

Dynamic coordination with timing frictions: Theory and applications

Bernardo Guimaraes¹  | Caio Machado²  | Ana E. Pereira³ 

¹Sao Paulo School of Economics—FGV, Sao Paulo, Brazil

²Instituto de Economía, Pontificia Universidad Católica de Chile, Santiago, Chile

³School of Business and Economics, Universidad de los Andes, Santiago, Chile

Correspondence

Ana E. Pereira, School of Business and Economics, Universidad de los Andes, Santiago, Chile.

Email: apereira@uandes.cl

Funding information

Fondo Nacional de Desarrollo Científico y Tecnológico, Grant/Award Numbers: 11180046, 11190202; Conselho Nacional de Desenvolvimento Científico e Tecnológico

Abstract

We present a general framework of dynamic coordination with timing frictions. A continuum of agents receive random chances to choose between two actions and remain locked in the selected action until their next opportunity to reoptimize. The instantaneous utility from each action depends on an exogenous fundamental that moves stochastically and on the mass of agents currently playing each action. Agents' decisions are strategic complements and history matters. We review some key theoretical results and show a general method to solve the social planner's problem. We then review applications of this framework to different economic problems: network externalities, statistical discrimination, and business cycles. The positive implications of these models are very similar, but the social planner's solution points to very different results for efficiency in each case. Last, we review extensions of the framework that allow for endogenous hazard rates and ex ante heterogeneous agents.

1 | INTRODUCTION

Several problems in economics exhibit strategic complementarities. For example, in a scenario of bank runs, withdrawing deposits from the bank might be the optimal action only if others also do so.¹ The decision about joining and posting on Facebook depends crucially on whether other people are doing the same.² For firms considering whether to invest or not, one important factor is the demand for their goods, which in turn depends on whether other firms choose to

¹See Diamond and Dybvig (1983).

²See Katz and Shapiro (1985) for a model of network externalities.



invest.³ Adopting a new technology may not be the best decision if others in the production chain will keep working with an old technology.⁴

Strategic complementarities are also common in the field of public economics. Evading taxes might be the optimal choice only if many others do so because the stigma associated with evasion and the probability of punishment are both smaller when many others are also evading.⁵ More generally, complying with social norms is more important if most people follow them.⁶ Donating for a charitable project may pay off only if the charity is able to attract enough donations to realize its projects.⁷ The incentives for a high-skill and honest citizen to become a politician depend on whether there are sufficiently many good politicians out there.⁸ Tax exemptions are particularly important for a firm if tax breaks are pervasive and the tax base is small. This yields strategic complementarities in lobbying activity.⁹

A variety of models that capture these economic problems give rise to multiple equilibria. In a range of parameters, different outcomes can arise in equilibrium depending on what agents expect others will do. For example, in models of tax evasion, for an intermediate range of fundamentals, there is an equilibrium where agents evade taxes and another where they do not; in models of network externalities, equilibria predicting the prevalence of different networks coexist; in macroeconomic models, economic activity might depend on arbitrary shifts on expectations.

One important question left unanswered by models with multiple equilibria is what determines which equilibrium will be played. Will firms evade taxes? Will competent and honest citizens choose a career as politicians? Will people coordinate on Facebook, Orkut or Google+? Will economic activity recover next year?

Different approaches to deal with equilibrium multiplicity have been proposed. This survey focuses on the literature of dynamic games with timing frictions, which follows the seminal contribution of Frankel and Pauzner (2000). Time is continuous. Agents choose between two states (say, low and high) and get opportunities to revise their behavior according to a Poisson clock. Their instantaneous utility gain from being in the high state increases in an exogenous fundamental variable and in the fraction of agents in the high state. The fundamental moves according to a Brownian motion.¹⁰

As an illustration, firms might choose procedures that comply with tax regulations or not. The net benefit from evading taxes depends on exogenous factors and on how many people are complying with tax rules. At some random points in time, firms decide whether they want to change their procedures. Here, the Poisson clock can be seen as an attention friction modeled in a reduced-form way. In other applications, the Poisson clock could be related to the maturity of a bond or an investment, or to the obsolescence of a machine.

Many economic problems are naturally described as dynamic environments with coordination motives and evolving fundamentals. Examples include compliance with rules, investment dynamics, technology adoption, bank runs, political behavior, and many others. This modeling approach provides a suitable framework for all these problems. As an alternative,

³See Cooper and John (1988) and Kiyotaki (1988). Investment in human capital can also feature coordination motives (see Palivos & Varvarigos, 2013).

⁴See Matsuyama (1991).

⁵See the models in Kim (2003), Nyborg and Telle (2004), and Sanchez-Villalba (2015).

⁶See the models in Bethencourt and Kunze (2019) and Meunier and Schumacher (2019).

⁷See the model in Andreoni (1998).

⁸See Caselli and Morelli (2004) for a model.

⁹See Ilzetzki (2018) for a model.

¹⁰See also Burdzy, Frankel, and Pauzner (2001).

one could introduce dynamics in global games. However, in general, dynamic global games are not very tractable—and might exhibit equilibrium multiplicity.¹¹ An alternative stream of the literature incorporates dynamics in coordination games through learning, but typically fundamentals are assumed to be fixed.¹²

One key aspect of the class of models studied here is that history matters. Incentives to pick one state are positively affected by the stock of agents in that state. Hence agents might stick to an action for longer periods of time simply because this is what they have been doing in the past. This feature of the model captures some well known dynamic coordination problems. One classic example is the prevalence of QWERTY keyboards (David, 1985), but this is likely to be important in a wide range of applications. The model also fits well cases where we coordinate on a pattern of behavior for such a long time that we may fail to notice a coordination problem.¹³

One important question is about the economic inefficiencies in these problems. We show a general formulation of the planner's problem that allows for a clean comparison between the socially optimal solution and the decentralized equilibrium. As it turns out, the planner's solution is given by a problem very similar to the agents' one, as if the planner were playing a game with its future selves, but with a different utility flow that takes into account the externalities of an agent on others.

The literature has used this framework to study several economic settings. Naturally, in all these applications, the positive implications are qualitatively similar. However, interestingly, the welfare implications are substantially different across applications and are often not obvious at all. For example, in a model of network externalities, a situation with agents stuck in a low-quality network is actually efficient: in the decentralized equilibrium, agents move too soon to a higher-quality network. At the moment agents start to switch, the social future gains from being in a higher-quality network are strictly smaller than the social transition costs. In a business cycle model with a fixed investment cost, firms' investment decisions depend on whether others have been investing, so the economy might get stuck in a situation where firms choose not to invest, but would invest if only others had been doing so. Nevertheless, the planner is not particularly concerned with stimulating investment in these occasions.

Section 2 presents the framework proposed by Frankel and Pauzner (2000) and some results for the general model. First, it presents results on equilibrium multiplicity and equilibrium uniqueness. Then, it shows a method to solve the social planner's problem and applies mathematical results about bifurcation probabilities from Burdzy, Frankel, and Pauzner (1998) to derive expressions for the equilibrium threshold in the limiting cases of vanishing shocks and vanishing frictions. We compare this framework to two popular approaches that deal with equilibrium selection in coordination games: global games and dynamic models with perfect foresight.

With this toolkit in hand, Section 3 shows analytical results for a case with linear preferences and presents applications of the basic framework to a variety of settings such as network externalities, statistical discrimination and stimulus policies in macroeconomics. Besides generating insights for specific questions, the applications illustrate the potential of the model to accommodate a large set of economic problems and the different results regarding efficiency.

¹¹For a class of tractable dynamic global games, see Mathevet and Steiner (2013).

¹²See, for example, Ennis and Keister (2005), Amir and Lazzati (2011), and Amir, Gabszewicz, and Resende (2014).

¹³For example, in 1918, the trade publication Earnshaw's Infants' Department stated that "the generally accepted rule is pink for the boys, and blue for the girls. The reason is that pink, being a more decided and stronger color, is more suitable for the boy, while blue, which is more delicate and dainty, is prettier for the girl." A hundred years later, we might have the impression that we have always coordinated on blue for boys—and some argue that this coordination pattern might be ripe for change.



Last, Section 4 discusses extensions of the framework that allow for endogenous hazard rates and ex ante heterogeneous agents.

Some of the material in this survey is not in the existing literature. Examples include the general formulation of the social planner's problem (Section 2.3); the model with asymmetric network externalities (Section 3.1.4); and the study of efficiency in the model of statistical discrimination (Section 3.2.2).

2 | THE FRAMEWORK

Time is continuous and runs forever. There is a continuum of agents with unit mass indexed by i . There are two possible actions $a_i \in \{0, 1\}$.

Payoffs. An agent's payoff depends on her own action, the action of others and on some exogenous fundamental θ_t . The action of other players is summarized by the proportion of players in action 1, denoted by n_t . The instantaneous payoff of action 0 is given by a function $u_0(\theta_t, n_t)$ and the payoff of action 1 is written as $u_1(\theta_t, n_t)$, both continuously differentiable. We define the instantaneous relative gain of action 1 as

$$\Delta u(\theta_t, n_t) = u_1(\theta_t, n_t) - u_0(\theta_t, n_t).$$

Throughout the text, we assume that $\Delta u(\cdot)$ is increasing in both arguments. It means that agents' incentives to choose action 1 are increasing in the exogenous fundamental and in the action of others (there are strategic complementarities).

Timing frictions. Agents cannot choose their actions at every period. They receive opportunities to revise their actions according to a Poisson clock (independent across agents) with arrival rate δ . Once they choose an action they are locked in the chosen action until the Poisson shock hits again.

Simple examples. In a model where agents choose whether to participate in a given network or not, action 1 corresponds to taking part in the network and θ is a variable summarizing the intrinsic quality of the network. Thus $u_1(\theta_t, n_t)$ is the instantaneous utility from participating, increasing in the quality of the network and in the number of people taking part, and $u_0(\theta_t, n_t)$ can be normalized to 0. Alternatively, actions 0 and 1 can refer to two competing networks, agents choose between them. In this case, $u_1(\theta_t, n_t)$ is increasing in n_t and $u_0(\theta_t, n_t)$ is decreasing in n_t . In a model of occupational choice along the lines of Matsuyama (1991), action 1 corresponds to an occupation that benefits from others choosing it as well (in the example, industry) and action 0 corresponds to an occupation with no such externalities (in the example, agriculture). In these examples, the timing frictions can be thought of as attention frictions.

The discounted expected gain of choosing 1 instead of zero at some date τ is given by

$$V_\tau = \int_\tau^\infty e^{-(\rho+\delta)(t-\tau)} \mathbb{E}[\Delta u(\theta_t, n_t)] dt, \quad (1)$$

where $\rho > 0$ is the discount rate of agents.¹⁴ Notice that an agent's decision today only affects her payoffs until she gets selected by the Poisson process again. That is why the probability of

¹⁴It is implicitly assumed here that $\Delta u(\theta, n)$ and the process for the fundamental are such that this integral is well defined for any initial conditions and continuous path for n .

not getting selected from τ to t , which is given by $e^{-\delta(t-\tau)}$, shows up in the expression above. Hence agents effectively discount the future at rate $\rho + \delta$. They choose 1 if $V_\tau > 0$, 0 if $V_\tau < 0$ and are indifferent between the two actions if $V_\tau = 0$.

An action is dominant if it is optimal for an agent to choose that action regardless of his beliefs about the actions of others. Throughout the text, we assume that the assumption below holds.

Assumption 1. (Existence of dominance regions) There exists $\tilde{\theta}$ and $\underline{\theta}$ such that: if $\theta_i > \tilde{\theta}$ choosing 1 is a strictly dominant action; if $\theta_i < \underline{\theta}$ choosing 0 is a strictly dominant action.

In this setting, a strategy is a map from the history of θ and all past actions to a probability distribution over $\{0, 1\}$. The equilibrium concept used throughout the text is Nash Equilibrium.

2.1 | Multiple equilibrium benchmark

Here we analyze the case where the fundamental is constant and deterministic, that is, $\theta_t = \theta$.

Take an agent deciding whether or not to choose action 1 when $n_t = n_\tau$. Suppose it is optimal for her to choose action 1 if she believes everyone that will choose after her will do the same. Then, players choosing after her have an even higher incentive to choose action 1 if they hold the same belief (because at any time $t > \tau$, n_t will be larger than n_τ). Therefore, everyone choosing action 1 forever is an equilibrium. Conversely, if it is optimal for an agent to choose 0 if she believes everyone will do the same, then everyone choosing action 0 forever is an equilibrium.

To find the values of θ for which everyone choosing action 1 or everyone choosing action 0 is an equilibrium, we compute two boundaries: (a) the boundary where an agent is indifferent between 0 and 1 if she believes that everyone will choose 1 in the future; (b) the boundary where an agent is indifferent if she believes no one will choose action 1 in the future. Those are given by the solutions of

$$\int_{\tau}^{\infty} e^{-(\rho+\delta)(t-\tau)} \Delta u(\theta, n_t^{\uparrow}) dt = 0 \quad (2)$$

and

$$\int_{\tau}^{\infty} e^{-(\rho+\delta)(t-\tau)} \Delta u(\theta, n_t^{\downarrow}) dt = 0, \quad (3)$$

where $n_t^{\uparrow} = 1 - (1 - n_\tau)e^{-\delta(t-\tau)}$ and $n_t^{\downarrow} = n_\tau e^{-\delta(t-\tau)}$.

For any $t > \tau$, $n_t^{\uparrow} > n_t^{\downarrow}$. Hence, for a given θ and n_τ , the LHS in (2) is larger than the LHS in (3). Therefore, for a given n_τ , the value of θ that satisfies (2) must be smaller than the value of θ that satisfies (3). Intuitively, if agents expect all others to choose action 1 in the future, the value of θ they require to be willing to choose 1 is lower than what it would be if they expected others to choose 0 forever.

Moreover, since $n_t^{\uparrow}$ is increasing in n_τ and $\Delta u(\theta, n_t^{\uparrow})$ is increasing in both arguments, the value of θ that satisfies (2) is decreasing in n_τ . Likewise, since $n_t^{\downarrow}$ is increasing in n_τ , the value of θ that satisfies (3) is also decreasing in n_τ .




FIGURE 1 Multiple equilibria with no shocks

The solutions of (2) and (3) are represented by the curves $\underline{\theta}$ and $\bar{\theta}$, respectively, in Figure 1. To the right of $\bar{\theta}$ choosing 1 is a dominant action; to the left of $\underline{\theta}$ choosing 0 is the dominant action; between the boundaries $\underline{\theta}$ and $\bar{\theta}$ we have multiple equilibria and agents' choices depend on their expectations about the actions of others.

In this setting, agents can perfectly anticipate the future actions of others. Hence there is no *strategic uncertainty*. As we see next, this has important implications.

