FUNDAÇÃO GETULIO VARGAS
ESCOLA DE ECONOMIA DE SÃO PAULO

GABRIEL DOLHNIKOFF JARDANOVSKI

# WHO MATTERS IN DYNAMIC COORDINATION PROBLEMS?

São Paulo

2020

GABRIEL DOLHNIKOFF JARDANOVSKI

# WHO MATTERS IN DYNAMIC COORDINATION PROBLEMS?

Dissertação apresentada à Escola de Economia de São Paulo como pré-requisito à obtenção de título de mestre em Economia de Empresas.

Orientador: Bernardo de Vasconcellos Guimarães.

.

São Paulo

2020

Jardanovski, Gabriel Dolhnikoff.

Who matters in dynamic coordination problems? / Gabriel Dolhnikoff Jardanovski. - 2020.

35 f.

Orientador: Bernardo de Vasconcellos Guimarães.

Dissertação (mestrado CMEE) – Fundação Getulio Vargas, Escola de Economia de São Paulo.

1. Economia. 2. Processo decisório. 3. Equilíbrio econômico. 4. Teoria dos jogos. I. Guimarães, Bernardo de Vasconcellos. II. Dissertação (mestrado CMEE) – Escola de Economia de São Paulo. III. Fundação Getulio Vargas. IV. Título.

CDU 33

Ficha Catalográfica elaborada por: Isabele Oliveira dos Santos Garcia CRB SP-010191/O
Biblioteca Karl A. Boedecker da Fundação Getulio Vargas - SP

GABRIEL DOLHNIKOFF JARDANOVSKI

# WHO MATTERS IN DYNAMIC COORDINATION PROBLEMS?

Dissertação apresentada à Escola de Economia de São Paulo como pré-requisito à obtenção de título de mestre em Economia de Empresas.

Data de aprovação: 06/05/2020

Banca examinadora:

Prof. Dr. Bernardo de Vasconcellos Guimarães
FGV-EESP (Orientador)

Prof. Dr. Braz Ministério Camargo
FGV-EESP

Profa. Dra. Ana Elisa Gonçalves Pereira
Universidad de los Andes

# Agradecimentos

Ao Bernardo Guimarães, meu orientador, pelo apoio no decorrer dessa dissertação, pela lucidez e atenção nos momentos de angústia e pelos muitos ensinamentos, tanto dentro quanto fora da sala de aula.

Aos membros da banca, Ana Elisa Pereira e Braz Camargo, pela leitura cuidadosa e comentários valiosos que muito enriqueceram esse trabalho. Agradeço ao Braz, ainda, pelas longas e instrutivas conversas sobre assuntos que por vezes iam além do mestrado.

Aos meus colegas de mestrado, por compartilharem comigo essa árdua jornada. Em especial agredeço ao Ângelo, Caio, Otávio Braga, Otávio Teixeira, João, José, Alexandre, Nicolas, Victor, Fernando, Ariel, Gabriel, Leonardo e Pedro por me ensinarem tanto. Das coisas que aprendi com vocês economia é a menos importante.

Sou grato aos meus amigos de longa data Luis, Rodrigo, Matias, Lucas, Vinicius, Ives, Gabriel, Gabriela, Michael e Victor por me darem juízo quando dei passos maiores que as pernas; e por abraçarem minhas ideias quando eu já não mais acreditava nelas. À Eridiane, por ser minha aliada de todas as horas. À Joana agradeço pelo longo caminho percorrido conjuntamente, na certeza de que me tornou mais sensível ao mundo. À Daniela, pela simplicidade do afeto. A sua alegria e riso despertam o que há de melhor em mim.

Aos meus pais, Marisa e Otavio, e à minha irmã, Laura, agradeço não só pelo amor e dedicação incondicionais mas, principalmente, por sua resiliência e entusiasmo serem motivo de enorme inspiração.

Por fim, agradeço à minha filha, Laura, pelo amor, paciência e compreensão. A felicidade que você me trouxe eu nem sabia que poderia existir.

O presente trabalho foi realizado com apoio da Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES) – Código de Financiamento 001.

