Dynamic Coordination and the Optimal Stimulus Policies*

Bernardo Guimaraes† Caio Machado‡

Forthcoming in the Economic Journal

Abstract

This paper studies stimulus policies in a simple macroeconomic model featuring a dynamic coordination problem that arises from demand externalities and fixed costs of investment. In times of low economic activity, firms face low demand and hence have lower incentives for investing, which reinforces their low-demand expectations. In a benchmark case with no shocks, the economy might get trapped in a low-output regime and a social planner would be particularly keen to incentivise investment at times of low economic activity. However, this result vanishes once shocks are considered.

JEL Classification: D84, E32, E62

Keywords: coordination failures, fiscal stimulus, timing frictions, demand externalities

Coordination is often said to play a role in recessions. This idea is captured by models with demand externalities that generate strategic complementarities in production.[1] In times of low economic activity, a firm faces low demand and thus has low incentives for

*Corresponding author: Bernardo Guimaraes. Rua Itapeva, 474, Sao Paulo-SP, Brazil. Email: bernardo.guimaraes@fgv.br.

We thank Marios Angeletos, Braz Camargo, Felipe Iachan, Nobuhiro Kiyotaki, V. Filipe Martins-da-Rocha, Stephen Morris, Ana Elisa Pereira, Mauro Rodrigues, Thijs van Rens and Mark Wright, the editor Morten Ravn, two anonymous referees, as well as seminar participants at U Carlos III, LACEA-LAMES 2014 (Sao Paulo), Spring 2015 Midwest Macro Meeting (St Louis), NASMES 2014 (Minneapolis), Princeton, PUC-Rio, Sao Paulo School of Economics – FGV, U Sao Paulo, SBE 2013 (Iguaçu), the XVII Workshop of International Economics and Finance (San Jose) and U Zurich. Guimaraes gratefully acknowledges financial support from CNPq. Machado gratefully acknowledges financial support from CAPES and the Sao Paulo Research Foundation (FAPESP), through grants 2012/23222-6 and 2013/22873-6. A previous version of the paper was entitled "Demand expectations and the timing of stimulus policies".

†Sao Paulo School of Economics - FGV

‡Instituto de Economía, Pontificia Universidad Católica de Chile

[1] Seminal papers in this literature are Shleifer (1986), Cooper and John (1988), Kiyotaki (1988) and Murphy *et al.* (1989).

investment. In a dynamic setting, this feedback effect may trap the economy in a low-output regime: lower investment today implies lower economic activity and lower investment tomorrow. What is the optimal stimulus policy for an economy subject to this dynamic problem? Is there a special reason for stimulus at times of low economic activity?

This paper develops a model to answer these questions. Investment is modeled as a fixed cost that increases production capacity. Investment decisions are staggered. Producers of each variety receive investment opportunities according to a Poisson clock, a simple way to capture production decisions that cannot adjust overnight. This assumption implies that investment decisions are not synchronised and economic activity is a state variable.

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, and producers have to form expectations about others' future decisions when deciding about investment.

Returns to investment depend not only on demand but also on productivity. If the increase in production resulting from investing is large enough, then investing is a dominant strategy. Likewise, if productivity is very low, investing is a dominated strategy. In an intermediate range, a producer's decision depends on his expectations about the actions of others: investing is the optimal decision if agents expect others to do so, but refraining from investment is the best choice in case of pessimistic beliefs.

We first consider a benchmark model with no shocks. In this case, there are multiple equilibria. In order to close the model, we need some assumption on beliefs. Assuming either 'pessimistic' or 'optimistic' beliefs, the solution to the planner's problem differs from the decentralised equilibrium in two ways: (i) the planner requires lower productivity to invest because it internalises the benefits to consumers from cheaper prices (monopoly distortion); and (ii) the difference between the planner's solution and the decentralised equilibrium is larger at times of low economic activity.

The second point highlights an inefficiency related to the dynamics of the economy. Agents might get stuck in a situation where economic activity is low, hence there is low demand and firms prefer not to invest even though productivity would be high enough to encourage investment if demand was high. In this situation, a firm would like others to invest, so that demand would increase and generate incentives for it to invest as well. However, nobody wants to be the first to invest and the economy is trapped in a situation with low economic activity. The planner would be particularly keen to stimulate investment in this situation.

We then consider the model with shocks to aggregate productivity. As in Frankel and Pauzner (2000), a unique equilibrium arises in the model. The model generates a unique set of rationalizable beliefs about others' actions. Intuitively, fully pessimistic beliefs are not rationalizable in a region where a small shock to productivity would make it dominant for all firms to invest. Likewise, fully optimistic beliefs are not rationalizable in a region where a small shock to productivity takes the economy to a region where investing is a dominated strategy. Agents know all others will reason like this and try to anticipate what others will do. This process yields a unique rationalizable set of strategies and beliefs.[2]

The main result of this paper is that differently from the case with multiple equilibria, there is no special reason for subsidies at times of low economic activity. The maximum amount of investment subsidies the planner is willing to provide at times of high and low economic activity is exactly the same.

The result holds even when the variance of productivity shocks is arbitrarily small. The only meaningful difference between arbitrarily small shocks and no shocks at all is that a unique set of beliefs is pinned down by the model in the former case. Hence the beliefs that arise in equilibrium exactly offset the dynamic inefficiency, so that there is no special reason for stimulus at times of low economic activity.

What are equilibrium beliefs like? Consider an agent indifferent between investing or not in a state of low economic activity. She understands that if fundamentals get a bit worse, firms will still be refraining from investing but there will be no major change in the state of the economy. Conversely, a slight improvement in fundamentals will trigger a recovery because firms will choose to invest and that will push the economy to a situation where investing is profitable for everyone. Owing to larger demand, firms will then have more incentives to invest, so it will take a large negative productivity shock to offset the benefit from increased demand and stop the recovery. The fundamental asymmetry is that bad news basically leaves the economy parked in a region of inaction, while good news drives the economy to a different state.

That does not mean that agents are usually more optimistic at times of low economic activity; it means that an agent that is indifferent between investing or not expects a positive change at times of low economic activity. Likewise, in a situation with high economic activity and relatively low productivity, agents understand that small shocks might trigger an investment slump.

In the benchmark case with multiple equilibria, the beliefs of agents that are indifferent between investing or not might be 'optimistic' (they expect others will invest) or 'pessimistic' (they expect others will not invest), depending on which equilibrium is considered. The key feature of the model with shocks is that beliefs in the indifference region are 'optimistic' when economic activity is low and productivity is relatively high and 'pessimistic' when

[2] Uniqueness of equilibrium does not preclude coordination failures – actually, the equilibrium is typically inefficient in static and dynamic coordination games with a unique equilibrium (examples include Carlsson and Van Damme (1993), Morris and Shin (1998) and most papers in the literature of global games). The equilibrium is also inefficient in our model.

economic activity is high and productivity is relatively low.

Related literature. This paper is related to the theoretical contributions in Frankel and Pauzner (2000) and Frankel and Burdzy (2005).[3] They show there is a unique rationalizable equilibrium in a class of dynamic models with time-varying fundamentals and timing frictions.[4] Our macroeconomic model fits in their framework so their results can be used to show equilibrium uniqueness in our model. However, neither of these papers solve the social planner's problem.

The demand externalities that play a key role in this paper are in the seminal contributions by Blanchard and Kiyotaki (1987), Kiyotaki (1988) and Murphy *et al.* (1989). When others produce more, the demand for a particular variety shifts to the right, and its producer finds it optimal to increase production. In Kiyotaki (1988), multiple equilibria arise because of increasing returns to scale. The model in this paper also gives rise to multiple equilibria in the absence of shocks to fundamentals, 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. Early examples include Cooper and John (1988), Benhabib and Farmer (1994) and Farmer and Guo (1994). Recent research on business cycles has explored the implications of equilibrium multiplicity in a variety of settings.

Benhabib *et al.* (2001) show that once the zero lower bound is considered, Taylor rules lead to multiple equilibria. Building on this insight, Evans and Honkapohja (2005) study the implications of replacing perfect foresight with learning in this environment; Mertens and Ravn (2014) show that government spending succeeds in stimulating output in case of a fundamental-driven liquidity trap, but fails to do so in case of a confidence-driven liquidity trap; and Aruoba *et al.* (2017) quantitatively assess the importance of sunspot shocks for the recent recessions in the US and in Japan. In Benigno and Fornaro (2015), the zero lower bound also play a key role and the model features two steady states owing to the interplay between productivity growth and aggregate demand.

The literature has shown other possible channels that generate multiplicity of equilibria, exploring a variety of feedback loops: in a search and matching model, Farmer (2012) re-

[3] See also Burdzy *et al.* (2001).

