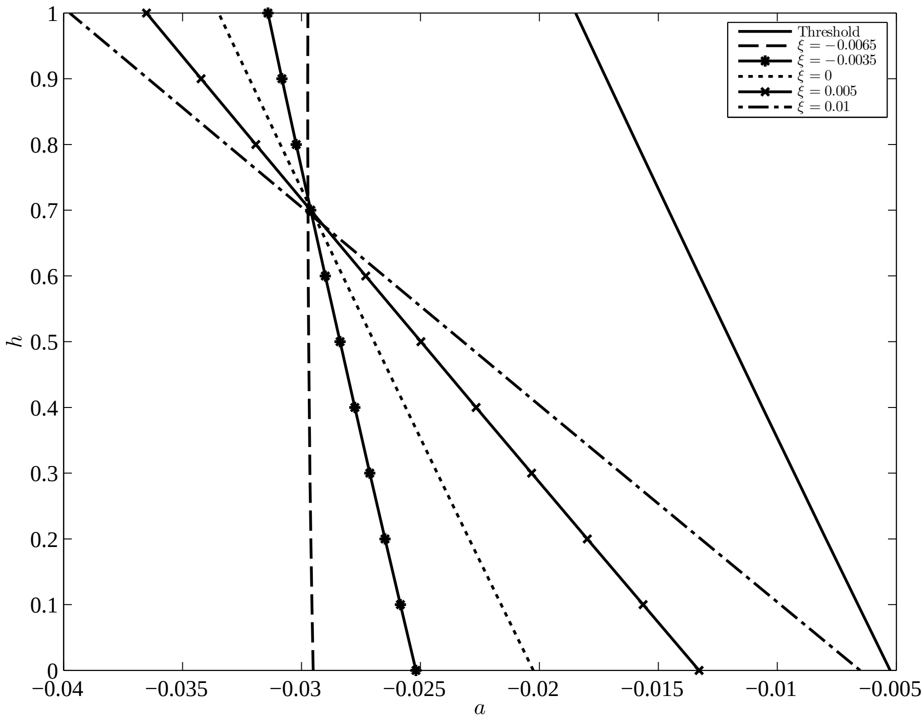

Figure 9 helps understanding the different effects of each policy. Consider an economy in a recession, with $h = 0$ and $a = -0.045$, and a_t is moving up towards zero. The stimulus policy that implements the threshold with higher slope (the almost vertical line in Figure 9) kicks in as soon as fundamentals 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 $a = -0.0075$. 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. The key difference is that now, if

[18] A comparison between Figures 9 and 8 shows that the planner's threshold is very far from the policies simulated in this section, so that the economy experiences almost no recession under the planner's solution.

productivity keeps going down, the policy that implements the high-slope threshold soon gives up investing 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 point.[19]