# Resumo

O presente artigo estuda um modelo de coordenação dinâmica em que agentes recebem oportunidades aleatórias para tomarem suas decisões. Há heterogeneidade em diversas dimensões. Em particular, cada agente pode afetar e ser afetado pelos demais de diferentes formas, e a frequência com que cada agente toma suas decisões também pode variar. O modelo apresenta um equilíbrio único. Por vezes a economia pode ficar presa em um regime ineficiente de baixa atividade e subsídios podem aumentar o bem-estar. Nós encontramos que o subsídio ótimo não depende da fricção temporal de cada tipo: a cada momento, o planejador deve simplesmente compensar a externalidade gerada por cada agente naquele intervalo de tempo. Mostramos esse resultado em um ambiente geral que pode ser usado para estudar focalização ótima de diversas intervenções em uma ampla classe de problemas de coordenação dinâmica envolvendo agentes heterogêneos.

**Palavras-chave**: coordenação, subsídio ótimo, agentes heterogêneos, jogos dinâmicos.

Abstract

This paper studies a dynamic model of coordination with timing frictions and heterogeneity in several dimensions. In particular, each agent might affect and be affected by others in different ways, and the frequency of their decisions might differ. There is a unique equilibrium in the model. At times, the economy might be stuck in an inefficient low-output equilibrium and subsidies can improve welfare. We find that the optimal subsidy does not depend on each type's timing frictions: at each point in time, the planner should simply compensate each agent for its externality on others at that particular moment. We show this in a rather general environment that can be used to study optimal targets for a variety of interventions in a large class of dynamic coordination problems with heterogeneous agents.

**Keywords**: coordination, optimal subsidies, timing frictions, heterogeneous agents, dynamic games.

# List of Figures

# Contents

# 1 Introduction

Several economic problems exhibit strategic complementarities, where decisions of agents mutually reinforce one another. For example, firms' investment decisions depend on the expected demand for their goods, which in turn depends on whether other firms decide to invest or not. In these settings, coordination failures may arise, and effective intervention could raise welfare. What is the optimal stimulus policy for an economy subject to this problem? In particular, to which firms should a government intervention be targeted?

Sákovics and Steiner (2012) study this problem in a static environment. They find that the optimal subsidy scheme should stimulate those who impose high externalities on others and on whom others impose low externalities. In this paper, we study the problem of optimal stimulus policies in a dynamic coordination game. In our setting coordination failures may trap the economy in a low-output regime: less investment today implies lower economic activity and lower investment tomorrow.

Firms differ not only on how they affect others but also on how often they make decisions. We introduce this to capture differences in adjustment costs to labor and capital inputs, both of which significantly vary across sectors of the economy (Hall, 2004). Hence, a dynamic setting adds a novel dimension to the question about who matters in coordination problems: should government intervention be targeted to those who make frequent decisions and would quickly react to subsidies or to those who make long-lasting investment decisions?

Subsidizing firms with low adjustment costs might be cheaper: they are more prone to investing since doing so does not entail a strong commitment for a long period. Also, once the subsidies are in place, it would not take much time for them to effectively start investing and generate positive spillovers to other firms. On the other hand, their decision to invest today tell us little about their decisions on the future as they can easily move back-and-forth between investment regimes. Thus, targeting firms with low adjustment costs could prove to be a waste: the subsidy would alter their decisions for a brief period of time and they would rapidly return to the low-output regime.

This contrasts with firms with high adjustment costs in that, once having decided to invest, those firms are expected to keep investing for longer periods. This would improve

other firms' expectations regarding the aggregate demand in future periods and could thus be more effective in pushing the economy to a high-output regime. At the same time, it would take longer for firms with high adjustment costs not currently investing to undertake more ambitious projects. They would not, then, immediately feed back in the investment loop. Without a formal model it is hard to make sense of all those trade-offs and identify what are the combination of characteristics of a firm that optimal interventions should target at.

We consider a dynamic environment with timing frictions as in Frankel and Pauzner (2000). Agents make a binary choice between two actions (say, investing or not). Their instantaneous utility flow depends on an exogenously moving fundamental (which captures, for example, the overall economy's productivity) and on how many others are in the network. Agents get opportunities to revise their behavior according to a Poisson clock, which can be seen as a simple way to capture production decisions that cannot adjust overnight.