[4] 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 analyse 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) build a model of debt runs. Guimaraes and Pereira (2017) analyse how payoff heterogeneity affects these models. Guimaraes and Pereira (2016) study welfare in a model of network externalities.

places the assumption of Nash bargaining over the wage with the assumption that firms produce as many goods as are demanded;[5] Chamley (2014) presents a model with decentralised trade and credit constraints where pessimistic expectations lead to precautionary savings; in Benhabib *et al.* (2015), firms make production decisions based on expected demand and households choose consumption and labour supply based on their expected income; and in Kaplan and Menzio (2016), larger unemployment implies people spend more time searching for lower prices, which reduces firms' incentives to produce more. In Heathcote and Perri (2016), multiple equilibria arise only when the level of household wealth is low, because demand is more sensitive to unemployment expectations in this case, owing to a stronger precautionary motive. Bacchetta and van Wincoop (2016) argue that when economic integration is large enough, self-fulfilling panics cannot be limited to one country. Rendahl (2016) considers an economy at the zero lower bound with search and matching frictions and studies how government spending can put a halt to a downward spiral of self-reinforcing thrift.

Another branch of this literature considers coordination and strategic complementarities in business cycles models with a unique equilibrium. Angeletos and La'O (2009b) and Angeletos and La'O (2013) show in an environment with noisy and dispersed information how self-fulfilling fluctuations can emerge. Their model has a unique equilibrium, but features some key aspects of sunspot models. Angeletos *et al.* (2014) attempt to quantitatively assess the role of confidence in business cycles. 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 (2009a)). In Schaal and Taschereau-Dumouchel (2016), if firms expect low aggregate demand, they post fewer vacancies, produce less and aggregate demand is indeed low (as in Howitt and McAfee (1992)). Firms heterogeneity restores equilibrium uniqueness.

Some of this work uses the global games methodology to understand the effects of stimulus packages on coordination.[6] Sákovics and Steiner (2012) build a model to understand who matters in coordination problems: in a recession, who should benefit from government subsidies? They conclude that the government should subsidise sectors that have a large externality on others but that are not much affected by others' actions. Guimaraes *et al.* (2016) study how government spending affects coordination in a static model. Closer to our paper, Schaal and Taschereau-Dumouchel (2015) build a dynamic macroeconomic model with coordination failures and a unique equilibrium. Firms' choices of capacity utilization at every period is subject to coordination failures and is modelled as a global game.

[5] See also Howitt and McAfee (1987).

[6] The seminal contributions in the global games literature are Carlsson and Van Damme (1993) and Morris and Shin (1998). For a survey, see Morris and Shin (2003). Chamley (1999), another early contribution to this literature, studies regime switches in a dynamic model with strategic complementarities and learning.

Households' consumption-saving decisions affect coordination among firms and the dynamics of the economy. Although their environment is substantially different from ours, simple subsidies can also implement the first best in their paper.

Differently from models that employ the global games methodology to obtain equilibrium uniqueness in macroeconomic settings with strategic complementarities, our results do not rely on noisy heterogeneous information – all information is common knowledge here. The key ingredients to resolve indeterminacy in this model are timing frictions and shocks to fundamentals. Our framework is particularly suitable to understand the dynamic interplay between economic activity, productivity and beliefs that arise in equilibrium.[7]

The paper is organised as follows. Section 1 presents the model. Section 2 shows results for the benchmark case without shocks, while Section 3 considers the model with shocks and explain the beliefs that arise in equilibrium. Section 4 drops the assumptions of fixed cost of investing and random switching opportunities and lets firms choose a switching rate at every moment. Our main result also holds in this extension: there is no special reason for stimulus at times of low economic activity. Section 5 concludes. All proofs are in Appendix A.

1 Model

1.1 Environment

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

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

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

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

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

There is a measure-one continuum of agents who discount utility at rate ρ. An agent's

[7]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.

instantaneous utility at time t is given by $U_t = C_t$, where C_t is her instantaneous consumption of the composite good. The assumption of linear utility implies that policies will be concerned with inefficiencies in production but will not aim at providing insurance to the household.

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

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

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

The assumptions on technology aim at modelling staggered investment decisions in a simple and tractable way. There are 2 production regimes, a *High*-capacity regime and a *Low*-capacity regime. Agents get a chance to switch regimes according to a Poisson process with arrival rate α.[8] Once an individual is picked up, she chooses a regime and will be locked in this regime until she is selected again. Choosing the *Low* regime is costless. Choosing the *High* regime costs ψ units of the composite good.[9]

An agent in the *Low* regime can produce up to 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.

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

Investment requires agents to acquire a stock ψ of composite goods, which cannot be funded by their instantaneous income, so we assume agents can trade assets and borrow to

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

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

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

invest. Owing to the assumption of linear utility, any asset with present value equal to ψ is worth ψ in equilibrium. For example, an agent might issue an asset that pays $(\rho+\alpha)\psi dt$ at every interval dt until the investment depreciates ($\rho\psi dt$ would be the interest payment and $\alpha\psi dt$ can be seen as an amortization payment since debt is reduced from ψ to 0 with probability αdt). Since agents are risk neutral, other types of assets would deliver the same results.

Let $a_t = \log(A_t)$ vary in time according to[11]

$$da_t = \sigma dZ_t,$$

where $\sigma > 0$ and Z_t is a standard Brownian motion. In order to ensure that agents face a coordination problem, we assume that $\sigma^2 < 2(\rho+\alpha)$.[12]

1.2 The Agent's Problem

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

$$p_{it} = y_{it}^{-1/\theta} Y_t^{1/\theta}, \tag{3}$$

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

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

and

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

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.

Let $\pi(a_t, h_t)$ be the difference between instantaneous income of agents locked in the *High* regime (w_{Ht}) and income of agents locked in the *Low* regime (w_{Lt}) when the economy is at

[11]The main results also hold if a_t has a constant drift and, under some additional technical assumptions, if a_t follows a process with mean reversion. See the Online Appendix.

[12]Intuitively, an agent's expected gain from investing is infinite if a_t is too volatile. The expected growth rate of the utility flow from investing in this model is $\sigma^2/2$ (Ito's Lemma). This has to be smaller than agents' effective discount rate (which is $\rho+\alpha$).

(a_t, h_t). Combining (4), (5), (6) and using $y_{Lt} = e^{a_t} x_L$ and $y_{Ht} = e^{a_t} x_H$ leads to

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

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 investment decisions are strategic complements: 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 $s(a_t, h_t) \mapsto \{Low, High\}$.[13] An agent at time $t = \tau$ that has to decide whether to invest will do so if

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

and will not invest if the inequality is reversed. In words, investing pays off if the discounted expected additional profits from choosing the *High* regime are larger than the fixed cost ψ. Future profits $\pi(a_t, h_t)$ are discounted by the sum of the discount rate and depreciation rate $(\rho + \alpha)$.

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.

1.3 The Planner's Problem

The planner maximises expected welfare, given by

$$E_\tau(W) = E_\tau \int_\tau^\infty e^{-\rho(t-\tau)} \left[Y(a_t, h_t) - \alpha\psi I(t) \right] dt, \tag{9}$$

where $Y(a, h)$ is given by (6) and $I(t) \in [0, 1]$ is the decision of the planner about investing at time t (the proportion of those who got an investment opportunity at t that will invest).

[13] To simplify the exposition, we present the definition of a Markovian strategy, but our results do not rely on that restriction on the strategy space.

Suppose the optimal investment decision at date τ implies $I(\tau) < 1$ and consider the following deviation: the planner chooses $I(\tau) = 1$ today and keeps investment choices for every realization of the Brownian path in the future unchanged.[14] This deviation cannot be profitable. Investing extra dI units today raises h_τ by αdI, but this depreciates at rate α. Hence

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

Therefore, this deviation is not profitable if

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

Since

$$\frac{\partial Y(h_t, a_t)}{\partial h} = e^{a_t} \frac{\theta}{\theta - 1} \left[h_t x_H^{\frac{\theta-1}{\theta}} + (1 - h_t) x_L^{\frac{\theta-1}{\theta}} \right]^{\frac{1}{\theta-1}} \left(x_H^{\frac{\theta-1}{\theta}} - x_L^{\frac{\theta-1}{\theta}} \right), \tag{10}$$

if the planner chooses not to invest $(I(\tau) = 0)$, it must be that

$$\int_\tau^\infty e^{-(\rho+\alpha)(t-\tau)} E_\tau \left[\frac{\theta}{\theta - 1} \pi(h_t, a_t) \right] dt \leq \psi. \tag{11}$$

An analogous reasoning shows that if at date τ the planner chooses to invest $(I(\tau) = 1)$, it must be that:

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

The expressions in (11) and (12) are necessary conditions for optimality. Note they are analogous to the corresponding necessary conditions for a Nash Equilibrium in the agents' game (which are also sufficient conditions in that case). The only difference between (8) and (12) is the constant multiplying $\pi(a, h)$ in (12). Thus, finding candidates for the planner's solution is equivalent to finding the equilibrium set of a modified game. It will be shown that in some cases, these necessary conditions are also sufficient.[15]

2 Benchmark Case: No Shocks

We first consider the case where the fundamental a does not vary over time, $\sigma = 0$.