Figure 10: One-off and Anticipated Subsidies

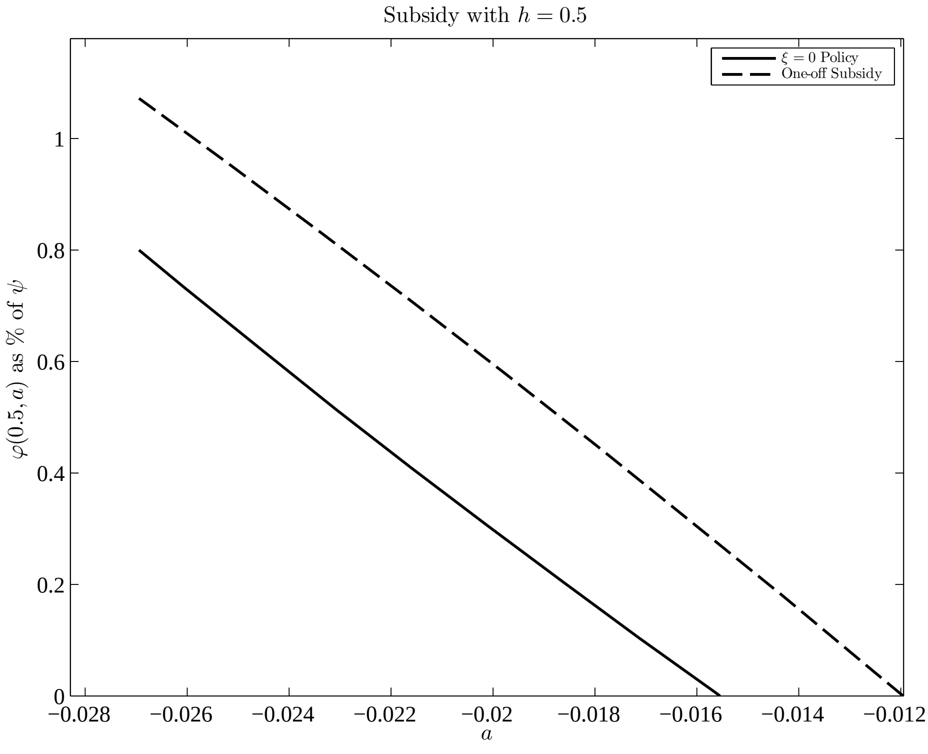

Figure 10 shows the amount of subsidies required to coax agents to invest in two situa-

[19]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 case of minimal spending policies, 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.

tions: when the government is not intervening in the economy, and when it is implementing the stimulus policy with $\xi = 0$ in Figure 9. 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 10 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 owing to the strategic complementarities in investment decisions.

Figure 11 shows the amount of spending needed to implement each of the 5 policies depicted in Figure 9 and confirms the main result of this paper. The amount of subsidies required for the obtention of 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.

Figure 11: Government spending

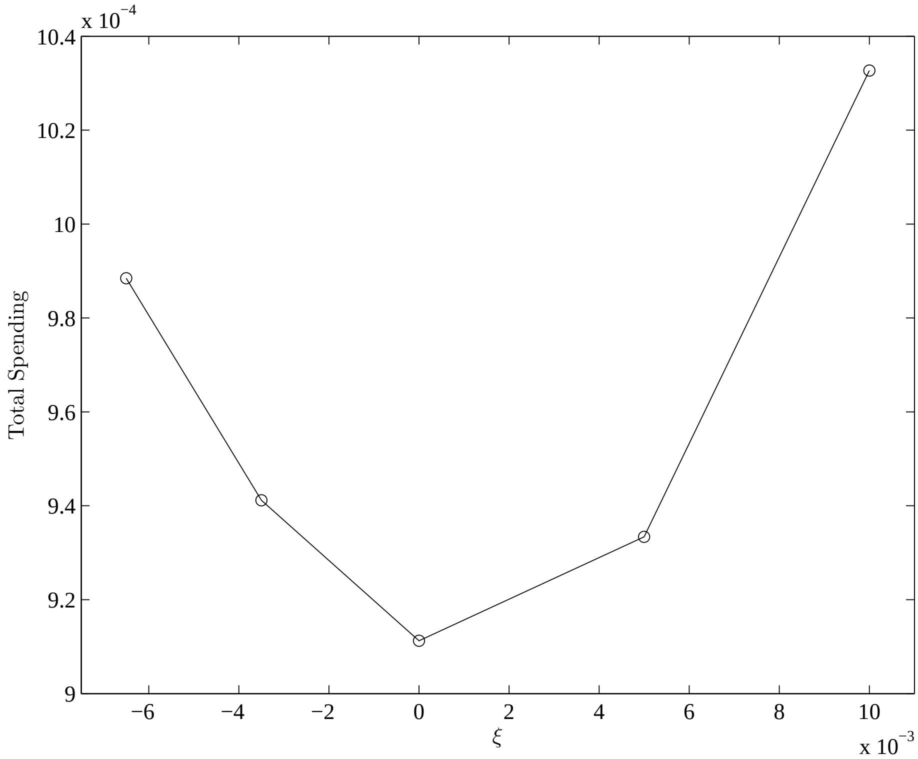

5 Concluding remarks

This paper proposes a tractable dynamic macroeconomic model with staggered investment decisions where demand expectations affect investment and might lead to coordination failures. Stimulus policies affect beliefs about the probability of an investment slump in a continuous and intuitive way. 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, agents are risk-neutral and expectations are pinned down by only 2 variables. 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. In other words, 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. Too much emphasis on preventing an investment slump is sub-optimal, but so is focusing exclusively on productivity.