2.2 | Equilibrium uniqueness

Now assume that the fundamentals are subject to stochastic shocks. The key difference from the case without shocks is that now agents may not be able to perfectly anticipate the actions of others in the future. Assume that

$$d\theta_t = \mu dt + \sigma dZ_t,$$

where dZ_t is a standard Brownian motion, μ is a constant drift and $\sigma > 0$ is the volatility. The process above has two properties that are convenient. First, shocks are frequent: for any interval of length $dt > 0$ we have that $\text{Prob}(\theta_t = \theta_{t+dt}) = 0$. Second, their increments are independent of history, that is, $\theta_{t+dt} - \theta_t$ is independent of $(\theta_s)_{s \leq t}$. With shocks, multiplicity disappears.

Theorem 1 (Frankel and Pauzner, 2000). *In the model with shocks, there is an (essentially) unique equilibrium. Agents choose 1 when $\theta_t > \theta^*(n_t)$ and 0 when $\theta_t < \theta^*(n_t)$, where $\theta^*(\cdot)$ is a decreasing function.¹⁵*

Figure 2 illustrates the result. An agent choosing an action at some point on the equilibrium threshold must be indifferent between both actions. Formally, for every $n_\tau \in [0, 1]$, $\theta^*(n_\tau)$ must solve:

$$\int_{\tau}^{\infty} e^{-(\rho+\delta)(t-\tau)} \mathbb{E}[\Delta u(\theta_t, n_t) | \theta^*, \theta^*(n_\tau), n_\tau] dt = 0, \quad (4)$$

¹⁵We say the equilibrium is *essentially* unique because when $\theta_t = \theta^*(n_t)$, agents are indifferent between the two actions and any action would be consistent with Nash equilibrium. Agents' actions at $\theta_t = \theta^*(n_t)$ are irrelevant for the path of n_t for any realization of the Brownian path. Throughout the text, we omit the specification of strategies on the zero-measure set of states that constitute the equilibrium threshold.




FIGURE 2 Unique equilibrium with shocks

where the operator $\mathbb{E}[\cdot | \tilde{\theta}, \theta_\tau, n_\tau]$ denotes agents' expectation when the current state is (θ_τ, n_τ) and they expect others to play according to the threshold $\tilde{\theta}$.

To understand the intuition behind the result, it is useful to go through the proof. The proof here is slightly different from the original proof in Frankel and Pauzner (2000) and a bit more similar to the proof of equilibrium uniqueness in Frankel, Morris, and Pauzner (2003), even though that is a static global-game model.

2.2.1 | Proof of equilibrium uniqueness

A key assumption is that we have dominance regions, as represented by the boundaries $\underline{\theta}_0$ and $\bar{\theta}_0$ in Figure 3. Instead of looking for a strategy profile that is a Nash equilibrium, we are going to look for strategy profiles that survive the iterated elimination of strictly dominated strategies (hereafter, IESDS). This is a less restrictive equilibrium concept, since every Nash Equilibrium also survives IESDS (but not every strategy profile that survives IESDS is a Nash Equilibrium).

A strategy here is a map from every possible history to a probability of choosing action 1 (i.e., a map that prescribes what an agent selected by the Poisson process will choose in every contingency).¹⁶

Iterations from the left. Let's start a process of elimination of strictly dominated strategies. Notice that any strategy that prescribes playing 1 to the left of $\underline{\theta}_0$ is strictly dominated. Therefore, we can remove those strategies from our game (we can also remove strategies that prescribe playing 0 to the right of $\bar{\theta}_0$, but let's leave them there for a while).

We can now ask the same question we asked to construct the boundary $\underline{\theta}_0$: when is an agent indifferent between actions 0 and 1 under the most optimistic belief about $(n_t)_{t \geq \tau}$? The difference now is that, after eliminating the strictly dominated strategies, the meaning of "most optimistic belief" changed. It is no longer possible to assume that agents will play action 1 always, since we eliminated strategies that prescribed playing 1 to the left of $\underline{\theta}_0$ and the probability of θ_t crossing the boundary $\underline{\theta}_0$ at some point is positive. Thus, the most optimistic belief now is the one in which other agents play 1 in any state (θ_t, n_t) , except in those states to the left of $\underline{\theta}_0$. Under those beliefs, an agent choosing at the point P in Figure 3 will no longer be indifferent between action 1 and 0. She will strictly prefer to choose action 0 (if she was indifferent under a more optimistic belief, now she is strictly preferring to choose action 0). Thus, the boundary where an agent is indifferent

¹⁶As it turns out, the strategy played in the unique equilibrium of this game will depend only on current values of θ and n almost everywhere, even though agents are allowed to choose richer strategies.




FIGURE 3 Dominance regions

under the most optimistic belief shifts to the right, as illustrated by the boundary $\underline{\theta}_1$ in Figure 4.

Now we can move to the second round of elimination of strictly dominated strategies. After eliminating the strategies that imply playing 1 in the dark gray area in Figure 4, the strategies that imply playing 1 in the light gray area also become strictly dominated strategies (agents would not choose it even under the most optimistic belief possible). Thus, we further eliminate the strategies that imply playing 1 to the left of $\underline{\theta}_1$ and we get a boundary $\underline{\theta}_2$. Continuing this process indefinitely we get a boundary $\underline{\theta}_\infty$, as illustrated in Figure 5. Notice that $\underline{\theta}_\infty$ must be to the left of $\bar{\theta}_0$, since the latter was constructed assuming that agents believe everyone will always choose 0, while the former assumed they believe that everyone will choose 0 to the left of $\underline{\theta}_\infty$ and 1 to the right of it (and thus they require a lower initial θ to be indifferent).

So far we know that in any strategy that survives IESDS agents play 0 to the left of $\underline{\theta}_\infty$. Notice also that an agent making a decision on the threshold $\underline{\theta}_\infty$ must be indifferent between 0 and 1 (otherwise the iterations could continue). Thus, everyone playing 1 to the right of $\underline{\theta}_\infty$ and 0 to the left of it must be a Nash equilibrium. We have just proved equilibrium existence.

Iterations from the right. We can repeat an analogous procedure starting from the boundary $\bar{\theta}_0$, by first eliminating the strategies in which agents play 0 to the right of $\bar{\theta}_0$. The only difference is that now we will always consider the most *pessimistic* belief possible about n_t in each iteration, given the strategies that we have already eliminated. This process will yield a boundary $\bar{\theta}_\infty$. In any equilibrium that survives IESDS, agents must play 1 to the right of $\bar{\theta}_\infty$. As was the case for $\underline{\theta}_\infty$, notice that everyone playing according to $\bar{\theta}_\infty$ is an equilibrium, since agents are indifferent between both actions on the threshold $\bar{\theta}_\infty$ if they believe others will play according to that threshold.




FIGURE 4 First round of elimination of dominated strategies




FIGURE 5 Convergence of iterations from the left

Limit of iterations coincide. Notice that $\bar{\theta}_\infty$ cannot be to the left of $\underline{\theta}_\infty$, otherwise there would be no strategy that survives IESDS (we know it cannot be true since a Nash Equilibrium exists and every Nash equilibrium survives IESDS).¹⁷

If $\bar{\theta}_\infty$ and $\underline{\theta}_\infty$ coincide for every n , then we have an essentially unique equilibrium. Suppose by contradiction that this is not the case, as exemplified in Figure 6. Then, we can always get a translation of $\bar{\theta}_\infty$ that lies entirely to the left of $\underline{\theta}_\infty$, but touches it in at least one point, as represented by the curve $\bar{\theta}'_\infty$ in the figure.

Consider a player that is choosing at point A and believes everyone will play according to $\bar{\theta}'_\infty$ (hereafter, “player A ”). Her payoff of choosing 1 must be at least zero, since it would be zero if she believed everyone would play according to $\underline{\theta}_\infty$ and therefore choose action 1 (weakly) less often for any Brownian path (remember that $\underline{\theta}_\infty$ is an equilibrium).¹⁸

Consider now the payoff of a player that is choosing at point B and believes everyone will play according to $\bar{\theta}_\infty$ (“player B ”). Player B faces the same initial n as player A , but a larger initial θ . For player B , the increments $(\theta_{\tau+s} - \theta_\tau)$ follow the same distribution as the increments for player A . In other words, for any Brownian path $(Z_t)_{t \geq \tau}$, player B observes the same fundamentals as player A plus a positive constant (i.e., the distance between their thresholds). Since $\bar{\theta}'_\infty$ is a translation of $\bar{\theta}_\infty$, for any Brownian path, player A will have θ_t to the left (right) of her threshold if, and only if, player B also has her θ_t to the left (right) of her threshold. Thus, they will always experience the same dynamics of n_t , but player B always experiences a higher fundamental. Thus, the relative payoff of player B of choosing action 1 (which we know to be zero, since $\bar{\theta}_\infty$ is an equilibrium) must be higher than that of player A . Thus we get the contradiction:

$$0 = (\text{Player } B\text{'s gain of choosing 1}) > (\text{Player } A\text{'s gain of choosing 1}) \geq 0. \quad \square$$

2.3 | The social planner's problem

We now study the social planner's problem. Let's say welfare is given by the discounted sum of individual agents' payoffs (which depend on others' actions as well). Thus, the planner maximizes:

$$\mathbb{E} \left[\int_{\tau}^{\infty} e^{-\rho(t-\tau)} W(\theta_t, n_t) dt \right], \quad (5)$$

¹⁷If $\bar{\theta}_\infty$ is to the left of $\underline{\theta}_\infty$, then there is a subset V of the state space such that: (a) every strategy that survives IESDS requires playing 1 when $(\theta_t, n_t) \in V$; (b) every strategy that survives IESDS requires playing 0 when $(\theta_t, n_t) \in V$. Obviously, no strategy can satisfy (a) and (b) at the same time.

¹⁸Notice that the definition of $\bar{\theta}'_\infty$ implies that it might coincide with $\bar{\theta}_\infty$.




FIGURE 6 Translations

where

$$W(\theta, n) = nu_1(\theta, n) + (1 - n)u_0(\theta, n). \quad (6)$$

At each date t , the planner chooses the proportion $\phi_t \in [0, 1]$ of agents that received a chance to switch actions that will pick action 1.

Suppose that at a given date τ it is optimal for the planner to choose $\phi_\tau < 1$ and consider the following deviation: the planner increases ϕ_τ in $\Delta\phi > 0$ units today, but keeps the future values of ϕ_t unchanged, for any realization of the Brownian path. Increasing ϕ_t by $\Delta\phi$ today, raises n_τ by $\delta\Delta\phi dt \equiv d\phi$. But at a given date $t > \tau$ a proportion $e^{-\delta(t-\tau)}$ of the additional $d\phi$ agents who switched to action 1 will have already been selected by the Poisson process again. Therefore, the resulting change in n at time $t \geq \tau$ is:

$$dn_t = d\phi e^{-\delta(t-\tau)}.$$

This deviation is not profitable if

$$\int_{\tau}^{\infty} \mathbb{E}_{\tau} \left[e^{-\rho(t-\tau)} \frac{\partial W(\theta_t, n_t)}{\partial n} e^{-\delta(t-\tau)} d\phi \right] dt \leq 0,$$

which becomes

$$\int_{\tau}^{\infty} e^{-(\rho+\delta)(t-\tau)} \mathbb{E}_{\tau} \left[\frac{\partial W(\theta_t, n_t)}{\partial n} \right] dt \leq 0. \quad (7)$$

Now assume that $\phi_\tau > 0$ and the planner chooses a similar deviation, but with $d\phi < 0$. The same reasoning implies that this deviation is not profitable if

$$\int_{\tau}^{\infty} e^{-(\rho+\delta)(t-\tau)} \mathbb{E}_{\tau} \left[\frac{\partial W(\theta_t, n_t)}{\partial n} \right] dt \geq 0. \quad (8)$$

This implies the following necessary conditions for optimality: if $\phi_\tau = 0$ then (7) holds; if $\phi_\tau = 1$ then (8) holds; if $\phi_\tau \in (0, 1)$ then (7) holds with equality. But those are exactly the necessary (and sufficient) conditions for a Nash Equilibrium in the game where the relative payoff $\Delta u(\theta, n)$ is replaced by $\frac{\partial W(\theta, n)}{\partial n}$ (see Equation (1)).

Hence, if we find the set of Nash Equilibria in this modified game, we have found all the candidates for the planner solution. But as long as $\frac{\partial W(\theta, n)}{\partial n}$ satisfies the same conditions we imposed on $\Delta u(\theta, n)$ and we have shocks, the equilibrium is unique, and therefore, these necessary conditions are also sufficient for optimality.

Therefore, we can think that the planner plays a modified game with its future selves in which the flow gain of choosing 1 is given by

$$\frac{dW(\theta, n)}{dn} = \underbrace{[u_1(\theta, n) - u_0(\theta, n)]}_{\Delta u(\theta, n)} + \underbrace{\left[n \frac{\partial u_1(\theta, n)}{\partial n} + (1 - n) \frac{\partial u_0(\theta, n)}{\partial n} \right]}_{\text{Externality}}. \quad (9)$$

The agent considers the first term in brackets, but does not consider the externality on others. The difference between the planner's problem and the agents' problem is only the second term in brackets. There are no other fundamental differences between the planner's solution and the decentralized equilibrium. In particular, the planner in its modified game and the agents in the original game have the same effective discount rate, $\rho + \delta$. The externality in (9) is the obstacle to efficiency in the Frankel-Pauzner framework.

This result allows us to understand the inefficiencies in a dynamic coordination model even if we cannot characterize the equilibrium threshold. In Section 3, we will use this result to compare the decentralized equilibrium and the social optimum in different applications of the general framework.

2.4 | Bifurcation probabilities

We now discuss other results that will be useful when explicitly characterizing the equilibrium in some particular cases. For simplicity of the exposition, assume $\mu = 0$ hereafter, unless stated otherwise. Suppose agents play according to some decreasing threshold $\theta^*(n)$ (as we have seen, they do that in equilibrium).¹⁹ Let's look at the limiting case where shocks are very small, that is, $\sigma \rightarrow 0$.²⁰ What is the dynamics of n_t when the economy starts off at some point on the threshold $\theta^*(n)$?

This mathematical problem is studied by Burdzy et al. (1998). They show that, as $\sigma \rightarrow 0$, the economy either bifurcates in the direction of $n = 0$ or in the direction of $n = 1$, and never comes back. Moreover, the time it takes for the economy to start heading off in one direction or the other goes to zero. The economy instantaneously moves in one of the directions indicated by the arrows in Figure 7.