[14]Since there is no interaction with other players and no intrinsic time-inconsistency in the planner's preferences, there is no commitment issue in the planner's problem.

[15]Agents takes others' strategies as given and anticipate other agents will be choosing according to (8). The planner can decide on the path of h, but it also anticipates its future selves will satisfy (12).

2.1 Equilibria

Consider an agent with optimistic beliefs, i.e., she expects all others will choose to invest in the future. An agent is indifferent between investing or not at $(a^*_{opt}(h_0), h_0)$ if $a^*_{opt}(h_0)$ solves

$$\int_0^\infty e^{-(\rho+\alpha)t}[\pi(a^*_{opt}(h_0), h_t^U)]dt = \psi, \tag{13}$$

where $h_t^U = 1 - (1 - h_0)e^{-\alpha t}$. If the agent holds pessimistic beliefs (if she expects all others will choose not to invest in the future), she is indifferent between investing or not at $(a^*_{pes}(h_0), h_0)$ if $a^*_{pes}(h_0)$ solves

$$\int_0^\infty e^{-(\rho+\alpha)t}[\pi(a^*_{pes}(h_0), h_t^D)]dt = \psi, \tag{14}$$

where $h_t^D = h_0 e^{-\alpha t}$.

Proposition 1 characterises the conditions for multiple equilibria.

PROPOSITION 1 (No Shocks). *Suppose $\sigma = 0$. There are strictly decreasing functions $a^*_{opt} : [0,1] \mapsto \Re$ and $a^*_{pes} : [0,1] \mapsto \Re$ with $a^*_{opt}(h) < a^*_{pes}(h)$ for all $h \in [0,1]$ such that*

1. *If $a < a^*_{opt}(h_0)$ there is a unique equilibrium, agents always choose the Low regime;*
2. *If $a > a^*_{pes}(h_0)$ there is a unique equilibrium, agents always choose the High regime;*
3. *If $a^*_{opt}(h_0) < a < a^*_{pes}(h_0)$ there are multiple equilibria, that is, both strategies High and Low can be long-run outcomes.*

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 regime where $h = 1$ (and there it will rest). If the productivity differential is sufficiently low, the gains from investing are offset by the fixed cost, so not investing is a dominant strategy. In an intermediate area, there are no dominant strategies, the optimal investment decision depends on expectations about what others will do and there are multiple equilibria.

Figure 1: Equilibria without Shocks

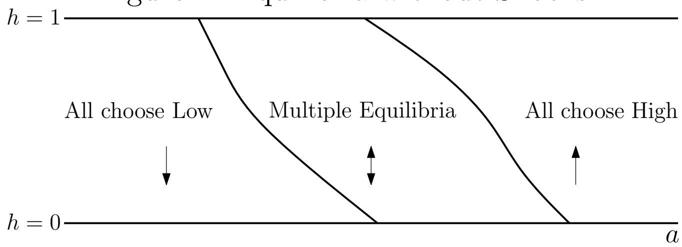

We say the economy is in the 'good equilibrium' when agents choose *High* and expect others to do so whenever that is rationalizable. Conversely, we say the economy is in the 'bad equilibrium' when agents choose *Low* and expect others to do so whenever that is rationalizable.

The left threshold (good equilibrium) is the set of (a, h) where an agent is indifferent between investing or not assuming all others will invest, and it is given by (13). For the right threshold (bad equilibrium), the assumption is that no other agent will ever choose to invest, as in (14). The good equilibrium Pareto dominates the bad equilibrium.

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 characterise the economy and its current state.

2.2 Optimal Policy

One important question is about whether inefficiencies are more pronounced at times of low economic activity. For instance, suppose the economy is stuck in a regime with low h. Is this situation particularly inefficient? Would a social planner be particularly inclined to stimulate investment in this case?

In the case with $\sigma = 0$, the planner chooses between always investing and never investing (any other option is dominated by one of these alternatives). Thus the planner is indifferent between investing and not investing at $(a_P^*(h_0), h_0)$ when $a_P^*(h_0)$ is given by

$$\int_0^\infty e^{-\rho t} \left[Y(a_P^*(h_0), h_t^U) - \alpha\psi \right] dt = \int_0^\infty e^{-\rho t} \left[Y(a_P^*(h_0), h_t^D) \right] dt. \tag{15}$$

This expression pins down a decreasing threshold a_P^* such that the planner will invest when $a > a_P^*(h_0)$ and not invest when $a < a_P^*(h_0)$.

In order to compare the social planners' solution with the decentralised economy and characterise optimal policies, we need to select an equilibrium in the model. One possible way is to assume agents coordinate in the good equilibrium whenever possible.[16] An alternative way is to assume the worst equilibrium is played whenever it exists and search for policies that eliminate this equilibrium.[17] Proposition 2 shows both ways lead to similar conclusions.

PROPOSITION 2. *Suppose a is fixed and the economy is either in the 'good' or in the 'bad' equilibrium. In both cases:*

1. *[Planner's choices depend less on h] The distance between the planner's threshold at $h = 0$ and at $h = 1$ is smaller than the equivalent distance in both decentralised*

[16] This assumption is common in the literature. Examples include Allen and Gale (1998), Zawadowski (2013) and Boissay *et al.* (2016).

[17] Policy analysis in models with multiple equilibria often take this route. Examples include Chang and Velasco (2000), Aghion *et al.* (2004), Benhabib *et al.* (2002) and some of the work building on the latter (e.g., Schmitt-Grohe and Uribe (2017) and Benigno and Fornaro (2015)).

equilibria:

$$a_P^*(0) - a_P^*(1) < a_{opt}^*(0) - a_{opt}^*(1) \quad and \quad a_P^*(0) - a_P^*(1) < a_{pes}^*(0) - a_{pes}^*(1).$$

2. *[More subsidies at low h] The planner's solution can be implemented by investment subsidies. At the planner's threshold, the amount of subsidies required to coax agents to invest is higher when $h = 0$ than when $h = 1$.*

The equilibria of the model and the planner's threshold are illustrated in Figure 2.

Figure 2: The Case with No Shocks

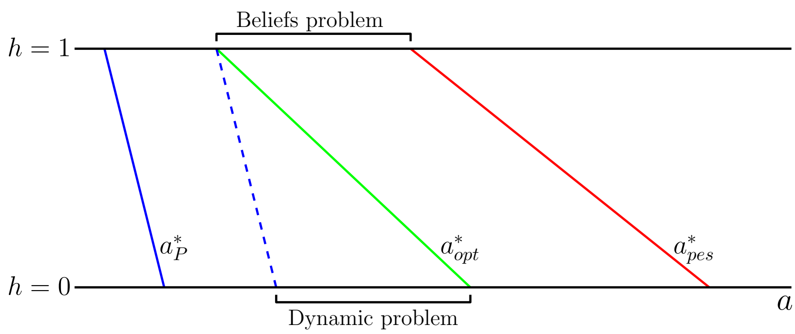

There are important differences between the planner's solution and the decentralised equilibrium. First, the solution for the planner's problem is unique but there are multiple self-fulfilling equilibria. More importantly, regardless of whether we assume optimistic or pessimistic beliefs, the equilibrium threshold is further away from the planner's threshold for low values of h. The difference in slopes reflects the planner's willingness to pay higher subsidies when h is low.

At the heart of this problem lies an inefficiency related to the dynamics of the economy. Agents only take into account the benefit of investment until their next decision, which means they effectively discount the future at rate $\alpha + \rho$. In contrast, the planner takes into account the effect that investing has on the whole path of the economy, and thus discounts the future at the much smaller rate ρ. In a region with low h and a just below the threshold, investment in the short run would drive the economy to the high regime and the planner takes that into account. Agents don't find it profitable to invest while demand is still low, but would be happy to sign a contract forcing everyone to take investment opportunities in the short run, as these losses for some would imply gains for all in the future.