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.

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.

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.

He, Z. and 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. 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.

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

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

Woodford, M. (1999). Commentary: How should monetary policy be conducted in an era of price stability?, *New challenges for monetary policy* **277**: 316.

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 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. □

A.2 Proof of Proposition 2

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

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

And thus a_t conditional on a_0 is normally distributed with mean

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

and variance

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

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

$$E_0\left[e^{a_t}\right] = \exp\left\{\mu + e^{-\eta t}(a_0 - \mu) + \frac{1}{4}\frac{\sigma^2}{\eta}\left(1 - e^{-2\eta t}\right)\right\}. \tag{16}$$

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 (16), 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

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

By (16), 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 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.

□

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*

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

where a_t is given by $da_t = \eta_t(\mu - a_t)dt + \sigma dZ_t$, with η_t given by (10). 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 $r_\tau(t) \equiv E_\tau[\hat{a}_t]$, 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

$$E_\tau\left[e^{\hat{a}_t}\right] = \frac{1}{\Sigma_t\sqrt{2\pi}}\int_{-\infty}^{M} \exp\left\{a_t - \frac{1}{2}\left(\frac{a_t - \mu - e^{-\eta t}(a_\tau - \mu)}{\Sigma_t}\right)^2\right\} da_t + \left(1 - \Phi\left(\frac{M - \mu - e^{-\eta t}(a_\tau - \mu)}{\Sigma_t}\right)\right) e^M, \quad (17)$$

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,

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

where the last equality follows from (A.2). Moreover,

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

We can show that $a_t|a_\tau$ follows a normal distribution, [20] and so,

$$E_\tau\left[e^{\hat{a}_t}\right] = \frac{1}{\sqrt{\left(\sigma^2(t-T) + \Sigma_T^2\right)2\pi}}\int_{-\infty}^{M} \exp\left\{a_t - \frac{1}{2}\left(\frac{a_t - \mu - e^{-\eta T}(a_\tau - \mu)}{\sqrt{\left(\sigma^2(t-T) + \Sigma_T^2\right)}}\right)^2\right\} da_t + \left(1 - \Phi\left(\frac{M - \mu - e^{-\eta T}(a_\tau - \mu)}{\sqrt{\left(\sigma^2(t-T) + \Sigma_T^2\right)}}\right)\right) e^M. \quad (18)$$

Notice that both (17) and (18) 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 (17) and (18) completes the proof for the case where

[20] 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 $E_\tau[a_t|a_T]$ is linear on a_T and $Var_\tau[a_t|a_T]$ 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.

$\tau < T$. The proof for the case where $\tau \geq T$ is very similar and therefore omitted. The last sentence of the Lemma comes directly from inspection of (17) and (18). □

Lemma 2. *Suppose $\sigma > 0$, the mean reversion parameter η_t is given by (10) and the relative payoff of investing is given by (11). 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 and Burdzy (2005) (since our model is a special case of their model).

□

A.4 Proof of Proposition 5

First statement

Suppose the planner is investing according to a threshold $a_p^*(h)$. For any initial state, there is always a positive probability that the economy will reach states where $h_t < 1$ or $h_t > 0$. That yields the following inequalities:

$$\int_\tau^\infty e^{-(\rho+\alpha)(t-\tau)} \frac{\theta}{\theta-1} E\left[\pi(1,a_t)\Big| a_\tau = a_p^*(1)\right] dt > \int_\tau^\infty e^{-(\rho+\alpha)(t-\tau)} \frac{\theta}{\theta-1} E\left[\pi(h_t,a_t) \mid h_\tau = 1, a_\tau = a_p^*(1)\right] dt = \psi \quad (19)$$