Crucially, our setting adds much heterogeneity to the model: some agents might matter more than others for each agent in the economy; and the parameter determining how often agents make decisions might also differ across groups, a feature that allows us to incorporate the dynamic dimension in the planner's problem. Guimaraes and Pereira (2017) have extended the framework of Frankel and Pauzner (2000) to include some heterogeneity in preferences. However, in their setting, there is no heterogeneity regarding how often agents make decisions and how a given type affects the payoff of others.

We first show there is a unique rationalizable equilibrium characterized by a threshold for each agent. Each threshold depends on the number of agents of each type in a network and on the exogenous fundamental. We then solve the planner's problem to understand the inefficiencies that arise in equilibrium and assess the effects of policies targeting distinct types.

Our planner is financially unconstrained and is endowed with full commitment. Also, it is capable of providing *flow* subsidies, i.e. subsidies on the flow gains of each agent. Those are rather strong assumptions which enormously facilitate our analysis. We thus see our model as a benchmark, providing a starting point for future research addressing budgetary, commitment and implementation constraints.

In this environment, we find that the optimal flow subsidy schedule that implements

the planner's solution is insensitive to each type's revision parameter and only depends on network externalities. At each point in time, the planner should reward agents by their externality on others at that particular moment. Once that is done, no dynamic consideration affects the planner solution. One could think that a policy targeting firms that quickly react to subsidies could be a more effective way to take the economy out of a low-output regime. Our results show that rewarding firms for their current externalities is enough: firms form expectations about what other firms will do exactly as the planner forms expectations about the actions of its future selves.

When compensating externalities at each point in time the planner must bear in mind not only how influential is a given type on others' decision but also the size and composition of each type's network. When it comes to network size, the planner would rather allocate resources to a type who has a mild spillover on a large fraction of the population than to target types whose decisions have a big effect on a small group.

As of composition concerns, the planner must take into account what is the fraction of types choosing each action at each period. In our model agents derive gains from choosing the same action as others. Consider again the example of firms deciding to invest or not. A firm would only benefit from other firms investing if it has capacity to increase production and retain gains from the expanded demand - i.e. it only benefits if it also decides to invest. Economic sectors whose performance has large spillovers on the rest of the economy are natural targets of stimulus. But if this sector is only influential for firms who are currently not investing, then the gains from pushing it to invest are not as big. The planner would rather focus on another sector which does not generate as much externalities but whose network is comprised of a large share of firms currently investing. Because of this, the distribution of subsidies features *history dependence*: the initial composition of each network is crucial in determining which types should be prioritized.

This paper contributes to the study of how subsidies induce efficient coordination. A recent literature, of which prominent examples are Langer and Lemoine (2018) and Groote and Verboven (2019), uses dynamic models to study the role of government intervention in the adoption of new durable technologies (solar panels, more specifically). They allow for heterogeneity in tastes, but focus on the timing of subsidies and remain agnostic about the welfare implications of each policy. Moreover, in their models, technology adoption is an irreversible choice.

More closely related to ours, Guimaraes and Machado (2018) and Schaal and

Taschereau-Dumouchel (2018) study optimal subsidy policies in macroeconomic models subject to coordination failures, but do not allow for any heterogeneity. The first employs an homogeneous-agents model - otherwise similar to ours - to examine the impact of investment subsidies in regime change. Unlike this paper, their main result argues in favor of a constant stimulus policies, irrespective of which regime - low or high-activity - the economy is stuck at. As we introduce heterogeneous externalities, the optimal subsidy becomes a function of the share of each type stuck at each one of the two possible actions. Schaal and Taschereau-Dumouchel (2018) model a global-games coordination problem that is akin to our example of firms deciding to invest. Similar to Guimaraes and Machado (2018), they find that simple constant subsidies may increase welfare when coordination failures are present. Differently from them - and from the entire global games literature - we do not rely on noisy heterogeneous information to achieve equilibrium uniqueness. Instead, the key ingredients to resolve indeterminacy in our model are timing frictions and shocks to fundamentals.

Studying optimal stimulus policies over the business cycle is one of the possible applications of our model. We extend the results of Frankel and Pauzner (2000) to problems involving the interaction of many players, of many types. Our conclusions hold independently of the functional forms of preferences or network structure. Thanks to its robustness, our model can be applied to a large class of setups and allows us to characterize the coordination outcome and specify the optimal policy for a broad range of payoffs and shocks structures.