In the absence of investment subsidies, the economy might be stuck in a recession trap (low h, not so low a). The root of the problem is a self-reinforcing lack of demand (and, consequently, investment) when h is low: no investment in the past leads to low demand and no investment today, which in turn leads to low demand and no investment in the future and so on.

The recession trap does not stem from the multiplicity of rationalizable beliefs, since it holds under the assumption that agents coordinate on the good equilibrium whenever possible. This result thus suggests that in a world with demand externalities and an inaction region for investment, subsidies to take the economy out of a recession are warranted. However, beliefs are exogenous in this reasoning. We now consider the model with $\sigma > 0$.

3 The Model with Shocks

We now turn to the general case where productivity varies over time. 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.

3.1 Equilibrium

The model satisfies the assumptions of the framework in Frankel and Pauzner (2000). Hence we can apply their results to show there is a unique rationalizable equilibrium where agents play according to a decreasing threshold $a^*(h)$.

PROPOSITION 3 (Frankel and Pauzner, 2000)**.** *Suppose $\sigma > 0$. There is a unique rationalizable equilibrium in the model. Agents invest if $a > a^*(h)$ and do not invest if $a < a^*(h)$, where a^* is a decreasing function.*

In order to understand how shocks affect the set of rationalizable strategies, consider a situation where productivity is relatively low, so a firm is only willing to invest if the probability the following firms will also invest is very high. In Figure 1, that would correspond to a point in the multiplicity region but close to its left boundary. In a world with shocks, the economy might cross to the region where investing is a dominated strategy. That imposes a cap on the probability that others will invest in the near future – the belief that they will certainly invest is not rationalizable. In consequence, some dominated strategies are eliminated, which imposes further limits on beliefs agents can hold. Iterating on this process leads to a unique equilibrium.[18]

This reasoning highlights a fundamental difference between the benchmark model with no shocks and this one. The 'good equilibrium' and the 'bad equilibrium' of the model in

[18]Since a follows a Brownian motion, shocks to a in a small period of time are potentially unbounded, but this is not important for the results. Uniqueness stems from the iterative elimination of dominated strategies, not from unlikely large shocks.

Section 2 are derived under the assumption that agents know what others will do (different beliefs would be rationalizable). Those equilibria would not survive the inclusion of some uncertainty about others' future actions. Here, in contrast, uncertainty about the path of a opens the door to uncertainty about the actions of others. This is arguably an important component of an economy prone to dynamic coordination failures and plays a key role in determining the equilibrium. The iterative process that leads to the elimination of a large set of strategies can be interpreted as agents trying to forecast the forecast of others.[19]

The equilibrium is characterised by a threshold. A larger h implies that agents are willing to invest for lower values of a, as in Figure 3. Beliefs about others' investment decisions are pinned down by fundamentals (a) and history (h). Shocks to a_t and movements in h_t might affect expectations about others' actions.

Figure 3: Equilibrium with Shocks

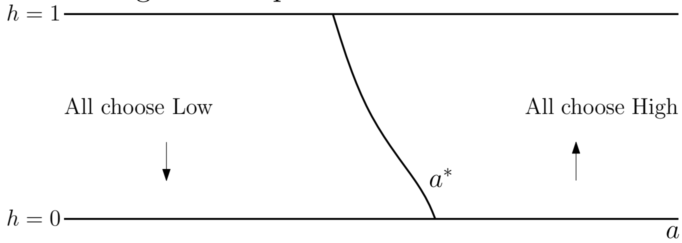

Let $V(a, h, \tilde{a})$ be the utility gain from choosing *High* obtained by an agent in state (a, h) that believes others will play according to threshold $\tilde{a}$. Then

$$V(a, h, \tilde{a}) = \int_0^\infty e^{-(\rho+\alpha)t} E[\pi(a_t, h_t)|a, h, \tilde{a}]dt - \psi, \tag{16}$$

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

3.2 Optimal Policy

Proposition 3 shows there is a unique rationalizable equilibrium in the model. Although agents face a dynamic coordination problem, a unique set of rationalizable beliefs emerges and, from the point of view of an individual firm, pins down the optimal decision. It is then natural to ask about the beliefs that arise in equilibrium and, in particular, about the inefficiencies that might exist in the model.

The key implications for the optimal stimulus policies are in Proposition 4.

[19]For more on higher order beliefs in dynamic coordination games with timing frictions, see Morris (2014).

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] 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 decentralised equilibrium.

Figure 4: Planner's Problem

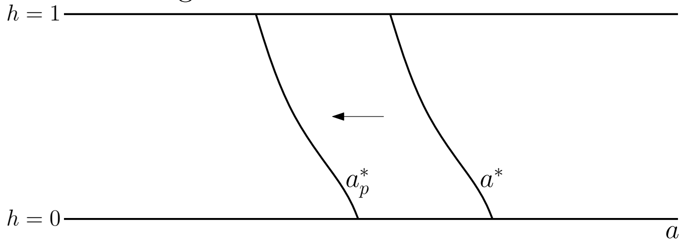

In principle, it is difficult to characterise the planner's solution because expectations about the path of (a, h) have to be taken into account when solving for the optimal decision, but the path of h will be optimally chosen by the planner. However, mathematically, the planner's problem is similar to the agent's problem in the decentralised equilibrium. At every point in time, there is investment if (12) holds, taking into account that the path of h in the future will be determined by a similar choice. The planner thus chooses according to a threshold a_P^* such that (12) holds with equality at $a_P^*(h)$ for $h \in [0, 1]$. The only difference is that the planner and agents follow different decision rules.

In the decentralised equilibrium, agents choose according to (8). Optimal policy boils down to making agents decide according to the expression in (12). The only difference between these expressions is the term $\theta/(\theta - 1)$ multiplying the benefit from investing in (12). Hence, the extra incentive for investment the planner would like to provide is a constant proportion of the flow payoff. Alternatively, the planner would like to reduce the investment cost, multiplying it by $(\theta - 1)/\theta$.

Since the investment cost is fixed, a constant investment subsidy equal to ψ/θ implements the planner's solution. As an alternative, the planner could top up firms' revenues, paying $\pi(h_t, a_t)/(\theta - 1)$ for firms in the *High* regime at every time t. The bottom line is that there is no special reason for incentivizing investment at times of low economic activity (low h) – regardless of how investments are incentivised. These policies disregard any costs imposed by taxation, required to fund subsidies, but in Section 3.5 we assume that every unit of subsidy has a small welfare cost and obtain similar results.[20]

From a social point of view, the problem in the decentralised equilibrium is that investment requires an excessively high benefit. The key result in Proposition 4 is that this problem is not more severe when h is low (or high). Hence there is no special reason for incentivizing investment at times of low (or high) economic activity.

Proposition 4 also shows that the planner's threshold is a translation of the equilibrium threshold, as in Figure 4. The slope of the threshold affects the likelihood of a recession and its expected duration.[21] Hence Proposition 4 also shows that the planner has no reason to affect the expected duration of recessions.

The planner's solution prescribes no extra stimulus for investment when h is low because equilibrium beliefs offset the dynamic inefficiency highlighted in Section 2. In order to understand this point, it is instructive to look at the case with very small shocks. In the case with no shocks, there are multiple equilibria, so beliefs outside the dominance regions are not determined by the model. In the case with arbitrarily small shocks, productivity a behaves in a very similar way, but a unique set of beliefs is pinned down by the model. This comparison allows us to understand the beliefs that arise in equilibrium and how they affect policy.

3.3 The Case with Very Small Shocks

The uniqueness result and the expressions for the equilibrium and planner's thresholds hold for any $\sigma > 0$. However, the expressions for the thresholds depend on beliefs about the path of (a, h). In general, these are complicated objects, but in case $\sigma \to 0_+$, these beliefs can be determined.

For any $h_0 \in [0, 1]$, suppose the economy is at the threshold, i.e., at $(a^*(h_0), h_0)$. Where will the economy go? This mathematical problem is studied by Burdzy *et al.* (1998) and their main result is that when $\sigma \to 0_+$, the economy will instantaneously move up in the direction of $(a^*(h_0), 1)$ with probability $1 - h_0$ and will move down in the direction of $(a^*(h_0), 0)$ with probability h_0. This result determines agents' beliefs at the equilibrium threshold.

In order to understand this result, suppose $h_0 = 0.1$. If the economy is just at the right of the threshold, 90% of the agents that get an opportunity to switch will change from the

[20]The result is also robust to the inclusion of some monitoring costs. If the planner faces a cost c to monitor the investment it subsidises, with $c < \psi/\theta$, an argument similar to the one in Proposition 4 shows that the optimal policy can be implemented by a constant subsidy and leads to a different translation of the threshold, but no rotation.