When the economy is at the threshold, a small positive shock pushes the economy to the right of it and n_t goes up. Negative shocks push the economy to the left of the threshold and thus n_t goes down. Since the threshold is negatively sloped, once the economy stays away for a while to the right of the threshold, as n_t goes up, it drifts further away from the threshold and never comes back. The same logic applies when a sequence of negative shocks hit. Since shocks are very small and frequent, the bifurcation happens very quickly. Burdzy et al. (1998) shows that as $\sigma \rightarrow 0$, the time until it bifurcates converges to zero.

¹⁹There are additional technical requirements that are always satisfied by an equilibrium threshold. In particular, $\theta^*(\cdot)$ must be a Lipschitz function.

²⁰The result presented in this section holds for $\mu, \sigma \rightarrow 0$, regardless of the relative speed at which μ and σ go to zero.




FIGURE 7 Bifurcation probabilities

But what is the probability of the economy moving in either direction? Burdzy et al. (1998) show that it is proportional to the speed at which the economy moves in each direction once it bifurcates. More specifically:

$$\frac{\text{Prob}(\text{bifurcate up})}{\text{Prob}(\text{bifurcate down})} = \frac{\text{Initial speed of } n_t \text{ if bifurcates up}}{\text{Initial speed of } n_t \text{ if bifurcates down}} = \frac{|\dot{n}_0^\uparrow|}{|\dot{n}_0^\downarrow|}, \quad (10)$$

which implies that

$$\text{Prob}(\text{bifurcate up}) = 1 - n_0 \text{ and } \text{Prob}(\text{bifurcate down}) = n_0. \quad (11)$$

To understand the intuition, suppose $\sigma \approx 0$ and that n_t moves up very quickly when to the right of the threshold and goes down very slowly when to the left of it. Then, a small sequence positive shocks moves the economy very far from the threshold, but we need a very large sequence of negative shocks to get far away from it. If positive and negative shocks are equally likely, it makes sense to think that it is more likely that economy will bifurcate up.

2.5 | Results for limiting cases

Although in general it is not easy to compute the equilibrium threshold $\theta^*(n)$, in two limiting cases the model is especially tractable: (a) when $\sigma \rightarrow 0$ and (b) when $\delta \rightarrow \infty$. Besides providing approximate results for the cases with small shocks and small frictions, these limiting cases are interesting for helping us to understand how small departures from the case with no shocks and/or no timing frictions affect equilibrium outcomes.

When analyzing the cases of vanishing shocks or vanishing frictions, there is no need to assume a constant drift term μ . One can allow the trend in fundamentals to depend on t and θ . Specifically, all results stated for the limiting cases continue to hold if instead of μ the drift term of the stochastic process is $\mu \cdot \varphi(t, \theta)$, for any continuously differentiable function $\varphi(\cdot)$ that is bounded on t .

2.5.1 | Vanishing shocks

When $\mu, \sigma \rightarrow 0$, we can use the results from Burdzy et al. (1998), namely that the time until bifurcation converges to zero and the bifurcation probabilities in (11), to compute the beliefs of an agent deciding on the threshold. The indifference condition in (4) becomes:




FIGURE 8 Equilibrium with vanishing shocks

$$(1 - n_0) \int_0^\infty e^{-(\rho+\delta)t} \Delta u(\theta^*(n_0), n_t^\uparrow) dt + n_0 \int_0^\infty e^{-(\rho+\delta)t} \Delta u(\theta^*(n_0), n_t^\downarrow) dt = 0, \quad (12)$$

where $n_t^\uparrow = 1 - (1 - n_0)e^{-\delta t}$ and $n_t^\downarrow = n_0 e^{-\delta t}$ (we normalize the decision period to $\tau = 0$). Solving the expression above for θ^* we get the equilibrium threshold. Alternatively, we can apply the changes of variables $v = n_t^\uparrow$ and $v = n_t^\downarrow$ in the integrals above to get:

$$\int_0^{n_0} \left(\frac{v}{n_0} \right)^{\frac{\rho}{\delta}} \frac{1}{\delta} \Delta u(\theta^*(n_0), v) dv + \int_{n_0}^1 \left(\frac{1-v}{1-n_0} \right)^{\frac{\rho}{\delta}} \frac{1}{\delta} \Delta u(\theta^*(n_0), v) dv = 0. \quad (13)$$

Figure 8 depicts the equilibrium. The curves θ^{opt} and θ^{pes} represent the two extreme equilibria of the game with no shocks—(“*opt*” and “*pes*” stand for optimistic and pessimistic beliefs, respectively). Notice that the threshold θ^* touches θ^{opt} when $n = 0$, since the bifurcation probabilities imply that agents expect n_t to go up forever with probability one at the state $(\theta^*(0), 0)$. Similarly, θ^* touches θ^{pes} when $n = 1$, since the system bifurcates down with probability one at the state $(\theta^*(1), 1)$.

2.5.2 | Vanishing frictions

With vanishing frictions ($\delta \rightarrow \infty$) agents are allowed to switch actions at almost every period. To derive the equilibrium in this case, it is useful to write down our environment in a different unit of time. Suppose that date t in the new unit of time is represented by $\tilde{t} = \delta t$. For instance, if we were initially measuring time in years and $\delta = 12$, it means that we are now measuring time in months. The arrival rate of the Poisson process becomes $\tilde{\delta} = 1$; the discount rate becomes $\tilde{\rho} = \rho/\delta$; the variance of the change in fundamentals between $\tilde{t}$ and $\tilde{t} + 1$ becomes $\tilde{\sigma}^2 = \sigma^2/\delta$.²¹ Sending δ to infinity is thus equivalent to taking the limit when $\rho \rightarrow 0$ and $\sigma \rightarrow 0$.

We know from the previous section that when $\sigma \rightarrow 0$, the threshold is characterized by (13). Taking the limit when $\rho \rightarrow 0$ in (13), we get that the equilibrium threshold $\theta^*(n_0)$ under vanishing frictions is given by:

$$\int_0^1 \Delta u(\theta^*(n_0), v) dv = 0. \quad (14)$$

²¹ $(\theta_{t+s} - \theta_t) \sim N(0, \sigma^2 s)$, for every t and s . Thus, σ^2 represents the variance of $(\theta_{t+1} - \theta_t)$.




FIGURE 9 Equilibrium with vanishing frictions

Notice that the expression above does not depend on n_0 , and thus $\theta^*(n_0)$ is constant. Figure 9 shows the equilibrium in this case.

The equilibrium threshold does not depend on the current value of n , just on the fundamental. The result may seem to follow immediately from the fact that n_t moves very quickly in this economy (true), which would imply that the current value of n_t is irrelevant (false). This intuition is incorrect because although the economy can move very quickly from a low n to a high n , the next opportunity to change behavior also comes very quickly (the agent effectively discounts the future at rate $\rho + \delta$). Indeed, the dominance region boundaries are not vertical.²²

The intuition for the irrelevance of the current value of n is the following. An agent starting on the threshold at $n = 0$ will experience n from 0 to $n^\dagger$, where $n^\dagger$ is uniformly distributed in $[0, 1]$, while an agent starting on the threshold at $n = 1$ will experience n from 1 to $1 - n^\dagger$. Hence an agent starting at $n = 0$ expects to face lower values of n than an agent starting at $n = 1$ (before they receive another opportunity to choose actions). However, for the agent starting at $n = 0$, the economy moves up very fast at lower values of n , but very slowly as n grows toward 1. For the agent choosing at $n = 1$, on the other hand, the economy moves down very fast initially, and then slowly as it approaches 0. It turns out that these effects cancel out, so agents at $n = 0$ and $n = 1$ are indifferent between actions 0 and 1 for the same value of θ .²³

2.6 | The planner's problem with vanishing shocks

For the planner, it makes no difference whether there are no shocks or shocks are very small (the planner does not face a coordination problem with itself). Hence the model can be solved in two different ways: assuming vanishing shocks ($\sigma \rightarrow 0$) or no shocks ($\sigma = 0$).

The case $\sigma \rightarrow 0$. As shown in Section 2.3, as long as $\partial W(\theta, n)/\partial n$ satisfies the same conditions we imposed on $\Delta u(\theta, n)$ and we have shocks, the planner's problem can be seen as analogous to the agents' problem, only with different payoffs. Considering vanishing shocks, we can then apply the bifurcation probabilities to solve for the planner's threshold. The indifference condition characterizing the planner's threshold is given by

$$(1 - n_0) \int_0^\infty e^{-(\rho+\delta)t} \frac{\partial W(\theta^p, n_t^\uparrow)}{\partial n} dt + n_0 \int_0^\infty e^{-(\rho+\delta)t} \frac{\partial W(\theta^p, n_t^\downarrow)}{\partial n} dt = 0,$$

²²For each initial $n = n_0$ the dominance region boundaries satisfy $\int_0^\infty e^{-(\rho+\delta)t} \Delta u(\bar{\theta}, n_t^\uparrow) dt = 0$ and $\int_0^\infty e^{-(\rho+\delta)t} \Delta u(\underline{\theta}, n_t^\downarrow) dt = 0$, where $n_t^\uparrow = 1 - (1 - n_0)e^{-\delta t}$ and $n_t^\downarrow = n_0 e^{-\delta t}$. Applying the changes of variables $v = n_t^\uparrow$ and $v = n_t^\downarrow$ to the integrals, taking the limit when $\delta \rightarrow \infty$ and rearranging, those conditions become $\int_{n_0}^1 \Delta u(\bar{\theta}, v) dv = 0$ and $\int_0^{n_0} \Delta u(\underline{\theta}, v) dv = 0$. Hence $\bar{\theta}$ and $\underline{\theta}$ are downward sloping under vanishing frictions.

²³See also the explanation in Burdzy et al. (2001).

where again $n_t^\uparrow = 1 - (1 - n_0)e^{-\delta t}$ and $n_t^\downarrow = n_0e^{-\delta t}$, and $\partial W(\theta, n)/\partial n$ is given by Equation (9). Solving this equation for θ^P gives us the planner's threshold as a function of n_0 .

The case $\sigma = 0$. We can also compute the planner's threshold by searching for the curve $\theta^P(n_0)$ along which the planner is indifferent between sending agents to network 1 forever, or to network 0 forever, assuming there are no shocks to fundamentals. Assuming $\sigma = 0$, for a given n_0 the planner is indifferent between an upward or a downward path for n_t when θ^P satisfies:

$$\begin{aligned} & \int_0^\infty e^{-\rho t} \left[n_t^\uparrow u_1(\theta^P, n_t^\uparrow) + (1 - n_t^\uparrow) u_0(\theta^P, n_t^\uparrow) \right] dt \\ &= \int_0^\infty e^{-\rho t} \left[n_t^\downarrow u_1(\theta^P, n_t^\downarrow) + (1 - n_t^\downarrow) u_0(\theta^P, n_t^\downarrow) \right] dt. \end{aligned}$$

The two approaches must lead to the same resulting threshold $\theta^P(n_0)$.

2.7 | Relation to the literature

This section compares the framework studied in this paper with alternative ways to deal with equilibrium selection in coordination games. We focus on two approaches that are closely related to this one: global games and dynamic models with perfect foresight.²⁴

2.7.1 | Relation to “global games”

Among the alternative ways to understand agents' behavior in settings with strategic complementarities, the so-called global games literature is a particularly popular one (see Carlsson & Damme, 1993; Morris & Shin, 1998, and Morris & Shin, 2003). The basic global-game models are static coordination games where agents observe a noisy idiosyncratic signal on a fundamental variable that affects payoffs (θ).²⁵ Building on the theory of global games, a large applied literature has evolved.²⁶

In its simplest versions, agents simultaneously take a binary action, and this noisy idiosyncratic signal about θ is the only (meaningful) information they have. As long as there are dominance regions, it is never common knowledge that the fundamental variable lies in a region where the equilibrium could be driven by self-fulfilling beliefs. Moreover, agents do not know what others are doing in equilibrium. Based on their own signals, agents form expectations about others' information and expectations about others' expectations (higher-order expectations). One key result is that a unique equilibrium arises. The intuition for uniqueness is similar to the intuition for Theorem 1. An agent with a signal about fundamentals very close to the region where, say, action 1 is dominant considers that others might think that fundamentals lie in the dominance region and will choose action 1. Hence action 1 is the best choice for this agent. This triggers a similar process of iterated elimination of strictly dominated strategies that ends up with a unique rationalizable strategy.

²⁴The literature on evolutionary game theory also offers insights to equilibrium selection. See, for example, Kandori, George, and Rob (1993) and Fudenberg and Harris (1992).

²⁵There are several dynamic models employing the global-game methodology. Examples include Angeletos, Hellwig, and Pavan (2007), Dasgupta (2007), Steiner (2008), Chassang (2010), Dasgupta, Steiner, and Stewart (2012), Kováč and Steiner (2013), and Mathevet and Steiner (2013). However, dynamics is not needed to generate equilibrium uniqueness and the basic insights from the methodology—actually, dynamics sometimes even restores multiplicity in those frameworks. This section highlights the analogies between the dynamic framework considered here and static global-game models.

²⁶Angeletos and Lian (2016) offer a comprehensive survey of this literature.



When agents also have access to public information, multiplicity might be restored. Intuitively, public information works as an anchor for higher-order beliefs, as agents' beliefs about others' expectations about θ will be strongly influenced by the information shared by everyone. In the limit of very accurate public information, we are close to the model with complete information about fundamentals. However, in cases with normal noise, as long as the idiosyncratic information about fundamentals is sufficiently more accurate than any public information agents might have, global-game models still yield a unique equilibrium.

Section 2.5.2 shows that in case $\delta \rightarrow \infty$, history plays no role. This case can thus be seen as an equilibrium selection device and can be more easily compared to static global games. Interestingly, the expression in (14) coincides with the equilibrium condition in a static global game where agents have a diffuse prior, get a noisy signal about θ and their payoffs are $\Delta u(\theta, n)$.²⁷ Both methodologies “select” the risk-dominant Nash equilibrium of the corresponding complete-information static game.

This dynamic framework seems different from global games because all information is common knowledge. Morris (2014) shows that this intuition is misleading: what matters is not players' actual beliefs but what their beliefs were at their last opportunity to switch actions. Now consider that an event is *effectively* known by a player if she knew it the last time she had an opportunity to change behavior. Then there is a tight connection between the lack of *effective* common knowledge here and the lack of common knowledge in global-game models (See Morris, 2014).