$$\int_\tau^\infty e^{-(\rho+\alpha)(t-\tau)} \frac{\theta}{\theta-1} E\left[\pi(0,a_t)\Big| a_\tau = a_p^*(0)\right] dt < \int_\tau^\infty e^{-(\rho+\alpha)(t-\tau)} \frac{\theta}{\theta-1} E\left[\pi(h_t,a_t) \mid h_\tau = 0, a_\tau = a_p^*(0)\right] dt = \psi, \quad (20)$$

where the last equality in both equations comes from the fact that equation (13) is satisfied with equality on the planner's threshold. Combining (19) and (20), rewriting $\pi(h,a) = e^a g(h)$ and rearranging we get

$$\int_\tau^\infty e^{-(\rho+\alpha)(t-\tau)} g(1) E\left[e^{a_t}\Big| a_\tau = a_p^*(1)\right] dt > \int_\tau^\infty e^{-(\rho+\alpha)(t-\tau)} g(0) E\left[e^{a_t}\Big| a_\tau = a_p^*(0)\right] dt.$$

Using the fact that for any initial condition a_0, a_t can be written as $a_t = a_0 + B_t$, where B_t is a Brownian motion with $B_0 = 0$,

$$g(1)e^{a_p^*(1)} \int_\tau^\infty e^{-(\rho+\alpha)(t-\tau)} E\left[e^{B_t}\right] dt > g(0)e^{a_p^*(0)} \int_\tau^\infty e^{-(\rho+\alpha)(t-\tau)} E\left[e^{B_t}\right] dt.$$

which implies that $g(1)e^{a_p^*(1)} = \pi(1, a_p^*(1)) > g(0)e^{a_p^*(0)} = \pi(0, a_p^*(0))$. Using the fact that $\frac{\partial Y(h,a)}{\partial h} = \frac{\theta}{\theta-1}\pi(h,a)$ concludes the proof.

Second statement

Doing some algebra, one can show that:

$$\frac{Y(1,a_p^*(1))/\pi(1,a_p^*(1))}{Y(0,a_p^*(0))/\pi(0,a_p^*(0))} = \left(\frac{x_H}{x_L}\right)^{\frac{\theta-1}{\theta}} > 1$$

and since $\pi(1, a_p^*(1)) > \pi(0, a_p^*(0))$, that yields the claim.

B Robustness

We ran our policy exercise using different sets of parameters. In order to get sufficiently different statistics from those in the baseline calibration, we did not try to match any business cycle features. In all of the following specifications, the economy experiences both crisis and recessions and minimal spending policies do 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 values of ξ were chosen to contemplate the cases of an almost vertical threshold and a threshold were the government barely pays subsidies when fundamentals are picking up and the economy is leaving from a situation with $h = 0$, as in Figure 9.

Table 2: Robustness check parameters

		Specification					
Parameter	**Symbol**	**1**	**2**	**3**	**4**	**5**	**6**
Production regime *High*	x_H	1.1	1.1	1.1	1.05	2	1.1
Production regime *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 with stimulus policies (%)	-	10.15	35.96	14.03	4.79	16.88	24.61

The results were the following:

1. **The planner's threshold**: The maximum difference between the slope (measured here as the distance in the horizontal axis of a threshold extreme points, i.e., $|a(1) - a(0)|$) of the planner's threshold and the agent's threshold was 4.4% of the slope of the agent's threshold. The minimum was below 0.1%. The mean of this difference was 2.14%. In some calibrations the planner's threshold had a higher slope and in others we obtained the opposite. Since it is very close to zero we cannot say precisely the direction of the change in the slope, since it could be only due to a measurement error. The fact is that the slope did not change very much.

2. **Minimal spending policies**: Figure 12 reports the results for government spending. The stimulus policies with $\xi = 0$ are the cheapest in all cases we tried.

Figure 12: Government spending in each specification