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

Low to the *High* regime (and the remaining 10% will stay in the *High* regime); while if the economy is just at the left of the threshold, 10% of the agents that get an opportunity to switch will change from the *High* to the *Low* regime (and the remaining 90% will stay in the *Low* regime). Hence, at the right of the threshold, the economy moves up with speed proportional to $90\% = 1 - h_0$, and at the left of the threshold the economy moves down with speed proportional to $10\% = h_0$. Burdzy *et al.* (1998) show that the probabilities the economy move up or down are proportional to the speeds at each side of $(a^*(h_0), h_0)$. Intuitively, in a very short period of time, shocks to a will make the economy move around the threshold. Since the threshold is negatively sloped, an economy that moves up very quickly when it is at the right of the threshold is likely to find itself sufficiently above the threshold very soon, so that negative shocks to a cannot bring it back to the left side of the threshold.

Using the beliefs implied by the result in Burdzy *et al.* (1998) and the equilibrium condition in (16), we get that for any $h_0 \in [0, 1]$, $a^*(h_0)$ solves

$$(1 - h_0) \int_0^\infty e^{-(\rho+\alpha)t} [\pi(a^*(h_0), h_t^U)] dt + h_0 \int_0^\infty e^{-(\rho+\alpha)t} [\pi(a^*(h_0), h_t^D)] dt = \psi. \tag{17}$$

The solution to the planner's problem is similar. Using (12), we get that for any $h_0 \in [0, 1]$, the planner's threshold $a_P^*(h_0)$ solves

$$(1 - h_0) \int_0^\infty e^{-(\rho+\alpha)t} [\pi(a_P^*(h_0), h_t^U)] dt + h_0 \int_0^\infty e^{-(\rho+\alpha)t} [\pi(a_P^*(h_0), h_t^D)] dt = \psi - \frac{\psi}{\theta}. \tag{18}$$

This expression seems very different from (15) but yields the same results. For the planner, there is no difference between the cases with no shocks ($\sigma = 0$) or vanishing shocks ($\sigma \to 0_+$).[22] Beliefs about the future are basically the same and the planner can effectively choose the path of h.

Figure 5: The Case with Very Small Shocks

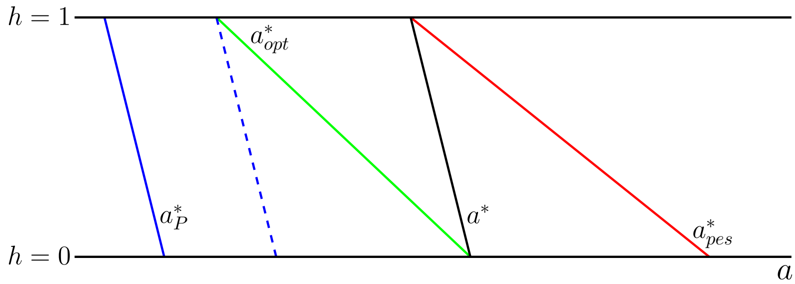

[22]The irrelevance of vanishing shocks for the planner's solution highlights the point that very small fluctuations in a are not intrinsically important. Their effects on the decentralised equilibrium stem from the determination of beliefs in case $\sigma \to 0_+$.

Figure 5 summarises the main results of this paper. The planner chooses to invest if the economy is at the right of a^*_P. In case $\sigma \to 0_+$, agents choose according to the threshold a^*, which solves (17). In case $\sigma = 0$, there are multiple equilibria. In the 'good equilibrium', agents believe all others will choose to invest. This is an equilibrium as long as the economy is at the right of a^*_{opt}. In the 'bad equilibrium', agents believe all others will not invest. This is an equilibrium as long as the economy is at the left of a^*_{pes}.

The result in case $\sigma \to 0_+$ is thus completely different from the case with $\sigma = 0$. In case $\sigma \to 0_+$, the slopes of the planner's and the agents' thresholds are the same. The slopes of both a^*_{opt} and a^*_{pes} are very different.

Since the productivity parameter moves very slowly when $\sigma \to 0_+$, the difference between the cases $\sigma \to 0_+$ and $\sigma = 0$ must stem from the difference in beliefs around the thresholds. As shown in Figure 5, $a^*(0)$ and $a^*_{opt}(0)$ coincide. Hence, beliefs at $(a^*(0), 0)$ and $(a^*_{opt}(0), 0)$ must be the same. Indeed, in the model with shocks, when the economy is at $(a^*(0), 0)$, agents believe others will choose to invest (when given an opportunity). Likewise, $a^*(1)$ and $a^*_{pes}(1)$ also coincide, as agents believe others will not invest when the economy is at $(a^*(1), 1)$ in case $\sigma \to 0_+$.[23]

Intuitively, in a neighborhood of the equilibrium threshold a^*, optimistic beliefs (i.e., beliefs that agents with a switching opportunity will invest) make perfect sense at $h = 0$, but no sense whatsoever at $h = 1$. Suppose the economy is at $(a^*(0), 0)$. The economy will stay around there as long as $a < a^*(0)$, but any shock that moves a above $a^*(0)$ leads agents to invest and drives the economy up in Figure 5. Since the slope of the threshold is negative, as soon as the economy is at $h > 0$, it is at the right of the threshold and hence agents have more incentives to invest. Thus the economy moves up in the direction of $(a^*(0), 1)$. The fundamental asymmetry here is that a small negative shock basically leaves the economy where it is, while a small positive shock drives the economy up in the direction of $h = 1$. The same reasoning implies that in a neighborhood of the equilibrium threshold at $(a^*(1), 1)$, a regime switch is also expected (beliefs are pessimistic).

Agents are indifferent between investing or not at $(a^*(0), 0)$ and at $(a^*(1), 1)$, but for different reasons. Around $(a^*(0), 0)$, productivity is relatively high but the economy is parked in a region of no investment; around $(a^*(1), 1)$, productivity is relatively low but everyone is producing at full capacity. The key feature of the model with shocks is that beliefs at the equilibrium threshold are 'optimistic' when economic activity is low and productivity is relatively high, and 'pessimistic' when economic activity is high and productivity is relatively low. This point does not apply to the model with multiple equilibria and no shocks, where beliefs of agents that are indifferent between investing or not might be 'optimistic' or

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

'pessimistic' for any value of h, depending on which equilibrium is considered. When shocks are considered, a unique set of beliefs arises in equilibrium and exactly offsets the dynamic inefficiency, so that there is no special reason for stimulus at times of low economic activity.

The explanation so far has considered the case $\sigma \to 0_+$ but Proposition 4 shows the result holds for any $\sigma > 0$. The intuition for the case with σ bounded away from zero is very similar. When economic activity (h) is low and agents are around the equilibrium threshold, it is likely that others will soon start to invest. Again, the key asymmetry here is that a movement of a to the left does not significantly affect the state of the economy, but a movement of a to the right affects the mass of agents investing, raising demand in the economy and incentives for the following agents with investment opportunities to take them. The recovery is just waiting for a small piece of good news.[24]

That does not mean that agents are more optimistic in recessions, in general. The empirical counterpart of the theoretical implication about beliefs is that 'confidence' (as measured by surveys) is a leading indicator, as agents anticipate the economy is about to leave the inaction region (or about to enter an investment slump).[25]

The binary set of actions is a tractable way to capture non-convexities in firms production. The fundamental assumption here is that these non-convexities generate an inaction region.[26] Different (non-convex) technologies should also imply that good news will have large effects on a firm's production only if the current level of production is low.

The cyclical behavior of beliefs offsets the dynamic inefficiency shown in the case with no shocks (Figure 2). For $\sigma > 0$, the distance between the planner's and the agents' threshold is independent of h, so subsidies do not depend on h as well. The only difference between the planner's solution and the decentralised equilibrium is the monopoly distortion in the investment decision. When beliefs are uniquely determined by the model, planner and agents solve a very similar problem. At every (a, h), investment is undertaken if its expected return pays off and the equilibrium (or planner's) threshold is a fixed point. "Pays off" means different things for agents and planner but the ratio is constant since the only difference is the externality from market power.

3.4 The Case with Vanishing Frictions

In case of vanishing frictions ($\alpha \to \infty$), the economy moves very fast from $h = 0$ to $h = 1$ but agents' horizons also become very short.[27] In the model with no shocks, agents take

[24]Graphically, for σ bounded away from 0, the planner's and agents' equilibrium thresholds would be parallel to each other as in Figure 5, but a^* would not touch a^*_{opt} and a^*_{pes} at any point.

[25]Confidence variables appear in many leading indicators indexes, such as those in the OECD System of Composite Leading Indicators.

[26]For a model where this inaction region arises as a result from fixed adjustment costs and may give rise to coordination failures, see Guimaraes *et al.* (2016).