2.7.2 | Relation to dynamic models with perfect foresight

Another approach to equilibrium selection is used by the seminal paper of Matsui and Matsuyama (1995) and the literature that followed (see, e.g., Hofbauer & Sorger, 1999; Oyama & Tercieux, 2009, and Oyama, Takahashi, & Hofbauer, 2008). They study 2×2 games with two strict Nash equilibria. Every period each player is randomly matched to another player drawn from a large population. As in the game studied here, there are timing frictions: players' opportunities to switch actions are determined by a Poisson process. Differently from the framework studied here, there are no shocks.²⁸

To predict the long-run outcome of such interactions, they introduce two stability concepts, informally explained in what follows. First, a Nash equilibrium of the original static game is said to be *globally accessible* if for any initial action distribution, there exists an equilibrium path of the associated dynamic game that converges to that Nash equilibrium. Second, a Nash equilibrium is said to be *absorbing* if any equilibrium path converges to that Nash equilibrium, whenever the initial action distribution is in a neighborhood of that Nash equilibrium. In other words, if an equilibrium is *globally accessible* it *can* be reached in equilibrium, *regardless* of the initial conditions. If an equilibrium is *absorbing* it *must* be reached for *some* initial conditions.

When timing frictions are small, they show that one Nash equilibria becomes *uniquely absorbing* (i.e., it is the unique absorbing equilibrium). Moreover the uniquely absorbing equilibrium is *globally accessible*. This makes the uniquely absorbing equilibrium a natural choice to predict long run outcomes among different strict Nash equilibria. Although our framework is slightly different from theirs, we can check if our game admits a uniquely absorbing equilibrium that is also globally accessible.

²⁷This equivalence also holds in cases with heterogeneous agents (see Guimaraes & Pereira, 2017).

²⁸Burdzy et al. (2001) study the Matsui and Matsuyama (1995) model with shocks and obtain results similar to Frankel and Pauzner (2000).




FIGURE 10 Perfect foresight selection (a) Perfect foresight selection with vanishing frictions; (b) Perfect foresight selection with non-vanishing frictions

Suppose θ is constant ($\sigma = 0$) and let n_0 denote the initial proportion of players locked in action 1. In the context of our model, global accessibility and absorption mean the following: an $n \in [0, 1]$ is said to be globally accessible if for any $n_0 \in [0, 1]$ there exists an equilibrium path that converges to n as time goes to infinity. An $n \in [0, 1]$ is absorbing if there is a neighborhood U of n , such that for any $n_0 \in U$, any equilibrium path converges to n .

Figure 10a shows the dominance regions boundaries $\underline{\theta}$ and $\bar{\theta}$ when $\sigma = 0$ and $\delta \rightarrow \infty$ (small frictions, no shocks). The equilibrium threshold with $\sigma \rightarrow 0$ and $\delta \rightarrow \infty$ is represented by θ^* . As argued in Section 2.5.2, the dominance regions boundaries are not vertical. Recall that $\underline{\theta}$ must touch θ^* when $n = 0$ and $\bar{\theta}$ must touch θ^* when $n = 1$. Figure 10a illustrates that, in the game with $\sigma = 0$, $n = 1$ is uniquely absorbing and globally accessible to the right of θ^* , while $n = 0$ is uniquely absorbing and globally accessible to the left of θ^* . To see that, suppose θ is slightly above θ^* . When n_0 is sufficiently high, any equilibrium path must converge to $n = 1$ (since (θ, n_0) would be in the dominance region), but there is no n_0 that makes $n = 0$ absorbing—and thus $n = 1$ is uniquely absorbing. Also, $n = 1$ is globally accessible, since everyone playing action 1 is an equilibrium in that region for any n_0 . An analogous reasoning applies to the area where $n = 0$ is uniquely absorbing and globally accessible. Hence, the unique Nash equilibrium with vanishing shocks and frictions in the Frankel and Pauzner (2000) framework is the same as the equilibrium selected by the approach of Matsui and Matsuyama (1995) with vanishing frictions.

With finite frictions, the Matsui and Matsuyama (1995) approach fails to deliver a unique selection device for any θ . This is illustrated in Figure 10b, which reproduces Figure 10a with non-vanishing frictions. In the white region, both $n = 0$ and $n = 1$ are absorbing and none is globally accessible. Hence, in that region the Matsui and Matsuyama (1995) approach does not yield a unique prediction on which equilibrium will be played, while the Frankel and Pauzner (2000) approach does.



3 | APPLICATIONS

We now consider applications of the framework presented in Section 2 to a variety of settings. All these models can be seen as particular cases of the basic framework, but their implications are substantially different. In particular, the efficiency results are quite different across models.

This section focuses on network externalities, statistical discrimination, and business cycles, but the Frankel-Pauzner framework has been employed in the study of several other questions. Plantin and Shin (2018) study carry trades and monetary spillovers across countries in a model where current trading profits of a trader are positively affected by future inflows. Crouzet, Gupta, and Mezzanotti (2018) consider an environment where agents choose between cash and electronic money. Using the Indian demonetization of 2016, they find empirical support for the model.

3.1 | Linear utility

We now consider a particular case of the model where the agents' utility is a linear function of θ_t and n_t , as in Guimaraes and Pereira (2016). This case is particularly tractable as it allows for closed form solutions in the limiting cases previously discussed.

Agents' utility functions are given by

$$u_t^0(\theta_t^0, n_t) = \theta_t^0 + \nu^0(1 - n_t) \quad \text{and} \quad u_t^1(\theta_t^1, n_t) = \theta_t^1 + \nu^1 n_t,$$

where θ_t^j represents the fundamentals affecting the flow-payoff of action j and follows a Brownian motion with drift μ_j and variance σ_j^2 . The flow-utility of those choosing action j increases linearly in the mass of agents taking the same action. We can write the relative payoff function as

$$\Delta u(\theta_t, n_t) = \theta_t + \gamma n_t,$$

where $\theta_t \equiv \theta_t^1 - \theta_t^0 - \nu^0$ follows a Brownian motion with drift $\mu = \mu_1 - \mu_0$ and variance $\sigma = \sigma_0^2 + \sigma_1^2$, and $\gamma \equiv \nu^0 + \nu^1$. Assume ν^0 and ν^1 are such that $\gamma > 0$. Notice $\Delta u(\cdot)$ satisfies all the assumptions listed in Section 2.

This is a general linear case. When $\mu_0 = \sigma_0^2 = \nu^0 = 0$, one of the actions yields a constant payoff. When $\nu^0 = \nu^1$, we have a model of network externalities that are the same for both networks.

3.1.1 | Equilibrium threshold in limiting cases

As in the general case, an agent who receives an opportunity to revise her choice at time τ will choose action $a_i = 1$ if the discounted relative payoff of doing so is positive:

$$\mathbb{E} \int_{\tau}^{\infty} e^{-(\rho+\delta)(t-\tau)} (\theta_t + \gamma n_t) dt > 0. \quad (15)$$

If the inequality is reversed, the agent will choose $a_i = 0$. Theorem 1 ensures there is a unique equilibrium in which agents play according to a downward sloping threshold. If we focus on the limiting case with $\mu, \sigma \rightarrow 0$, we can simply apply the bifurcation probabilities of Section 2.4 to

compute the equilibrium threshold. Substituting the expression for $\Delta u(\theta_t, n_t)$ in this tractable case into (12), we have that the indifference condition for agents is given by

$$(1 - n_0) \int_0^\infty e^{-(\rho+\delta)t} (\theta^l + \gamma n_t^\uparrow) dt + n_0 \int_0^\infty e^{-(\rho+\delta)t} (\theta^l + \gamma n_t^l) dt = 0.$$

Solving for θ^l , we get an explicit linear expression for the equilibrium threshold:

$$\theta^l(n_0) = -\frac{\gamma\delta}{\rho + 2\delta} - \frac{\gamma\rho}{\rho + 2\delta} n_0. \quad (16)$$

It is also possible to explicitly characterize the equilibrium threshold in the other tractable limiting case: when timing frictions vanish. It suffices to substitute $\Delta u(\tilde{\theta}^l, n) = \tilde{\theta}^l + \gamma n$ into (14) and solve for $\tilde{\theta}^l$. Analogously, one could simply take the limit of the right-hand side of (16) as $\delta \rightarrow \infty$ to find that the threshold is given by $\tilde{\theta}^l = -\gamma/2$.

3.1.2 | The planner's problem with network externalities

We now use the model with linear utility to study efficiency in an environment with network externalities. The expression in (6) and simple algebra yield the planner's instantaneous payoff:

$$W(\theta, n) = \theta^0 + \nu^0 + (\theta - \nu^0)n + \gamma n^2,$$

hence

$$\frac{dW(\theta, n)}{dn} = \underbrace{[\theta + \gamma n]}_{\Delta u(\theta, n)} + \underbrace{\gamma n - \nu^0}_{\text{Externality}}.$$

Notice that increasing n implies a positive externality term if and only if $\gamma n > \nu^0$, that is, if $\nu^1 n > \nu^0(1 - n)$. Intuitively, the choice of action 1 increases the payoff of agents locked in action 1 but reduces the payoff of those locked in action 0. Therefore, the sign of the externality term depends on which of those effects dominate, which in turn depends on the amount of agents in each action and on the relative importance of the externality for them. From Section 2.3, the planner recommends action 1 whenever

$$\int_0^\infty e^{-(\rho+\delta)t} [\theta_t - \nu^0 + 2\gamma n_t] dt > 0.$$

This expression for the planner is quite similar to the condition for an agent in (15). The only substantial difference is that γ is multiplied by 2, indicating that the planner gives more importance to externalities.

Limiting case. Assume $\mu, \sigma \rightarrow 0$. From Section 2.6, we know the indifference condition for the planner becomes

$$(1 - n_0) \int_0^\infty e^{-(\rho+\delta)t} [\theta^p - \nu^0 + 2\gamma n_t^\uparrow] dt + n_0 \int_0^\infty e^{-(\rho+\delta)t} [\theta^p - \nu^0 + 2\gamma n_t^l] dt = 0.$$

Solving for θ^P , we find that the planner plays according to the threshold:

$$\theta^P(n_0) = v^0 - \frac{2\gamma\delta}{\rho + 2\delta} - \frac{2\gamma\rho}{\rho + 2\delta}n_0. \quad (17)$$

3.1.3 | Symmetric network effects

We now assume $v \equiv v^0 = v^1$ so that externalities are the same for both networks. Hence increasing n implies a positive externality term whenever $n > 0.5$ and a negative term otherwise. Using $v^0 = \gamma/2$ in (17) and doing some algebra yields:

$$\theta^P(n_0) = -\frac{\gamma\delta}{\rho + 2\delta} + \frac{\gamma\rho}{2(\rho + 2\delta)} - \frac{2\gamma\rho}{\rho + 2\delta}n_0. \quad (18)$$

Guimaraes and Pereira (2016) use this model to describe consumers' choices between two competing standards, such as the QWERTY keyboard and the alternative (and allegedly better) option, the Dvorak keyboard.²⁹ The utility functions capture agents' preference for higher-quality products and for using the standard that is widespread. One can think that the QWERTY standard incidentally spread out before the higher-quality Dvorak standard became available. We have then a conflict between the two features of preferences: one product has the biggest network of consumers, but the other product is the highest-quality one. Is there room for interventions such as subsidies to eliminate inefficiencies in this environment? Comparing the decentralized equilibrium with the planner's solution sheds light on this matter.

Figure 11 shows how the decentralized equilibrium relates to the planner's solution. To the right of θ^* , agents are willing to pick action 1, and to the left of θ^* , action 0. The planner would mandate that action 1 is chosen whenever the economy is to the right of θ^P , instead. We can think of a vertical line crossing $(\theta^*(0.5), 0.5)$ as a dividing line between the regions where network 0 or network 1 are "intrinsically better." When $n = 0.5$, $\theta^* = -\gamma/2$, so along that vertical line $\theta^1 = \theta^0$.

The curve θ^P is half as steep as θ^* . The reason is that, in a decentralized economy, an agent at a revision opportunity considers the intrinsic quality of the two standards and the size of each network only to the extent that network effects affect her own payoff. The social planner dictating an agent's choice takes into account all factors that the agent considers plus the effect of that choice on other people. Hence, the planner will place higher weight on the coordination with the majority than the agent. Since the net externality from choosing action 1 is positive if (and only if) $n > 0.5$, the planner's threshold rotates around the agent's threshold so that if $n < 0.5$, the planner's threshold lies to the right of the agents', and if $n > 0.5$, it lies to the left.

As a consequence, the social planner is more conservative than agents regarding the transition to a better but smaller network. There exist quality differences such that the transition to the best standard (Dvorak) happens in the decentralized economy, but the social planner would choose to stay with the worst one (QWERTY). The shaded area represents such states where inefficient shifts to the intrinsically best network happen. If we observe that the transition to Dvorak has not happened in the decentralized economy, it must be that such transition would be inefficient (and thus should not be subsidized). The model predicts that whenever the prevailing standard is the worst-quality one, this is surely efficient.

²⁹This application fits well network goods that are incompatible, as in the case of firm-specific networks in Amir, Evstigneev, and Gama (2019).




FIGURE 11 Planner's solution with linear utilities

Common wisdom would tell us that the economy gets inefficiently stuck at low-quality networks. The results show this intuition is misleading and highlight the importance of the transition costs.

Angeli (2018) extends this model to consider present-biased agents. He shows that present bias induces agents to wait for larger increases in quality before switching to a better but smaller network. Hence, in this context, procrastination can be efficient. Intuitively, present-biased agents overvalue relative quality because transitioning is costly in the present and yields benefits only in the future.

3.1.4 | Asymmetric network effects

Now, consider the case with $\nu^1 \neq \nu^0$. The planner's threshold is given by (17). When the network effect is asymmetric, the planner not only rotates the threshold around $n = 0.5$, but it also shifts the threshold to enlarge the region in which agents choose the action that generates more externalities.

The next figure depicts the planner's solution when externalities are larger in network 0 than in network 1. When $\nu^0/\nu^1 > (\rho + \delta)/\delta$, the case of Figure 12a, the externality in network 0 is so large in comparison to the externality in 1 that the planner's threshold lies completely to the right of the agents' threshold.³⁰ The smaller the value of n , the larger the quality gap needed for the switch to network 1 to be efficient. This is because the planner takes into account that, when n is small, a lot of agents are stuck in action 0 (due to timing frictions) and they all would benefit from the network effects generated by an additional decrease in n .

Figure 12b depicts the case with $(\rho + \delta)/\delta \geq \nu^0/\nu^1 > 1$. Again, the region where the planner chooses network 0 is enlarged in comparison to the case with symmetric network effects, but in this case the planner's threshold crosses the agents' at some $n > 0.5$.

3.1.5 | Alternative interpretation: Negative externalities

This result is also applicable to a case where a certain action creates negative externalities, but those externalities end up generating strategic complementarities. Consider the following variation of the model: flow-utilities from states 0 and 1 are given by

$$u_t^0(\theta_t^0, n_t) = \theta_t^0 - \nu^0 n_t \quad \text{and} \quad u_t^1(\theta_t^1, n_t) = \theta_t^1 - \nu^1 n_t,$$

³⁰This is the case when $\theta^P(1) > \theta^*(1)$.




FIGURE 12 Asymmetric network effects (a) Large asymmetry; (b) Some asymmetry

with $v^0 > v^1 \geq 0$. The relative payoff function is given by

$$\Delta u(\theta_t, n_t) = \theta_t + \gamma n_t,$$

with $\theta_t \equiv \theta_t^1 - \theta_t^0$ and $\gamma \equiv v^0 - v^1 > 0$. These payoffs capture a situation in which action 1 harms everyone, but it is even more prejudicial to agents currently choosing action 0. As an example, imagine that agents can adopt an aggressive behavior (action 1) or an accommodating behavior (action 0). Action 1 is particularly harmful for those adopting an accommodating behavior. As an illustration, imagine that n is the proportion of drivers driving SUVs, and $1 - n$ is the proportion driving sedans and smaller cars. More SUVs means more severe car crashes, especially for sedan owners.

Solving the planner's problem in this case yields the same threshold as in (17), with the feature that $\theta^P(1) > \theta^*(1)$ as in Figure 12a. Agents do not internalize the harm imposed on others when they opt for an aggressive behavior. A larger number of agents choosing an accommodating behavior (a smaller n) implies a larger potential damage of an additional agent choosing to behave aggressively, and thus the larger the gap in fundamentals required for action 1 to be the right choice from a social perspective.

3.2 | Statistical discrimination

This section shows a simplified version of the model of statistical discrimination in Levin (2009). Individuals enter the economy at rate δ and decide whether or not to make an investment and become *skilled*. The cost of investment $C = -\theta$ is known at the time of the investment, where θ follows a Brownian motion.

Then, when a Poisson event occurs (again, rate δ), an individual is matched to an employer and receives her payoff. Employers observe (a) a noisy signal about the individual's skill; and (b) the fraction n_t of skilled individuals in the population. The signal about the individual's skill might be "good" or "bad." The signal is "good" with probability ϕ_H if the individual is skilled and ϕ_L if she is not, with $\phi_H > \phi_L$. We will later assume that skilled individuals always receive a good signal, that is, $\phi_H = 1$. An individual's payoff is equal to the probability attached to her being skilled. When the Poisson event hits, the individual receives her payoff and leaves the model.

The investment cost $-\theta$ can be negative. One interpretation is that θ is the expected benefit from investment received after the employment opportunity considered in the model net of the investment cost.

3.2.1 | Equilibrium

The probability an employer assigns to an individual being skilled is

$$\gamma_{Gt} = \frac{\phi_H n_t}{\phi_H n_t + \phi_L (1 - n_t)} \quad \text{and} \quad \gamma_{Bt} = \frac{(1 - \phi_H) n_t}{(1 - \phi_H) n_t + (1 - \phi_L)(1 - n_t)}$$

if the signal is "good" and if it is "bad," respectively.

The gain from investing is the increase in the probability of getting a good signal ($\phi_H - \phi_L$) times the present value of the gain from having a good signal ($\gamma_{Gt} - \gamma_{Bt}$). Hence an individual will decide to make the investment if

$$(\phi_H - \phi_L) \mathbb{E} \int_{\tau}^{\infty} e^{-(\rho+\delta)(t-\tau)} \delta (\gamma_{Gt} - \gamma_{Bt}) dt \geq -\theta_{\tau}.$$

Note that both γ_{Gt} and γ_{Bt} are increasing in n_t . Intuitively, if n_t is small, it is more likely that a good signal is due to luck and it is more likely that a bad signal reflects the individuals' true level of skill. However, the difference $\gamma_{Gt} - \gamma_{Bt}$ is not necessarily increasing in n_t .³¹ Hereafter assume that $\phi_H = 1$, which ensures that the payoff from investing is increasing in n_t because it implies $\gamma_{Bt} = 0$, meaning that individuals with a bad signal are surely unskilled. Individuals with a good signal might be either skilled or lucky (unskilled with a good signal) and the probability of the former is increasing in n_t . Hence agents' decisions are strategic complements. An individual will invest to become skilled if

$$\mathbb{E} \int_{\tau}^{\infty} e^{-(\rho+\delta)(t-\tau)} \delta \frac{n_t(1 - \phi_L)}{\phi_L + n_t(1 - \phi_L)} dt \geq -\theta_{\tau}. \quad (19)$$

Although the structure of this model is different from the framework presented in Section 2, mathematically this model can be seen as a particular case of that one. There is always a measure-one continuum of agents in the model, as agents enter and leave the economy at rate δ . Agents' decisions are given by (19), which is a particular case of (4) and the flow payoff is

³¹For example, suppose that $1 > \phi_H > \phi_L = 0$. Then $\gamma_{Gt} = 1$ and γ_{Bt} is increasing in n_t . Intuitively, an individual with a good signal must be skilled, but an individual with a bad signal might be either unskilled or unlucky, and the probability of the latter increases with n_t . In this case, $\gamma_{Gt} - \gamma_{Bt}$ is decreasing in n_t .



increasing in n_t . It is easy to write the flow payoff as increasing in θ_t as well and, finally, θ_t follows a Brownian motion.³²

Agents' decisions are thus determined by a unique downward sloping threshold θ^* . Levin (2009) highlights that for some values of θ , individuals would be investing if n_t were large, but may get stuck in a situation with low n_t , reflecting the persistence of statistical discrimination: employers are likely to believe that a good signal is likely to reflect luck, not skill, and individuals refrain from investing.

The model captures a situation where the rewards from investment depend on whether others believe the investment was indeed undertaken and this belief is affected by the choices of others in the population (or, in a particular group). Beyond labor markets, this model could be applied to other problems where statistical discrimination is an issue.

3.2.2 | Efficiency

Is the persistence of statistical discrimination inefficient? Should we be particularly concerned with economies with very low n_t but relatively high θ_t ?

Using the results derived in Section 2.3, we can ask how the decentralized equilibrium compares to the planner's solution. Consider a social planner whose flow benefit from investment is given by the sum of payoffs of agents that leave the game at a given point in time. This flow benefit is then given by:

$$W(n_t) = \delta(n_t\gamma_{Gt} + (1 - n_t)\phi_L\gamma_{Gt}) = \delta n_t.$$

Intuitively, at a given moment of time, the average payoff of agents will be n_t . If signals are perfectly informative ($\phi_L = 0$), skilled agents (that amount to a proportion n_t of individuals) will receive 1 and unskilled agents will receive 0. Less informative signals imply a lower payoff for skilled agents and a larger payoff for unskilled ones, but the average payoff is still the fraction of skilled agents, n_t . Hence, the planner faces no conflict between benefiting skilled or unskilled agents. An extra individual investing implies a gain equal to 1, that is distributed among the population.

Following the steps from Section 2.3, we get that the planner chooses to invest if

$$\mathbb{E} \int_{\tau}^{\infty} e^{-(\rho+\delta)(t-\tau)} \delta dt \geq -\theta_{\tau}.$$

Note that this is exactly what we would obtain from (19) in case $\phi_L = 0$. If $\phi_L = 0$ and $\phi_H = 1$, the signal about an individual is perfectly informative, so perceptions about an agent are not affected by the characteristics of the population (or group). Hence, an individual's action entails no externality.

The planner's threshold θ^P is given by

$$\theta^P = -\frac{\delta}{\delta + \rho}. \quad (20)$$

³²Specifically, (19) is equivalent to $\mathbb{E} \int_{\tau}^{\infty} e^{-(\rho+\delta)(t-\tau)} \left[(\rho + \delta)\theta_t + \delta \frac{n_t(1 - \phi_L)}{\phi_L + n_t(1 - \phi_L)} \right] dt \geq 0$.




FIGURE 13 Equilibrium in case of vanishing shocks

Investment is efficient as long as the expected present value of its benefit is larger than the cost $-\theta$. In the (θ, n) -space, the planner chooses according to a vertical threshold (θ^P is independent of n_t).

3.2.3 | Example

Using (13), we get the equilibrium threshold θ^* in the limit of vanishing shocks ($\sigma \rightarrow 0$):

$$\theta^*(n_0) = - \int_0^{n_0} \left(\frac{v}{n_0} \right)^{\frac{\rho}{\delta}} \frac{(1 - \phi_L)v}{\phi_L + v(1 - \phi_L)} dv - \int_{n_0}^1 \left(\frac{1 - v}{1 - n_0} \right)^{\frac{\rho}{\delta}} \frac{(1 - \phi_L)v}{\phi_L + v(1 - \phi_L)} dv.$$

Figure 13 shows an example with $\rho = 0.1$, $\delta = 0.2$. The solid vertical line at the left is the planner's threshold, from (20). The other curves are the equilibrium thresholds for $\phi_L = 0.1$, $\phi_L = 0.5$ and $\phi_L = 0.9$, starting from the left.

When $\phi_L = 0.9$, the private benefit from investing is small, as unskilled agents are likely to get a good signal. Hence, in equilibrium, investment pays off only if its cost is close to zero. Interestingly, even when the signal is very informative, there are important inefficiencies. The leftmost dashed curve corresponds to the case $\phi_L = 0.1$ (and, as in all other cases, $\phi_H = 1$). Still, the equilibrium threshold is quite far from the efficient benchmark, especially for low n . When $n = 0$, agents are willing to pay no more than 0.46 for investing, but it would be efficient to pay up to 0.67. When $n = 1$, agents are willing to pay up to 0.57.

In sum, the persistence of statistical discrimination is inefficient and we should indeed be particularly concerned with economies with very low n_t but relatively high θ_t .

3.3 | A macroeconomic model

Kiyotaki (1988) shows that in a model with monopolistic competition and some locally increasing returns to scale, there are multiple equilibria. Guimaraes and Machado (2018) embed a simple model that captures those insights in a dynamic framework with timing frictions.



The demand side is standard. A continuum of risk-neutral agents produce each a differentiated good that are aggregated by a competitive final good producer according to:

$$Y_t = \left(\int_0^1 y_{i,t}^{\frac{\varepsilon-1}{\varepsilon}} di \right)^{\frac{\varepsilon}{\varepsilon-1}}, \quad (21)$$

where Y_t is the amount produced of the final good, $y_{i,t}$ is the amount purchased of the variety produced by agent i , and $\varepsilon > 1$ is the elasticity of substitution. The zero-profit condition for final good producers is

$$\int_0^1 p_{i,t} y_{i,t} di = P_t Y_t, \quad (22)$$

where $p_{i,t}$ is the price of variety i and P_t is the price of the final good.

Intermediate goods producers can operate in two regimes: 1 (*High*) and 0 (*Low*). Chances to switch regimes arrive according to a Poisson process with arrival rate δ . Agents in the *High* regime can produce up to $e^{\theta} x_H$ and agents in the low regime can produce up to $e^{\theta} x_L$, with $x_H > x_L$. The marginal cost is zero up to firms' capacity, but firms must pay a fixed cost $\psi > 0$ to choose the *High* regime when they get the chance. One should think of moving to regime *High* as buying some machine that decreases the marginal cost of production. When a producer in regime *High* gets the Poisson realization, the machine fully depreciates and she gets a chance to get another one. When a producer in regime *Low* gets the Poisson realization, it simply means that an investment opportunity has arrived.

Alternatively, one could say that producers rent a machine but must pay a competitive user cost $\tilde{\psi} = (\rho + \delta)\psi$ that equals the expected depreciation ($\delta\psi$) plus interest ($\rho\psi$). Since agents are risk-neutral, as long as they can borrow in frictionless credit markets, they are indifferent between buying or renting and both interpretations are equivalent.

3.3.1 | Equilibrium

Say producers of the final good will spend C on intermediate goods. At every instant, they will choose inputs y_i to maximize (21) subject to

$$\int_0^1 p_{i,t} y_{i,t} di = C,$$

taking prices of intermediate goods as given. This is a static constrained-maximization problem. Combining the first-order conditions for intermediate goods i and j and rearranging yield

$$\frac{p_i}{p_j} = \left(\frac{y_i}{y_j} \right)^{-1/\varepsilon}.$$

Plugging this expression into the zero-profit condition in (22) yields the usual demand for each variety:

$$p_{i,t} = \left(\frac{y_{i,t}}{Y_t} \right)^{\frac{1}{\varepsilon}} P_t. \quad (23)$$

Using (23), and the fact that it will always be optimal to produce at maximum capacity, we can write the profit of producers in regime *High* and *Low* as

$$u_1 = p_{i,t} y_{i,t} - \tilde{\psi} = \left(e^{\theta_i x_H} \right)^{\frac{\varepsilon-1}{\varepsilon}} Y_t^{\frac{1}{\varepsilon}} P_t - \tilde{\psi} \quad (24)$$

and

$$u_0 = p_{i,t} y_{i,t} = \left(e^{\theta_i x_L} \right)^{\frac{\varepsilon-1}{\varepsilon}} Y_t^{\frac{1}{\varepsilon}} P_t. \quad (25)$$

By market clearing we have

$$Y_t = \left(n_t \left(e^{\theta_i x_H} \right)^{\frac{\varepsilon-1}{\varepsilon}} + (1 - n_t) \left(e^{\theta_i x_L} \right)^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{\varepsilon}{\varepsilon-1}}. \quad (26)$$

Plugging (26) into (24) and (25), and normalizing P_t to 1, we get:

$$u_1(\theta_i, n_t) = e^{\theta_i x_H^{\frac{\varepsilon-1}{\varepsilon}}} \left(n_t x_H^{\frac{\varepsilon-1}{\varepsilon}} + (1 - n_t) x_L^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{1}{\varepsilon-1}} - \tilde{\psi}$$

and

$$u_0(\theta_i, n_t) = e^{\theta_i x_L^{\frac{\varepsilon-1}{\varepsilon}}} \left(n_t x_H^{\frac{\varepsilon-1}{\varepsilon}} + (1 - n_t) x_L^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{1}{\varepsilon-1}},$$

which implies that

$$\Delta u(\theta_i, n_t) = e^{\theta_i} \left(n_t x_H^{\frac{\varepsilon-1}{\varepsilon}} + (1 - n_t) x_L^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{1}{\varepsilon-1}} \left(x_H^{\frac{\varepsilon-1}{\varepsilon}} - x_L^{\frac{\varepsilon-1}{\varepsilon}} \right) - \tilde{\psi}.$$

Notice that the expression above is increasing in θ_i and n_t . To guarantee the existence of dominance regions (Assumption 1), we need to impose that $\sigma^2 < 2(\rho + \delta)$.³³ Then, this model can be seen as a particular case of the framework presented in Section 2.