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

h into account in their decisions, regardless of whether we assume fully optimistic or fully pessimistic beliefs. In contrast, the planner does not take h into consideration. Moving the economy to a different regime takes very little time and hence the transition is unimportant. The threshold from the planner's problem converges to a vertical line, as in Figure 6. This result holds for any $\sigma > 0$.

Figure 6: The Case with Vanishing Frictions

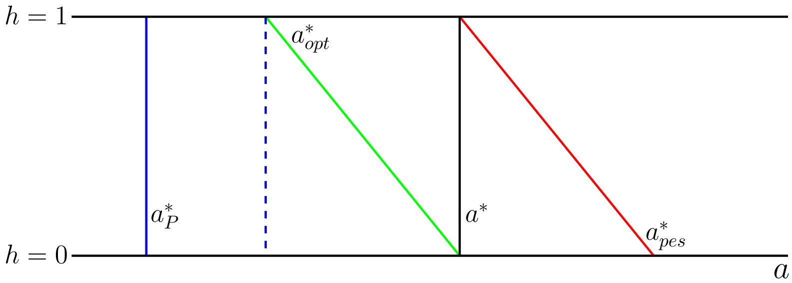

An implication of Proposition 4 is that in equilibrium agents also play according to a vertical threshold. History thus becomes irrelevant. For a large α, an agent at the equilibrium threshold and $h = 0$ knows the economy will move up with probability 1, while an agent at the equilibrium threshold and $h = 1$ is sure the economy will move down. They don't know their 'position in the queue', i.e., when they will have the next opportunity for revising their behavior (and thus how much time will have elapsed and the value of h when they can choose again).

On the one hand, the agent at $h = 0$ will experience a lower range of values of h than the agent at $h = 1$ (before their next opportunity to choose again): the one starting at $h = 0$ will experience h from 0 to $h^\dagger$, where $h^\dagger$ is uniformly distributed in $[0, 1]$; the agent starting at $h = 1$ will experience h from 1 to $1 - h^\dagger$. On the other hand, for the agent at $h = 0$, the economy moves up very quickly at lower values of h, but slowly as h approaches 1, so the last firms to change their decision will spend relatively more time at high values of h. As it turns out, both effects exactly cancel each other, so the agents at $h = 0$ and $h = 1$ are indifferent between investing and not investing for the same value of a.

This result might sound obvious owing to a simple but incorrect intuition: the economy can move very quickly to the $h = 1$ regime, so h does not affect the result. However, this intuition would also apply to the model with multiple equilibria, and the results in that case are completely different, h does matter as $\alpha \to \infty$. That is because although the economy can move very quickly from a low h to a high h, the next opportunity to invest also comes very quickly, so only the very short run matters for an agent's decision.

3.5 Implementation

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: what if every unit of subsidy has a small welfare cost $\varepsilon \approx 0$? It is not difficult to show that the government will use minimal spending policies, as in Definition 1.

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

$$\varphi(a,h) = \begin{cases} \psi - \int_0^\infty e^{-(\rho+\alpha)t} E[\pi(a_t,h_t)|a,h,a_p^*]dt & \textit{if } a_p^*(h) \leq a \leq \hat{a}(h), \\ 0 & \textit{otherwise.} \end{cases} \tag{19}$$

Figure 7 shows three 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 $h \in [0,1]$.

Figure 7: Example of Minimal Spending Policy

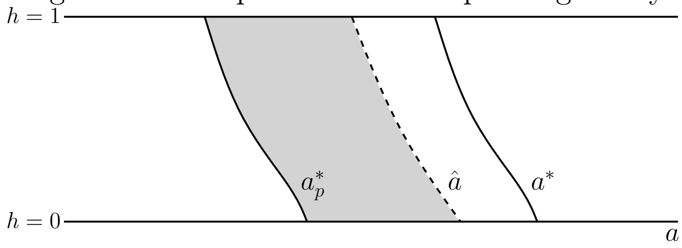

A government following a minimal spending policy is committed to give an investment subsidy to each agent in the region between a_p^* and $\hat{a}$ (the gray area in figure 7). The subsidy $\varphi(a,h)$ makes her indifferent between choosing *High* and *Low* given others will play according to a_p^*. 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^*.[28]

[28] The equilibrium under the minimal spending policy is no longer unique. If agents believe others will play according to a^* their best response is to play according to a^* as well, thus the policy has no effect at all. The amount of subsidies required to coax agents to invest depends on whether they expect others to respond to the stimulus policy. However, the government could implement this allocation through a contingent subsidy

Proposition 5 shows that minimum spending policies do not affect the main result of the paper.

PROPOSITION 5. *In the model of Section 1 with minimal spending policies, the maximum optimal subsidy is ψ/θ for all $h \in [0,1]$.*

The result is intuitive. Under minimal spending policies, investment subsidies are equal to the minimum between how much the planner is willing to pay and how much is required to coax agents to invest. The result from Proposition 5 thus follows from the planner's willingness to subsidise being independent of h (Proposition 4).

4 The Model with Endogenous Hazard Rates

The model so far has considered an exogenous rate α for the arrival of switching opportunities. We now modify the model in order to endogenise the switching rate.

As before, a firm produces $A_t x_H$ in the *High* regime and $A_t x_L$ in the *Low* regime. At each moment, a firm in the *Low* regime chooses the hazard rate of switching opportunities $\alpha_t^L \in [0, \overline{\alpha}]$, with $\overline{\alpha} < \infty$, subject to a cost $c(\alpha_t^L)$. The cost function is increasing, continuous and convex. This assumption replaces the fixed cost of investing in the previous model. A firm in the *High* regime switches to the *Low* regime with an exogenous hazard rate α^H.

In one interpretation, $c(\alpha_t^L)$ is the amount of resources a firm spends in the search for new ideas of production methods, and a useful idea appears at rate α_t^L. In another interpretation, equipments break or become obsolete at rate α^H and are immediately replaced, but the time until a new equipment can be used in production is stochastic and depends on the amount of resources allocated to this end, $c(\alpha_t^L)$. Alternatively, as in Howitt and McAfee (1992), a firm in state *Low* searches for a match and the cost $c(\alpha_t^L)$ is increasing in the search intensity. Owing to attrition, matches are broken at rate α^H.

To ease notation, define:

$$g(h_t) = \left[h_t x_H^{\frac{\theta-1}{\theta}} + (1-h_t) x_L^{\frac{\theta-1}{\theta}}\right]^{\frac{1}{\theta-1}} \left(x_H^{\frac{\theta-1}{\theta}} - x_L^{\frac{\theta-1}{\theta}}\right).$$

In order to find the decentralised equilibrium, we can apply the results in Frankel and Burdzy (2005). Their Theorems 4 and 5 imply that the relative value of being in the *High* regime is given by:

$$\Delta V_\tau = E_\tau \int_\tau^\infty e^{-\int_\tau^s \left(\rho + \alpha_v^L + \alpha^H\right) dv} \left[e^{a_s} g(h_s) + c(\alpha_s^L)\right] ds,$$

that would be essentially equivalent to a minimal spending policy: a large subsidy contingent on others not investing, and the subsidy prescribed by the minimal spending policy in case others invest as well.

agents choose α_τ^L in order to maximise

$$\alpha_\tau^L \left(E_\tau \int_\tau^\infty e^{-\int_\tau^s \left(\rho+\alpha_v^L+\alpha^H\right)dv} \left[e^{a_s} g(h_s) + c(\alpha_s^L)\right] ds \right) - c(\alpha_\tau^L), \tag{20}$$

and this problem yields essentially a unique equilibrium.

The expression in (20) shows that firms' optimal choice of hazard rates depends on the expected gains from switching. The relative value of being in the *High* regime is increasing in expected values of a and h, which depend positively on their current values. Hence, larger a and larger h will induce higher hazard rates.