Firms have higher incentives to invest if they expect others to do the same in the future. This is so because of the demand externality: when others increase their production, the demand for one's variety increases, as shown in (23). Hence firms face a dynamic coordination problem in their investment decisions.

In what follows, it is useful to define $\widetilde{\Delta u}(\theta, n) \equiv \Delta u(\theta, n) + \tilde{\psi}$. Notice that $\widetilde{\Delta u}(\theta, n)$ is simply the instantaneous gain in revenue of operating in regime *High* instead of *Low*. The instantaneous losses from higher costs is $\tilde{\psi}$. Using (4), agents choose to invest as long as:

$$\int_{\tau}^{\infty} e^{-(\rho+\delta)(t-\tau)} \mathbb{E}_{\tau} [\widetilde{\Delta u}(\theta_t, n_t)] dt \geq \psi. \quad (27)$$

The equilibrium is given by the unique threshold θ^* that satisfies (27).

The dynamic coordination problem faced by firms may lead to dynamic inefficiencies that persist over time: firms may not invest today because they are not confident others will invest

³³This is because the integral $\int_0^{\infty} e^{-(\rho+\delta)t} \mathbb{E}[e^{\theta_t}] dt = \int_0^{\infty} e^{-(\rho+\delta-0.5\sigma^2)t} e^{\theta_0} dt$ diverge to ∞ if $\rho + \delta - 0.5\sigma^2 < 0$, making it always optimal to choose regime *High*.

tomorrow. But that contributes to lower investment tomorrow, since economic activity will be low. In other words, the economy may get in a dynamic coordination trap.

3.3.2 | Optimal stimulus policies

Some natural questions that emerge are the following: what kind of investment subsidies can implement the first best? Does the dynamic coordination problem imply that a social planner should provide counter-cyclical subsidies, that is, higher subsidies in times of low economic activity (low n_t)? We can use this model to answer these questions.

Using (6), (24), (25), and (26), the social planner instantaneous payoff is given by

$$W(\theta, n) = Y_t - \tilde{\psi} n_t,$$

which yields

$$\frac{dW(\theta, n)}{dn} = \left(\frac{\varepsilon}{\varepsilon - 1} \right) \widetilde{\Delta u}(\theta_t, n_t) - \tilde{\psi}. \quad (28)$$

As shown in Section 2.3, the social planner's decision threshold must satisfy (5). The planner thus chooses to invest as long as:

$$\int_{\tau}^{\infty} e^{-(\rho+\delta)(t-\tau)} \mathbb{E}_{\tau} \left[\left(\frac{\varepsilon}{\varepsilon - 1} \right) \widetilde{\Delta u}(\theta_t, n_t) - \tilde{\psi} \right] dt \geq 0. \quad (29)$$

Rearranging and using $\tilde{\psi} = (\rho + \delta)\psi$ yield

$$\int_{\tau}^{\infty} e^{-(\rho+\delta)(t-\tau)} \mathbb{E}_{\tau} [\widetilde{\Delta u}(\theta_t, n_t)] dt \geq \psi - \frac{1}{\varepsilon} \psi.$$

This is very similar to (27). It implies that the agents' problem is identical to the planner's problem if they get an investment subsidy of $\frac{1}{\varepsilon}\psi$. The planner can, therefore, implement the first best by providing a constant investment subsidy. Alternatively, it could pay $\frac{1}{\varepsilon-1}$ dollars for each unit of revenue firms get (Equation (29)). The important thing here is that the planner is not more willing to incentivize investment in times of low economic activity.

Now notice that we can write agents' payoffs as $\Delta u(\theta_t, n_t) = e^{\theta_t} g(n_t) - \tilde{\psi}$. Therefore, the planner's flow payoff can be written as:

$$\frac{dW(\theta, n)}{dn} = e^{\theta_t + \log\left(\frac{\varepsilon}{\varepsilon-1}\right)} g(n_t) - \tilde{\psi}.$$

Let $b_t = \theta_t + \log\left(\frac{\varepsilon}{\varepsilon-1}\right)$. Notice that b_t follows the same law of motion as θ_t . Thus the planner's problem is identical to the agent's problem if the fundamental is b_t instead of θ_t . Hence, the planner's threshold plotted in the (b_t, n_t) -space is identical to the agent's threshold plotted in the (θ_t, n_t) -space, which implies that the planner's threshold is a left translation of the agent's threshold, and the distance between the two is $\log\left(\frac{\varepsilon}{\varepsilon-1}\right)$. This result is shown in Figure 14.




FIGURE 14 Planner's translation

A similar reasoning implies that the distance between the planner's and the agents' threshold (for a given n) is proportional to the maximum subsidy the planner is willing to pay (which is the subsidy required to make agents invest on the planner's threshold). A subsidy of x units is equivalent to increasing θ in $\log(1 - x/\psi)$.

Agents always benefit if others invest (action 1), regardless of their current action. Using (9), (28), and $\Delta u(\theta, n) = \widetilde{\Delta u}(\theta, n) - \widehat{\psi}$, we get that the externality term is always positive (and proportional to $\widetilde{\Delta u}(\theta, n)$). The planner always has higher incentives to choose action 1 and thus chooses a threshold that lies entirely to the left of the agent's threshold.

But why does not the dynamic coordination problem require higher subsidies in times of low economic activity? To better understand the results it is useful to look at the planner's problem in a world without shocks and with multiple equilibria.

3.3.3 | Optimal stimulus with no shocks

The solution to the planner's problem when $\sigma \rightarrow 0$ is the same as the solution with $\sigma = 0$. But we know that the agent's problem is very different. Figure 15 combines the results depicted in Figures 8 and 14, showing the agents' threshold with $\sigma \rightarrow 0$ (θ^*), the two extreme equilibria of the game without shocks (θ^{opt} and θ^{pes}) and the planner's threshold.

When agents play according to either θ^{opt} or θ^{pes} , the distance between the planner's and the agents' threshold is larger when $n = 0$ than when $n = 1$. Therefore, regardless of whether we choose the “good” or “bad” equilibrium, the model with multiple equilibria predicts that the planner is more willing to pay subsidies in times of low economic activity, while the model with a unique equilibrium prescribes a constant subsidy.

Figure 16 illustrates the intuition for this result. First, imagine we are in a world where agents play according to the good equilibrium. In that case, it is harder to coordinate when $n = 0$ for an intuitive reason. In times of low economic activity but reasonably good




FIGURE 15 Planner's problem when $\sigma = 0$




FIGURE 16 Dynamic problem and beliefs problem $\sigma = 0$

fundamentals everyone would be happy to sign a contract that forces everyone to invest, even though the instantaneous gains of doing so today are negative. This is because higher investment will force a regime switch that will benefit everyone in the future. The problem is that no one wants to be the first to invest, leading to a dynamic coordination trap where one firm keeps waiting for the other to move. This is reflected by the negative association between n and the distance between the planner's threshold and the good equilibrium in Figure 16. We dub it the *dynamic problem*.

Now suppose the economy is at some point between θ^{opt} and θ^{pes} and for some reason agents are not investing. We know that in that region agents would invest if they were as optimistic as possible. Thus, whenever agents are not investing in the region between θ^{opt} and θ^{pes} we say that there is a *beliefs problem*. For instance, if agents play according to the unique equilibrium with vanishing shocks the size of the beliefs problem for a given n is the gray area between θ^{opt} and θ^* in Figure 16. Notice that θ^* touches the good equilibrium when $n = 0$ and the bad equilibrium when $n = 1$.³⁴

When agents play according to the unique equilibrium with vanishing shocks, the dynamic problem is higher when $n = 0$, but the beliefs problem is nonexistent. Conversely, when $n = 1$ there is no dynamic problem, but the beliefs problem is very severe. It turns out that the different size of the beliefs problem across different values n exactly cancels the dynamic problem. In the equilibrium with small shocks, agents are very optimistic at the threshold when $n = 0$ and very pessimistic when $n = 1$. Why?

Consider an agent deciding on the threshold θ^* when $n = 0$. She knows that if negative shocks hit, that will not change the level of economic activity. But if positive shocks hit, the economy will bifurcate up and leave the investment slump. Good shocks create a boom, while bad shocks do not worsen economic activity. Conversely, when $n = 1$ and the economy is at the threshold θ^* , good shocks do not increase economic activity, while bad shocks trigger an investment slump. The beliefs that arise in equilibrium offset the dynamic coordination trap, eliminating the need for higher subsidies at times of low economic activity.

3.4 | Housing and coordination

The basic framework that we have been discussing considers a continuum of agents with measure one. However, it also fits cases where there is a fixed supply of positions and two large groups of agents that bid for those positions.

³⁴Note that even when there is no beliefs problem (agents play according to θ^{opt}) and no dynamic problem ($n = 1$), there is still an inefficiency that mechanically arises from monopoly power, which is represented by the distance between $\theta_{planner}^*(1)$ and $\theta^{opt}(1)$.

In the model of Frankel and Pauzner (2002), there is a measure-one continuum of houses in a neighborhood and a larger measure of agents of two groups (they call them “blacks” and “whites,” we will call them 0 and 1). An agent from group $i \in \{0, 1\}$ living in the neighborhood gets a flow utility $u_i(\theta_t, n_t)$ where θ_t indicates the relative attractiveness of the neighborhood for people in group 1 (the attractiveness of other neighborhoods is normalized to zero) and n_t is the measure of houses owned by people from group 1. It is assumed that u_0 is increasing in $(1 - n_t)$ and u_1 is increasing in n_t .

People living in the neighborhood get moving opportunities at rate δ (say they have to go somewhere else). When this happens, they sell their houses to whoever offers a larger price (there is a large number of people from both groups willing to buy the house).

For an individual from group i at time τ , a house in the neighborhood is worth

$$P_\tau^i = \int_\tau^\infty e^{-(\rho+\delta)(t-\tau)} \mathbb{E}[u_i(\theta_t, n_t)] dt + \int_\tau^\infty \delta e^{-(\rho+\delta)(t-\tau)} \mathbb{E}[P_t(\theta_t, n_t)] dt.$$

The first term is the utility the agent will get while living in the house until she has to move. The second term is the expected price the agent gets when she sells the house, where $P_t(\theta_t, n_t)$ is the price she sells a house when the state is (θ_t, n_t) . Note that the second term does not depend on the type of agent (the value of the house does not depend on the type of the seller, she has to sell it anyway).

An individual from group 1 will outbid individuals from group 0 if and only if $P_\tau^1 > P_\tau^0$. That happens if:

$$\int_\tau^\infty e^{-(\rho+\delta)(t-\tau)} \mathbb{E}[\Delta u(\theta_t, n_t)] dt \geq 0,$$

where $\Delta u(\theta_t, n_t) = u_1(\theta_t, n_t) - u_0(\theta_t, n_t)$. This is equivalent to the expression in (1), and since Δu is increasing in θ_t and in n_t , all results from Section 2 apply here.

Frankel and Pauzner (2002) show that with vanishing shocks (or no shocks and a deterministic positive trend), an initially segregated neighborhood with $n = 0$ will experience a transition toward $n = 1$ at the first moment in which the transition is self-fulfilling: whenever outsiders would outbid insiders if a transition were to happen at the fastest possible rate, the transition starts.

4 | EXTENSIONS

Frankel and Pauzner (2000) assume that fundamentals follow a Brownian motion and allow for a deterministic drift. Frankel and Burdzy (2005) generalize the uniqueness result allowing for seasonal and mean-reverting shocks to θ . Daniëls (2009) extends the uniqueness result to the case where θ follows a jump diffusion process and derives an expression analogous to (14) for the case of vanishing frictions.

Guimaraes (2006) proposes an extension of this framework to the case of currency attacks. A currency pegged to the dollar might depreciate if enough agents take short positions. While everybody is long, the peg is more likely to be kept, which increases incentives for going long; but if everyone decides to short the currency, the peg is more likely to be abandoned, which raises incentives for agents to go short. At this point, agents would like to preempt others.



Despite the similarities, this is not a particular case of Frankel and Pauzner (2000), mainly because strategic complementarities are not global in the model.³⁵

This section considers two other extensions. First, it allows agents to choose the switching rate and shows that, as in Section 2.3, the planner's problem is a modified version of the decentralized equilibrium. Second, it considers a case with ex ante heterogeneous agents and shows that the main insights from the basic framework remain unchanged.

4.1 | Endogenous timing

So far agents were not allowed to decide on the timing of their opportunities to switch actions. But in many applications this can be an important margin of adjustment. For instance, a firm looking for a worker may decide to increase its search intensity if it faces higher demand; a firm building a new plant may want to speed up its construction if the economy is booming; a consumer unsatisfied with his current cell phone may want to try to sell it on secondary markets. The framework of Frankel and Burdzy (2005) allows agents to choose the arrival rate of their switching opportunities.

Now agents choose the *switching rate* instead of actions 0 and 1 directly. Agents locked in action 0 choose an arrival rate δ^0 of the Poisson process at every instant. Once the Poisson shock hits, the agent automatically switches to action 1 (she cannot choose if she wants to go). Similarly, agents locked in action 1 choose the switching rate δ^1 of the Poisson process that will send them to regime 0. At date t , an agent in regime 1 that chooses some arrival rate δ_t^1 must pay a cost $c^1(\delta_t^1)$. Therefore, the flow payoff of this agent is given by $u_1(\theta_t, n_t) - c^1(\delta_t^1)$. Similarly, agents in regime 0 bear a cost $c^0(\delta_t^0)$ and their flow payoff is $u_0(\theta_t, n_t) - c^0(\delta_t^0)$. The cost functions $c^j(\cdot)$ are weakly increasing and left-continuous in δ^j . Moreover, $\delta_t^j \in K_j$, where K_j is some closed interval in $\mathbb{R}_+$. The relative gain of being in regime high is

$$D(\theta, n, \delta^0, \delta^1) = \Delta u(\theta, n) - [c^1(\delta^1) - c^0(\delta^0)].$$

It is assumed that $D(\cdot)$ is weakly increasing and Lipschitz in θ and n .³⁶ Moreover, the assumption below is imposed (it is simply the counterpart of Assumption 1).

Assumption 2. (Existence of dominance regions) There exists $\tilde{\theta}$ and $\underline{\theta}$ such that: if $\theta_t > \tilde{\theta}$, it is strictly dominant for players in regime 0 (1) to switch at the maximal (minimal) rate; if $\theta_t < \underline{\theta}$ it is strictly dominant for players in regime 0 (1) to switch at the minimal (maximal) rate.

This framework can easily accommodate the baseline model. We can set $K_0 = K_1 = [0, \delta]$ and $c^0(\delta^0) = c^1(\delta^1) = 0$, for every n , δ^0 , and δ^1 . Then, choosing action j with probability p conditional on getting the chance in the baseline model, is equivalent to choosing an arrival rate $\delta_t^j = (1 - p)\delta$ and $\delta_t^{1-j} = p\delta$ in this model.