We now proceed to solve the planner's problem. Since the cost function is convex, the planner chooses the same hazard rate for every firm in regime *Low*. Welfare in this economy is thus given by

$$E_\tau(W) = E_\tau \int_\tau^\infty e^{-\rho(t-\tau)} \left[Y(a_t, h_t) - (1-h_t)c(\alpha_t^L)\right] dt,$$

where $Y(a, h)$ is given by (6). Suppose the planner is following the optimal plan and consider the following deviation: change α_t^L to $\tilde{\alpha}$ at τ for an infinitesimal period dt and keep future choices for every realization of the Brownian path in the future unchanged. This affects current costs and output net of switching costs for all $s > \tau$. Since there are $1 - h_t$ agents at the *Low* state, costs change by

$$[c(\tilde{\alpha}) - c(\alpha_t^L)](1-h_t)dt \tag{21}$$

and the immediate effect on h_t is

$$dh_\tau = (\tilde{\alpha} - \alpha_t^L)(1-h_t)dt. \tag{22}$$

This effect dies out in time:

$$dh_s = dh_\tau - \int_\tau^s dh_v \left(\alpha_v^L + \alpha^H\right) dv,$$

which implies that

$$dh_s = dh_\tau e^{-\int_\tau^s \left(\alpha_v^L+\alpha^H\right)dv}. \tag{23}$$

The effect on output net of switching costs for $s > \tau$ is

$$E_\tau \int_\tau^\infty e^{-\rho(s-\tau)} \left\{ \left[\frac{\partial Y(h_s, a_s)}{\partial h_s} + c(\alpha_s^L) \right] dh_s \right\} ds. \tag{24}$$

This deviation cannot be profitable. Hence, putting together (21), (22), (23) and (24) and using (10) it must be that

$$(\tilde{\alpha} - \alpha_t^L) E_\tau \int_\tau^\infty e^{-\int_\tau^s (\rho + \alpha_v^L + \alpha^H) dv} \left[\frac{\theta}{\theta - 1} e^{a_s} g(h_s) + c(\alpha_s^L) \right] ds - \left[c(\tilde{\alpha}) - c(\alpha_t^L) \right] \leq 0,$$

for any $\tilde{\alpha} \in [0, \overline{\alpha}]$. That is equivalent to stating that a necessary condition for the planner's solution is that α_t^L must maximise

$$\alpha_\tau^L E_\tau \int_\tau^\infty e^{-\int_\tau^s (\rho + \alpha_v^L + \alpha^H) dv} \left[\frac{\theta}{\theta - 1} e^{a_s} g(h_s) + c(\alpha_s^L) \right] ds - c(\alpha_\tau^L). \tag{25}$$

This expression is very similar to (20), the only difference is the term $\theta/(\theta - 1)$ in the integral. Mathematically, finding the solution to the planner's problem is thus equivalent to finding a solution to a game played by agents that maximise (25). We can thus apply the results in Frankel and Burdzy (2005) and obtain a result analogous to Proposition 4.

PROPOSITION 6. *The planner's solution can be implemented by a subsidy equal to* $c(\alpha_\tau^L)/\theta$.

As in Proposition 4, the optimal policy prescribes a subsidy equal to a constant fraction of the switching cost (or, equivalent, a payment of $e^{a_s} g(h_s)/(\theta - 1)$ to all firms in the *High* regime at every t). The important implication is that the planner is not particularly concerned about incentivizing investment when h is low.

5 Concluding Remarks

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

From a methodological point of view, the paper highlights the importance of understanding beliefs that arise in equilibrium for policy analysis. The main result of the paper relies on the link from the business cycle to agents' beliefs about economic activity.

References

Aghion, P., Bacchetta, P. and Banerjee, A. (2004). 'A corporate balance-sheet approach to currency crises', *Journal of Economic Theory*, vol. 119, pp. 6–30.

Allen, F. and Gale, D. (1998). 'Optimal financial crises', *Journal of Finance*, vol. 53(4), pp. 1245–1284.

Angeletos, G.M., Collard, F. and Dellas, H. (2014). 'Quantifying confidence', Manuscript.

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

Angeletos, G.M. and La'O, J. (2009b). 'Noisy business cycles', *NBER Macroeconomics Annual*, vol. 24, pp. 319–378.

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

Aruoba, S.B., Cuba-Borda, P. and Schorfheide, F. (2017). 'Macroeconomic dynamics near the ZLB: A tale of two countries', *Review of Economic Studies*, vol. forthcoming.

Bacchetta, P. and van Wincoop, E. (2016). 'The great recession: A self-fulfilling global panic', *American Economic Journal: Macroeconomics*, vol. 8(4), pp. 177–198.

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

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

Benhabib, J., Schmitt-Grohé, S. and Uribe, M. (2001). 'The perils of taylor rules', *Journal of Economic Theory*, vol. 96(1), pp. 40–69.

Benhabib, J., Schmitt-Grohe, S. and Uribe, M. (2002). 'Avoiding liquidity traps', *Journal of Political Economy*, vol. 110(3), pp. 535–563.

Benhabib, J., Wang, P. and Wen, Y. (2015). 'Sentiments and aggregate demand fluctuations', *Econometrica*, vol. 83(2), pp. 549–585.

Benigno, G. and Fornaro, L. (2015). 'Stagnation traps', Manuscript.

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

Boissay, F., Collard, F. and Smets, F. (2016). 'Booms and banking crises', *Journal of Political Economy*, vol. 124(2), pp. 489–538.

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*, pp. 483 – 500, Amsterdam: North-Holland, ISBN 978-0-444-50083-0, doi:10.1016/B978-044450083-0/50032-0.

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

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

Chamley, C. (1999). 'Coordinating regime switches', *Quarterly Journal of Economics*, vol. 114(3), pp. 869–905.

Chamley, C. (2014). 'When demand creates its own supply: Savings traps', *Review of Economic Studies*, vol. 81(2), pp. 651–680.

Chang, R. and Velasco, A. (2000). 'Financial fragility and the exchange rate regime', *Journal of Economic Theory*, vol. 92(1), pp. 1–34.

Cooper, R. and Haltiwanger, J. (2006). 'On the nature of capital adjustment costs', *Review of Economic Studies*, vol. 73(3), pp. 611–633.

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

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

Evans, G. and Honkapohja, S. (2005). 'Policy interaction, expectations and the liquidity trap', *Review of Economic Dynamics*, vol. 8(2), pp. 303–323.

Farmer, R. (2012). 'Confidence, crashes and animal spirits', *Economic Journal*, vol. 122(559), pp. 155–172.

Farmer, R. and Guo, J.T. (1994). 'Real business cycles and the animal spirits hypothesis', *Journal of Economic Theory*, vol. 63(1), pp. 42–72.

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

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

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

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

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

Guimaraes, B., Machado, C. and Ribeiro, M. (2016). 'A model of the confidence channel of fiscal policy', *Journal of Money, Credit and Banking*, vol. 48(7), pp. 1363–1395.

Guimaraes, B. and Pereira, A.E. (2016). 'QWERTY is efficient', *Journal of Economic Theory*, vol. 163, pp. 819–825.

Guimaraes, B. and Pereira, A.E. (2017). 'Dynamic coordination among heterogeneous agents', Manuscript.

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

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

Heathcote, J. and Perri, F. (2016). 'Wealth and volatility', Manuscript.

Howitt, P. and McAfee, P. (1987). 'Costly search and recruiting', *International Economic Review*, vol. 28(1), pp. 89–107.

Howitt, P. and McAfee, P. (1992). 'Animal spirits', *American Economic Review*, vol. 82(3), pp. 493–507.

Kaplan, G. and Menzio, G. (2016). 'Shopping externalities and self-fulfilling unemployment fluctuations', *Journal of Political Economy*, vol. 124(3), pp. 000–000.

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

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

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

Mertens, K. and Ravn, M. (2014). ‘Fiscal policy in an expectations driven liquidity trap’, *Review of Economic Studies*, vol. 81(4), pp. 1637–1667.

Morris, S. (2014). ‘Coordination, timing and common knowledge’, *Research in Economics*, vol. 68(4), pp. 306–314.

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

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

Murphy, K., Shleifer, A. and Vishny, R. (1989). ‘Industrialization and the big push’, *Journal of Political Economy*, vol. 97(5), pp. 1003–1026.

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

Rendahl, P. (2016). ‘Fiscal policy in an unemployment crisis’, *Review of Economic Studies*, vol. 83(3), pp. 1189–1224.

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

Schaal, E. and Taschereau-Dumouchel, M. (2015). ‘Coordinating business cycles’, Manuscript.

Schaal, E. and Taschereau-Dumouchel, M. (2016). ‘Aggregate demand and the dynamics of unemployment’, Manuscript.

Schmitt-Grohe, S. and Uribe, M. (2017). ‘Liquidity traps and jobless recoveries’, *American Economic Journal: Macroeconomics*, vol. 9(1), pp. 165–204.

Shleifer, A. (1986). ‘Implementation cycles’, *Journal of Political Economy*, vol. 94(6), pp. 1163–1190.

Zawadowski, A. (2013). ‘Entangled financial systems’, *Review of Financial Studies*, vol. 26(5), pp. 1291–1323.