Let $\Delta \mathcal{V}_\tau$ denote the lifetime utility of being in regime 1 minus the lifetime utility of being in regime 0, at date τ . This is given by

³⁵In the global games literature, there are several results for equilibrium uniqueness when strategic complementarities are not global. There is room for research on this issue in the Frankel and Pauzner (2000) environment.

³⁶Under some additional assumptions, the cost $c(\cdot)$ can also depend on n_t . See Frankel and Burdzy (2005).

$$\Delta \mathcal{V}_\tau = \mathbb{E}_\tau \left[\int_\tau^\infty e^{-\int_\tau^t (\rho + \delta_t^0 + \delta_t^1) dt} D(\theta_t, n_t, \delta_t^0, \delta_t^1) dt \right].$$

Note that δ_t^0 and δ_t^1 appear on the discount factor. This is because the difference in utility between agents in regime 1 and 0 becomes zero whenever both agents are in the same regime. It happens either when an agent in regime 1 switches to regime 0; or when an agent in regime 0 switches to regime 1. The probability that none of those events have happened at some date $t > \tau$ is $\exp\{-\int_\tau^t (\delta_t^0 + \delta_t^1) dt\}$.

As shown in Frankel and Burdzy (2005), there is a unique equilibrium in this game. At each date τ , agents in regime 1 choose a switching rate $\delta^1 = \delta^{1*}(n_t, \theta_t)$ to maximize

$$\delta^1(-\Delta \mathcal{V}_\tau) - c^1(\delta^1).$$

Similarly, agents in regime 0 choose a switching rate $\delta^0 = \delta^{0*}(n_t, \theta_t)$ to maximize

$$\delta^0 \Delta \mathcal{V}_\tau - c^0(\delta^0). \quad (30)$$

If the solution is interior, these expressions are simply stating that agents equalize the marginal cost of increasing the probability of switching with the expected gain. In equilibrium, agents take as given their own switching rate in the future and the switching rate of others ($\Delta \mathcal{V}$ is taken as given).

4.1.1 | The planner's problem

We add the assumption that costs $c^j(\cdot)$ are convex. This implies that the planner chooses the same hazard rate for every agent. For simplicity, we assume a constant hazard rate for agents locked in action 1, equal to δ^1 , with $c^1(\delta^1) = 0$. Hence welfare in this economy is given by

$$\mathbb{E}_\tau \int_\tau^\infty e^{-\rho(t-\tau)} W(\theta_t, n_t, \delta_t^0) dt,$$

with

$$W(\theta_t, n_t, \delta_t^0) = n_t u_1(\theta_t, n_t) + (1 - n_t) \left[u_0(\theta_t, n_t) - c^0(\delta_t^0) \right].$$

The strategy here follows Guimaraes and Machado (2018) and is similar to the one employed in Section 2.3. Suppose the planner is following the optimal plan and consider the following deviation: change δ_τ^0 to $\tilde{\delta}$ at time 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 $1 - n_\tau$ agents locked in action 0, costs change by

$$\left[c^0(\tilde{\delta}) - c^0(\delta_\tau^0) \right] (1 - n_\tau) dt \quad (31)$$

and the immediate effect on n_τ is

$$dn_\tau = (\tilde{\delta} - \delta_\tau^0)(1 - n_\tau)dt. \quad (32)$$

This effect dies out in time:

$$dn_s = dn_\tau - \int_\tau^s dn_v (\delta_v^0 + \delta^1) dv,$$

which implies that

$$dn_s = dn_\tau e^{-\int_\tau^s (\delta_v^0 + \delta^1) dv}. \quad (33)$$

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

$$\mathbb{E}_\tau \int_\tau^\infty e^{-\rho(s-\tau)} \left(\frac{\partial W(\theta_s, n_s, \delta_s^0)}{\partial n_s} dn_s \right) ds. \quad (34)$$

This deviation cannot be profitable. Hence, putting together (31), (32), (33), and (34), it must be that

$$(\tilde{\delta} - \delta_\tau^0) \mathbb{E}_\tau \int_\tau^\infty e^{-\int_\tau^s (\rho + \delta_v^0 + \delta^1) dv} \left(\frac{\partial W(\theta_s, n_s, \delta_s^0)}{\partial n_s} \right) ds - [c^0(\tilde{\delta}) - c^0(\delta_\tau^0)] \leq 0$$

for any $\tilde{\delta} \in [0, \bar{\delta}]$. That is equivalent to stating that a necessary condition for the planner's solution is that δ_τ^0 must maximize

$$\delta_\tau^0 \mathbb{E}_\tau \int_\tau^\infty e^{-\int_\tau^s (\rho + \delta_v^0 + \delta^1) dv} \left(\frac{\partial W(\theta_s, n_s, \delta_s^0)}{\partial n_s} \right) ds - c^0(\delta_\tau^0). \quad (35)$$

This expression is similar to (30). The difference is the term $\partial W(\theta_s, n_s, \delta_s^0)/\partial n_s$ inside the integral instead of $D(\theta_s, n_s, \delta_s^0)$. As in (9), the planner considers the externality on other agents. Mathematically, finding the solution to the planner's problem is thus equivalent to finding a solution to a game played by agents that maximize (35).

4.1.2 | Examples

Let's say agents are firms and each firm has a job that can be either filled (locked in action 1) or vacant (locked in action 0). Firms with a vacant job choose the search intensity δ^0 and matches are destroyed at an exogenous rate δ^1 , with $c^1(\delta^1) = 0$ (so K_1 is a singleton). A vacant job yields a flow payoff $u_0(\theta_t, n_t) = 0$ and a filled job generates a flow payoff $u_1(\theta_t, n_t)$ that is increasing in both arguments— θ_t can be seen as a measure of productivity

and it is assumed that firms' profits are larger when output in the economy (proportional to n_t) is larger.

This is the essence of the model of Howitt and McAfee (1992). They show the existence of multiple equilibria in this setting with no shocks to θ . If θ follows a Brownian motion, we can apply the results in Frankel and Burdzy (2005), so the model has a unique equilibrium.

The model in Section 3.3 can also be extended to allow for endogenous hazard rates. Firms can be in two regimes: 0 (low capacity) and 1 (high capacity). Firms in regime 1 have some capital that fully depreciate at some exogenous rate. Once it depreciates, firms can decide how much resources (if any) they will spend building a fixed unit of capital (say, constructing a new plant). In other words, they choose the arrival rate of the new plant. The more resources spent, the more likely the new plant will be ready to operate. Guimaraes and Machado (2018) show that the constant-subsidy result also applies to this setting.

4.2 | Ex ante heterogeneous agents

A consumer deciding between Facebook and another social network takes into account what others have been choosing but also her own tastes; a firm's investment decisions might depend on other firms' investment (which affects aggregate demand), but also on idiosyncratic factors that affect the demand for its particular product. In these and other settings, both strategic complementarities and idiosyncratic features of preferences or technologies are important. While strategic complementarities lead agents to try to do what others are expected to do, idiosyncratic components of preferences might push agents in different directions. Guimaraes and Pereira (2017) study the interplay of complementarities and heterogeneity in payoffs in a dynamic setting with timing frictions as in Frankel and Pauzner (2000).

Denote agent i 's relative flow-payoff of choosing action 1 by $\Delta u_{q(i)}(\theta, n)$, where $\theta \in \mathbb{R}$ denotes the fundamentals of the economy, $n \equiv \int_0^1 a_i di$ is the fraction of agents currently committed to action 1, as before, and $q(i) \in \mathcal{Q} = \{1, \dots, Q\}$ is agent i 's type. All functions $\Delta u_q(\cdot)$ are continuously differentiable and strictly increasing in both arguments. If we let α_q denote the mass of type- q agents in the population and n_q the proportion of type- q agents currently playing 1, n can be written as $n = \sum_{q=1}^Q \alpha_q n_q$.

An agent who receives a revision opportunity at time τ will choose $a_i = 1$ whenever

$$\mathbb{E} \int_{\tau}^{\infty} e^{-(\rho+\delta)(t-\tau)} \Delta u_{q(i)}(\theta_t, n_t) dt > 0$$

and $a_i = 0$ if the inequality is reversed.

Assumption 1 is now replaced by the assumption that all payoff functions $\Delta u_q(\cdot)$ are such that there are dominance regions for all types of agents. That is, for each $q \in \mathcal{Q}$, there are values $\tilde{\theta}_q$ and $\underline{\theta}_q$ such that: if $\theta_t > \tilde{\theta}_q$, playing 1 is a dominant action, and if $\theta_t < \underline{\theta}_q$, playing 0 is a dominant action for type- q agents.

In principle, one could expect the dynamics of the system to depend on the proportion of each type of agent currently choosing each option. However, owing to the assumption of a Poisson process for the arrival rate of switching opportunities, we can deal with this problem in

a two-dimensional space. Guimaraes and Pereira (2017) show that for any given strategy profile, the dynamics of n_t (i.e., $\partial n_t / \partial t$) depends on each $n_{q,t}$ (for $q \in \mathcal{Q}$) only through n_t . Denote by $\mathcal{I}_t$ the set of types whose strategies prescribe action 1 at time t . We have that

$$\frac{\partial n_t}{\partial t} = \delta \left[\sum_{q \in \mathcal{I}_t} \alpha_q - n_t \right]. \quad (36)$$

Guimaraes and Pereira (2017) then show that the uniqueness result in Theorem 1 extends to this environment. Agents choose action 1 when $\theta_t > \theta_q^*(n_t)$ and action 0 when $\theta_t < \theta_q^*(n_t)$, where $\theta_q^*(\cdot)$ for $q \in \mathcal{Q}$ are decreasing functions. The argument employs a strategy of iterative elimination of strictly dominated strategies similar to the proof of Theorem 1.

4.2.1 | Bifurcation probabilities

From now on, assume there are only two types of agents, $\mathcal{Q} = \{q, \bar{q}\}$. Suppose there is a mass α of type- $\bar{q}$ agents, and assume that the utility functions are such that for any pair (θ, n) we have $\Delta u_{\bar{q}}(\theta, n) > \Delta u_q(\theta, n)$, that is type- $\bar{q}$ agents have a higher relative instantaneous payoff of choosing action 1 in every state.

The result in Burdzy et al. (1998) essentially states that as $\mu, \sigma \rightarrow 0$, the probabilities of the economy bifurcating up and down are proportional to the speed that the economy moves in either direction around the threshold, as in (10). Hence we can extend Burdzy, Frankel and Pauzner's (1998) result on bifurcation probabilities to the case with heterogeneous agents. Suppose agents play according to two arbitrary (downward sloping and Lipschitz) thresholds $\theta_{\bar{q}}(n) < \theta_q(n)$ for all n in some interval (n^1, n^2) . Consider a point along one of those thresholds with $n \in (n^1, n^2)$. As $\mu, \sigma \rightarrow 0$, the time it takes for the system to bifurcate either up or down converges to zero. The probabilities of an upward bifurcation along the thresholds are computed as follows (the probabilities of downward bifurcation are the complements):

1. For a point (θ_t, n_t) with $\theta_t = \theta_{\bar{q}}(n_t)$,

$$\text{Prob}(\text{bifurcate up}) = \begin{cases} 0, & \text{if } n \geq \alpha, \\ 1 - \frac{n}{\alpha}, & \text{if } n < \alpha. \end{cases}$$

2. For a point (θ_t, n_t) with $\theta_t = \theta_q(n_t)$.

$$\text{Prob}(\text{bifurcate up}) = \begin{cases} \frac{1-n}{1-\alpha}, & \text{if } n > \alpha, \\ 1, & \text{if } n \leq \alpha. \end{cases}$$

Figure 17 illustrates the dynamics around the two types' thresholds in case they do not intersect (see Equation (36)) and the implied bifurcation probabilities along the thresholds. As in the case with a single threshold, starting somewhere along the curves $\theta_{\bar{q}}$ or θ_q , once the economy has headed off in one direction, it does not revert to the threshold,




FIGURE 17 Bifurcation probabilities with two thresholds

since it is downward sloping and shocks to fundamentals are small. The direction of bifurcation depends on the realization of the Brownian motion in a tiny time span and on the speed at which n increases or decreases at each side of the threshold.

To fix ideas, consider a point $(\theta_{\bar{q}}(n), n)$ with $n > \alpha$. Small negative shocks bring the economy to the left of $\theta_{\bar{q}}$, a region in which n decreases. A tiny shock to the right also puts the economy in a region where n decreases—at a lower speed, but still.³⁷ Hence, regardless of the direction of the (small) shocks that hit, the economy immediately heads off in the direction of $n = 0$: the system bifurcates down with probability one. Now, suppose $n < \alpha$. Negative shocks that bring the economy to the left of $\theta_{\bar{q}}$ make n decrease at maximum speed, while positive shocks that lead the economy to the right of the threshold make n increase at rate $\delta(n - \alpha)$.

4.2.2 | Results for limiting cases

Once again we focus on the tractable limiting cases with vanishing shocks to fundamentals and vanishing timing frictions. Let $\theta_{\bar{q}}^*(n)$ and $\theta_{\underline{q}}^*(n)$ denote the two types' equilibrium thresholds when $\mu, \sigma \rightarrow 0$. The equilibrium properties in this case depend on the degree of payoff heterogeneity.

Denote by $\theta_{\bar{q}}^{pes}$ the boundary of the upper dominance region of a type- q agent (i.e., the curve along which such agent is indifferent between the two actions if she believes everyone after her will choose 0), and by $\theta_{\underline{q}}^{opt}$ the boundary of the lower dominance region for a type- q player. If heterogeneity is not too large, dominance regions are such that $\theta_{\underline{q}}^{opt}(n) < \theta_{\bar{q}}^{pes}(n)$ for all n , so there is a region in which, for both types, neither action is dominant. Figure 18 illustrates this case.

Proposition 2 shows the main equilibrium properties for the case of vanishing shocks when heterogeneity is not so large.³⁸

Proposition 2. *In the setting with two types of agents and vanishing shocks, in the unique rationalizable equilibrium, if $\theta_{\underline{q}}^{opt}(n) < \theta_{\bar{q}}^{pes}(n)$ for all n , then $\theta_{\bar{q}}^*(n) = \theta_{\bar{q}}^{pes}(n)$ for all n in an interval containing α . Moreover, there are neighborhoods around 0 and 1 in which $\theta_{\bar{q}}^*(n) < \theta_{\underline{q}}^*(n)$.*

³⁷Between the thresholds, all high-type agents who get the chance will play 1, but all low-types will play 0. When $n > \alpha$, there are more agents currently committed to action 1 than agents willing to choose 1, so n decreases at a rate $\delta(n - \alpha)$.

³⁸See Guimaraes and Pereira (2017) for the proof.




FIGURE 18 The case of not-so-large heterogeneity

Proposition 2 states that there is some conformity in agents' behavior as long as heterogeneity is not large enough to make it impossible for agents to play according to the same threshold for any (arbitrary) set of beliefs. Moreover, different types choose according to the same threshold for an intermediate range of n , but for extreme values of n , heterogeneity in preferences beats coordination motives and each type has a distinct threshold. Figure 19 illustrates this result.

To better understand the result, suppose agents play according to distinct thresholds as in Figure 17. For $n = \alpha$, an agent at the leftmost threshold holds the most pessimistic beliefs, that n will surely decrease from then on at maximum speed. Hence an agent will not choose action 1 unless it is dominant to do so. Conversely, an agent at the rightmost threshold holds the exact opposite beliefs, and thus will not choose 0 unless it is dominant to do so. This reasoning implies that an equilibrium with two distinct thresholds at $n = \alpha$ would require (a) high-type agents being indifferent between either choice for some $\tilde{\theta}$ under the most pessimistic beliefs; and (b) low-type agents being indifferent between either choice for some $\theta > \tilde{\theta}$ under the most optimistic beliefs, which cannot happen in the case of low heterogeneity. Owing to the large dispersion in beliefs offsetting idiosyncratic payoffs, both thresholds will coincide at $n = \alpha$.

An analogous reasoning helps understand why for extreme values of n thresholds for distinct types do not coincide. Consider $n = 0$. Notice beliefs at the two thresholds are not so different: both types expect the economy to bifurcate up with probability one. The speed n will move up is not the same in both cases, but that is not such a significant difference in beliefs. That is why even a small difference in preferences leads to the existence of two distinct thresholds.

Consider now the limiting case of $\delta \rightarrow \infty$. The next proposition emphasizes some properties of the equilibrium when distinct types' flow payoffs differ by a constant.³⁹

Proposition 3. Let $\Delta u_{\bar{q}}(\theta, n) = \Delta u(\theta, n) + \bar{\varepsilon}$ and $\Delta u_{\underline{q}}(\theta, n) = \Delta u(\theta, n) + \underline{\varepsilon}$, where $\Delta u(\cdot)$ is continuously differentiable and strictly increasing in both arguments and $\bar{\varepsilon} > \underline{\varepsilon}$. Define $\hat{\varepsilon} \equiv \alpha \bar{\varepsilon} + (1 - \alpha) \underline{\varepsilon}$ and $\hat{\theta}$ as satisfying $\int_0^1 \Delta u(\hat{\theta}, n) dn = -\hat{\varepsilon}$. In the limit as $\delta \rightarrow \infty$, the vertical line $\hat{\theta}$ divides the state space in two regions: whenever $\theta_i > \hat{\theta}$, $n_i \approx 1$ and whenever $\theta_i < \hat{\theta}$, $n_i \approx 0$.

Proposition 3 states that if heterogeneity is not so large, the economy under vanishing frictions behaves as if agents were identical and had an intermediate preference parameter $\hat{\varepsilon}$. Although agents' strategies do not coincide, agents of a given type immediately switch actions when fundamentals cross the vertical division line, and the change in n leads the opposite type

³⁹See Guimaraes and Pereira (2017) for the proof. For general payoff functions, one can use the bifurcation probabilities to compute the equilibrium analytically.




FIGURE 19 Equilibrium with small heterogeneity and $\sigma \rightarrow 0$




FIGURE 20 Equilibrium when $\delta \rightarrow \infty$

to switch actions as well. Since revision opportunities arrive at a very fast pace, the dynamics of the economy is basically the same as if agents were identical. Figure 20 illustrates the result.

The planner's problem when agents are ex ante heterogeneous can be solved in an analogous way to the one presented in Section 2.3. Guimaraes and Pereira (2017) show that the region of the state space where the planner would choose the same strategy for different types is always larger than the region where agents play the same strategy in the decentralized equilibrium. Therefore, from an efficiency perspective, there is not enough conformity. Agents do not internalize the spillovers from their action, and thus do not put enough weight on coordination relative to their own idiosyncratic tastes.

5 | FINAL REMARKS AND DIRECTIONS FOR FUTURE RESEARCH

The framework studied in this paper fits a wide range of economic problems, but many potential applications have not yet been explored. The tools presented here might be useful for applied theorists hoping to contribute to this literature. In particular, the comparison between the social planner's solution and the decentralized equilibrium can shed light on the nature of inefficiencies in dynamic coordination problems.

There is also room for theorists to expand the lack of possible applications. The literature on perfect foresight dynamics has generalized the results in Matsui and Matsuyama (1995) for games with more than two actions. The same is true for the global games literature.⁴⁰ However, there are no equivalent results for the Frankel and Pauzner (2000) framework. That would be desirable in many applications. For example, in a network formation game, allowing for

⁴⁰Applied research has employed global-game models with a continuous set of actions for a variety of applications, including how risk considerations driving portfolio choices affect currency attacks (see Guimaraes & Morris, 2007) and how revolutionary leaders inspire participation in political regime change (see Morris & Shadmehr, 2017).

intermediate level of commitment might affect results;⁴¹ in macroeconomics, incorporating coordination problems into real business cycle models would require allowing firms to choose investment not only in the extensive margin, but also in the intensive margin.

ACKNOWLEDGMENTS

We thank the editor Rabah Amir, the associate editor, an anonymous referee and the students of courses on Dynamic Coordination at the Sao Paulo School of Economics—FGV for their helpful comments and suggestions. Guimaraes gratefully acknowledges financial support from CNPq. Machado and Pereira gratefully acknowledge financial support from CONICYT through FONDECYT Iniciación grants #11180046 and #11190202, respectively.

DATA AVAILABILITY STATEMENT

Data sharing is not applicable to this article as no new data were created or analyzed in this study.

ORCID

Bernardo Guimaraes  <http://orcid.org/0000-0003-0098-2174>

Caio Machado  <http://orcid.org/0000-0002-5545-6270>

Ana E. Pereira  <http://orcid.org/0000-0002-8271-6699>

REFERENCES

- Amir, R., Evstigneev, I., & Gama, A. (2019). Oligopoly with network effects: Firm-specific versus single network. *Economic Theory*, <https://doi.org/10.1007/s00199-019-01229-0>
- Amir, R., Gabszewicz, J., & Resende, J. (2014). Thematic clubs and the supremacy of network externalities. *Journal of Public Economic Theory*, 16, 706–729.
- Amir, R., & Lazzati, N. (2011). Network effects, market structure and industry performance. *Journal of Economic Theory*, 146, 2389–2419.
- Andreoni, J. (1998). Toward a theory of charitable fund-raising. *Journal of Political Economy*, 106, 1186–1213.
- Angeletos, G.-M., Hellwig, C., & Pavan, A. (2007). Dynamic global games of regime change: Learning, multiplicity, and the timing of attacks. *Econometrica*, 75, 711–756.
- Angeletos, G.-M., & Lian, C. (2016). Incomplete information in macroeconomics: Accommodating frictions in coordination. *Handbook of Macroeconomics*, 2, 1065–1240.
- Angeli, D. (2018). *Dynamic coordination with network externalities: When procrastination can be efficient* (Working paper).
- Bethencourt, C., & Kunze, L. (2019). Tax evasion, social norms, and economic growth. *Journal of Public Economic Theory*, 21, 332–346.
- Burdzy, K., Frankel, D., & Pauzner, A. (1998). On the time and direction of stochastic bifurcation. In B. Szyszkowicz (Ed.), *Asymptotic methods in probability and statistics* (pp. 483–500). Amsterdam: North-Holland.
- Burdzy, K., Frankel, D., & Pauzner, A. (2001). Fast equilibrium selection by rational players living in a changing world. *Econometrica*, 69, 163–189.
- Carlsson, H., & Damme, E. V. (1993). Global games and equilibrium selection. *Econometrica*, 61, 989–1018.

⁴¹See the model in Dev (2018).

- Caselli, F., & Morelli, M. (2004). Bad politicians. *Journal of Public Economics*, 88, 759–782.
- Chassang, S. (2010). Fear of miscoordination and the robustness of cooperation in dynamic global games with exit. *Econometrica*, 78, 973–1006.
- Cooper, R., & John, A. (1988). Coordinating coordination failures in Keynesian models. *Quarterly Journal of Economics*, 103, 441–463.
- Crouzet, N., Gupta, A., & Mezzanotti, F. (2018). *Shocks and technology adoption: Evidence from electronic payment systems* (Working Paper).
- Daniëls, T. R. (2009). Unique equilibrium in a dynamic model of speculative attacks. *De Economist*, 157, 417–439.
- Dasgupta, A. (2007). Coordination and delay in global games. *Journal of Economic Theory*, 134, 195–225.
- Dasgupta, A., Steiner, J., & Stewart, C. (2012). Dynamic coordination with individual learning. *Games and Economic Behavior*, 74, 83–101.
- David, P. A. (1985). Clio and the economics of QWERTY. *American Economic Review*, 75, 332–337.
- Dev, P. (2018). Group identity in a network formation game with cost sharing. *Journal of Public Economic Theory*, 20, 390–415.
- Diamond, D. W., & Dybvig, P. H. (1983). Bank runs, deposit insurance, and liquidity. *Journal of Political Economy*, 91, 401–419.
- Ennis, H. M., & Keister, T. (2005). Government policy and the probability of coordination failures. *European Economic Review*, 49, 939–973.
- Frankel, D., & Burdzy, K. (2005). Shocks and business cycles. *Advances in Theoretical Economics*, 5, 1–86.
- Frankel, D., Morris, S., & Pauzner, A. (2003). Equilibrium selection in global games with strategic complementarities. *Journal of Economic Theory*, 108, 1–44.
- Frankel, D., & Pauzner, A. (2000). Resolving indeterminacy in dynamic settings: The role of shocks. *Quarterly Journal of Economics*, 115, 285–304.
- Frankel, D., & Pauzner, A. (2002). Expectations and the timing of neighborhood change. *Journal of Urban Economics*, 51, 295–314.
- Fudenberg, D., & Harris, C. (1992). Evolutionary dynamics with aggregate shocks. *Journal of Economic Theory*, 57, 420–441.
- Guimaraes, B. (2006). Dynamics of currency crises with asset market frictions. *Journal of International Economics*, 68, 141–158.
- Guimaraes, B., & Machado, C. (2018). Dynamic coordination and the optimal stimulus policies. *Economic Journal*, 128, 2785–2811.
- Guimaraes, B., & Morris, S. (2007). Risk and wealth in a model of self-fulfilling currency attacks. *Journal of Monetary Economics*, 54, 2205–2230.
- Guimaraes, B., & Pereira, A.E. (2016). QWERTY is efficient. *Journal of Economic Theory*, 163, 819–825.
- Guimaraes, B., & Pereira, A. E. (2017). Dynamic coordination among heterogeneous agents. *Journal of Mathematical Economics*, 73, 13–33.
- Hofbauer, J., & Sorger, G. (1999). Perfect foresight and equilibrium selection in symmetric potential games. *Journal of Economic Theory*, 85, 1–23.
- Howitt, P., & McAfee, R. (1992). Animal spirits. *American Economic Review*, 82, 493–507.
- Ilzetzki, E. (2018). Tax reform and the political economy of the tax base. *Journal of Public Economics*, 164, 197–210.
- Kandori, M., George, J. M., & Rob, R. (1993). Learning, mutation, and long run equilibria in games. *Econometrica*, 61, 29–56.
- Katz, M. L., & Shapiro, C. (1985). Network externalities, competition, and compatibility. *American Economic Review*, 75, 424–440.
- Kim, Y. (2003). Income distribution and equilibrium multiplicity in a stigma-based model of tax evasion. *Journal of Public Economics*, 87, 1591–1616.
- Kiyotaki, N. (1988). Multiple expectational equilibria under monopolistic competition. *Quarterly Journal of Economics*, 103, 695–713.
- Kováč, E., & Steiner, J. (2013). Reversibility in dynamic coordination problems. *Games and Economic Behavior*, 77, 298–320.
- Levin, J. (2009). The dynamics of collective reputation. *The BE Journal of Theoretical Economics*, 9, 27.



- Mathevet, L., & Steiner, J. (2013). Tractable dynamic global games and applications. *Journal of Economic Theory*, 148, 2583–2619.
- Matsui, A., & Matsuyama, K. (1995). An approach to equilibrium selection. *Journal of Economic Theory*, 65, 415–434.
- Matsuyama, K. (1991). Increasing returns, industrialization, and indeterminacy of equilibrium. *Quarterly Journal of Economics*, 106, 617–650.
- Meunier, G., & Schumacher, I. (2019). The importance of considering optimal government policy when social norms matter for the private provision of public goods. *Journal of Public Economic Theory*, <https://doi.org/10.1111/jpet.12418>
- Morris, S. (2014). Coordination, timing and common knowledge. *Research in Economics*, 68, 306–314.
- Morris, S., & Shadmehr, M. (2017). *Inspiring regime change* (Working Paper).
- Morris, S., & Shin, H. S. (1998). Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review*, 587–597.
- Morris, S., & Shin, H. S. (2003). Global games: Theory and applications. In M. Dewatripont, L. Hansen & S. Turnovsky (Eds.), *Advances in economics and econometrics: Theory and applications, Eighth World Congress*. Cambridge: Cambridge University Press.
- Nyborg, K., & Telle, K. (2004). The role of warnings in regulation: Keeping control with less punishment. *Journal of Public Economics*, 88, 2801–2816.
- Oyama, D., Takahashi, S., & Hofbauer, J. (2008). Monotone methods for equilibrium selection under perfect foresight dynamics. *Theoretical Economics*, 3, 155–192.
- Oyama, D., & Tercieux, O. (2009). Iterated potential and robustness of equilibria. *Journal of Economic Theory*, 144, 1726–1769.
- Palivos, T., & Varvarigos, D. (2013). Intergenerational complementarities in education, endogenous public policy, and the relation between growth and volatility. *Journal of Public Economic Theory*, 15, 249–272.
- Plantin, G., & Shin, H. S. (2018). Exchange rates and monetary spillovers. *Theoretical Economics*, 13, 637–666.
- Sanchez-Villalba, M. (2015). Global inspection games. *Journal of Public Economics*, 128, 59–72.
- Steiner, J. (2008). Coordination cycles. *Games and Economic Behavior*, 63, 308–327.

How to cite this article: Guimaraes B, Machado C, Pereira AE. Dynamic coordination with timing frictions: Theory and applications. *J Public Econ Theory*. 2020;22:656–697. <https://doi.org/10.1111/jpet.12427>