A Proofs

A.1 Proof of Proposition 1

Consider an agent deciding at time normalised 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 $h_t^D = h_0 e^{-\alpha t}$, which is independent of a. Thus, choosing *High* raises his payoff by

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

Therefore this agent will choose *High* if $\underline{U}(a, h_0) \geq 0$. Now, $\underline{U}(a, h_0)$ is continuous and strictly increasing in a, $\lim_{a\to\infty} \underline{U}(a, h_0) = \infty$, and $\lim_{a\to-\infty} \underline{U}(a, h_0) = -\psi$. Thus for any h_0, there is $a = a^*_{pes}(h_0)$ such that $\underline{U}(a, h_0) = 0$. Since $\underline{U}(a, h_0)$ is strictly increasing in a, for any $a' > a^*_{pes}(h_0)$ we have $\underline{U}(a', h_0) > 0$ 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 $\underline{U}(a, h_0)$ is strictly increasing in both a and h_0 and thus $a^*_{pes}(h_0)$ is strictly decreasing.

A similar argument proves that there exists a strictly decreasing threshold $a^*_{opt}(h_0)$ such that if $a < a^*_{opt}(h_0)$, *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 $h_t^U = 1 - (1-h_0)e^{-\alpha t}$, so choosing *High* instead of *Low* raises his payoff by

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

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

Take a pair (a, h_0) such that $a^*_{opt}(h_0) < a < a^*_{pes}(h_0)$. Since $a < a^*_{pes}(h_0)$, if an agent believes that the path of h_t will be h_t^D, then $\underline{U}(a, h_0) < 0$ 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 $a > a^*_{opt}(h_0)$ 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

Since the planner's expected discounted payoff (given a chosen threshold) is a continuous function of σ, solving the planner's problem with $\sigma = 0$ or solving it with $\sigma > 0$ and then taking the limit of the solution when $\sigma \to 0$ must yield the same planner's threshold. Proposition 4 (to be proved next) implies that for any $\sigma > 0$, the planner's threshold is equivalent to the agent's threshold if the investment cost were $\psi - \psi/\theta$ (instead of ψ). It follows from Theorem 2 in Burdzy *et al.* (1998) that when $\sigma \to 0$, the limit of the agents threshold is given by the indifference condition of an agent that believes h_t will either go up forever (i.e., $h_t = h_t^U$, for every t) with probability $1 - h_0$ or go down forever (i.e., $h_t = h_t^U$, for every t) with probability h_0. Therefore, the planner's threshold $a_P^*(h_0)$ is given by the solution to equation (A.1) in Section 3.3. Rearranging, that expression becomes

$$e^{a_P^*(h_0)+\log\left(\frac{\theta}{\theta-1}\right)}\left[(1-h_0)\int_0^\infty e^{-(\rho+\alpha)t}g(h_t^U)dt + h_0\int_0^\infty e^{-(\rho+\alpha)t}g(h_t^D)dt\right] = \psi, \qquad \text{(A.1)}$$

where $g(h)$ is such that $\pi(a,h) = e^a g(h)$.

First statement: The expression in (A.1) implies that the translation $\hat{a}_P^*(h_0) \equiv a_P^*(h_0) - \log(\theta/(\theta-1))$ of the planner's threshold is a curve that lies between $a_{opt}^*(h_0)$ and $a_{pes}^*(h_0)$. Moreover, $\hat{a}_P^*(h_0) = a_{opt}^*(h_0)$ only for $h_0 = 0$ and $\hat{a}_P^*(h_0) = a_{pes}^*(h_0)$ only for $h_0 = 1$.

Since $a_{opt}^*(h) < a_{pes}^*(h)$ for any $h \in [0,1]$,

$$\hat{a}_P^*(0) = a_{opt}^*(0) < a_{pes}^*(0) \text{ and } a_{opt}^*(1) < \hat{a}_P^*(1) = a_{pes}^*(1),$$

which yields

$$\hat{a}_P^*(0) - \hat{a}_P^*(1) < a_{opt}^*(0) - a_{opt}^*(1) \text{ and } \hat{a}_P^*(0) - \hat{a}_P^*(1) < a_{pes}^*(0) - a_{pes}^*(1),$$

and since $\hat{a}_P^*(h) = a_P^*(h_0) - \log[\theta/(\theta-1)]$, we get the claim.

Second statement: Since $\pi(a,h)$ is increasing in h and the planner cares about the discounted sum of agents wealth minus the investment cost, the planner's threshold must lie entirely to the left of the of $a_{opt}^*(h_0)$. If agents that believe everyone will invest in the future are willing to invest, so must be the planner, due to the positive externality of investment.

Suppose agents always play according to the best equilibrium (i.e., the economy is always in the 'good equilibrium'). Let ϵ be the subsidy required to coax agents to invest at some state (a_0, h_0) to the left of the agents threshold $a_{opt}^*(h_0)$. The subsidy ϵ must be such that it makes agents indifferent between investing or not:

$$e^{a_0}\int_0^\infty e^{-(\rho+\alpha)t}g(h_t^U)dt+\epsilon=\psi$$
$$=e^{a^*_{opt}(h_0)}\int_0^\infty e^{-(\rho+\alpha)t}g(h_t^U)dt,$$

where the last equality follows from the fact that agents are indifferent at the threshold $a^*_{opt}(h_0)$. From the equations above we get

$$\epsilon=\psi\left(1-e^{-\left|a_0-a^*_{opt}(h_0)\right|}\right),$$

and thus, the subsidy ϵ is a increasing function of the distance $\left|a_0-a^L(h_0)\right|$. Since the distance between the planner's threshold and $a^*_{opt}(h_0)$ is larger at $h=0$ than at $h=1$, we obtain the result.

The same argument applies if we assume agents will always play according to the 'bad equilibrium'. □

A.3 Proof of Proposition 3

Proposition 3 follows from Frankel and Pauzner (2000), but to apply their results we need to show that for sufficiently high (low) a choosing the *High* (*Low*) regime is the optimal choice, regardless of the actions of others. Consider an agent deciding at some state (a_0,h_0) that expects $h_t=0$, for every $t\geq 0$, with probability one. Her relative gain of choosing regime *High* is given by

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

Since $a_t|a_0\sim N(a_0,\sigma^2t)$, $E_0\left(e^{a_t}\right)=e^{a_0+0.5\sigma^2t}$. Therefore

$$\underline{\mathcal{U}}(a_0)=x_L^{\frac{1}{\theta}}\left(x_H^{\frac{\theta-1}{\theta}}-x_L^{\frac{\theta-1}{\theta}}\right)e^{a_0}\int_0^\infty e^{-(\rho+\alpha-0.5\sigma^2)t}dt-\psi,$$

which implies that there is $\tilde{a}$ such that $\underline{\mathcal{U}}(a_0)>0$ for every $a_0>\tilde{a}$ and every belief over h_t. A similar reasoning shows that there exists $\underset{\sim}{a}$ such that for every $a_0<\underset{\sim}{a}$ the relative gain of choosing regime *Low* is negative. □

A.4 Proof of Proposition 4

First statement: The solution to the planner's problem prescribes investment if the condition in (12) is satisfied (and no investment if the inequality in (12) is reversed). Multiplying both sides of (12) by $(\theta-1)/\theta$ yields the condition for an agent to invest in (16) in an economy where the cost for investing is $\psi - \psi/\theta$.

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

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

Define $b_t = a_t + \log(\theta/(\theta-1))$ and consider the planner's problem in the (b,h)-space. The expression for the planner's decisions is identical to the expression in (16) for the agents' decisions in the decentralised equilibrium (in the (a,h)-space). Moreover, the law of motion for b_t is exactly the same as the law of motion for a_t. Therefore, the solution for the problem must be the same as well.

We know there is a unique decentralised equilibrium given by a threshold a^*, hence $a^* = b^*$, which implies $a^*(h) = a^*_P(h) + \log[\theta/(\theta-1)]$ and yields the claim. $\square$

A.5 Proof of Proposition 5

For a given $h \in [0,1]$ the maximum amount of subsidies the planner have to pay to implement his threshold $a^*_p(h)$ is $\varphi(a^*_p(h), h)$, since agents payoffs are increasing in a. But from Proposition 4 we know that an agent will be indifferent between investing and not investing at $(a^*_p(h), h)$ if the cost is $\psi - \psi/\theta$ and he believes that the others will play according to $a^*_p(h)$. Thus, $\varphi(a^*_p(h), h) = \psi/\theta$, which is independent of h. $\square$

A.6 Proof of Proposition 6

The solution to the planner's problem prescribes a switching rate α^L_τ that maximises (25). Multiplying (25) by $(\theta-1)/\theta$ yields the agents' objective function for an economy where the cost is $c(\alpha^L_\tau)(\theta-1)/\theta$ (see the expression in (20)). $\square$