# 2 Model

The economy is populated by a continuum of infinitely-lived agents indexed by $i \in [0,1]$. Each agent $i$ is of a type $q(i) \in \mathcal{Q} := \{1, .., Q\}$. We let $n$ be the proportion of agents choosing action 1 and we denote by $\boldsymbol{n}$ the vector in $[0,1]^Q$ of shares of agents of each type currently playing 1, that is, $\boldsymbol{n} = (n_1, ..., n_Q)$. Moreover, we let $\alpha_q$ be the mass of type-$q$ players in the total population. This implies that the share of type-$q$ players currently choosing 1 among the total population is $\alpha_q \cdot n_{qt}$. The total mass of agents currently committed to 1 can be written as $n_t = \sum_{q=1}^{Q} \alpha_q n_{qt}$.

Time is continuous and agents discount the future at an homogeneous rate $\rho$. There are two possible actions, $a_i \in \{0,1\}$, but agents cannot switch from one to another at will. They receive chances to revise their actions according to a Poisson process with arrival rate $\delta^q$, which is type-specific, and stay committed to this choice until the arrival of another opportunity.

Agents' instantaneous payoff from either action depends on their own action, the action of others and some exogenous fundamental, summarized by $\theta$. For a type-$q$ agent the instantaneous payoff of action 1 is given by $u_1^q(\theta_t, \boldsymbol{n_t})$ and the payoff of action 0 is given by $u_0^q(\theta_t, \boldsymbol{n_t})$, both continuously differentiable for all $q \in \mathcal{Q}$. We define the instantaneous relative gain of action 1 for type-$q$ agents as

$$\triangle u^q(\theta_t, \boldsymbol{n_t}) = u_1^q(\theta_t, \boldsymbol{n_t}) - u_0^q(\theta_t, \boldsymbol{n_t}) \tag{2.1}$$

**Assumption 2.1.** *For all $q \in \mathcal{Q}$, the function $\triangle u^q$ is increasing in $\theta$ and in a subset of $\{n_1, ..., n_Q\}$.*

This is sufficiently general to allow for strategic complementarities arising due to one-sided externalities or two-sided externalities: either the payoff of choosing action 0 is independent of $\boldsymbol{n_t}$ but the payoff of choosing 1 is increasing in a subset of $\{n_1, ..., n_Q\}$ (as in the homogeneous case of Matsuyama (1991)); or both actions become more appealing the larger is the proportion of agents taking them (along the lines of Argenziano (2008)); or flow-payoffs from both actions can be increasing in each element of the relevant subset of $\{n_1, ..., n_Q\}$ but the difference in payoffs is also monotonically increasing in each one of those $n_q$ (as in the homogeneous case of Guimaraes and Machado (2018)).

A simple extension would also allow for strategic complementarities arising from congestion externalities. We would derive the same insights if we were to model a situation in which action 1 harms everyone but it is even more prejudicial to agents currently choosing action 0[1]. Our framework, however, is *not* suitable to model situations with mixed externalities, where agents' decisions generate positive spillovers to some types and negative spillovers to others.

A type is characterized by the triple $(u_1^q, u_0^q, \delta^q)$. At any given period $\tau$, a type-$q$ agent perfectly observes the fundamental, the current share of each type that is committed to playing 1 and forms expectations on the dynamics of $\boldsymbol{n}$. The expected discounted relative payoff of choosing action 1 at some date $\tau$ for type-$q$ agent is then

$$V_\tau^q = \int_\tau^\infty \mathrm{e}^{-(\rho+\delta^q)(t-\tau)} \mathbb{E}[\triangle u^q(\theta_t, \boldsymbol{n_t})] \, \mathrm{d}t \tag{2.2}$$

A type-$q$ agents chooses to play 1 whenever she gets the chance to revise their choice and $V_\tau^q > 0$, and to play 0 if $V_\tau^q < 0$. The agent discounts the future at rate $\rho$, but only until she gets selected again. She must then discount any future period $t$ by the probability of not being chosen by the Poisson process, $e^{-\delta^q(t-\tau)}$.

We further assume that payoff functions $\triangle u^q$ are such that there are dominance regions for all types of agents.

**Assumption 2.2.** *(Dominance regions) For all $q \in \mathcal{Q}$ there exists $P_q$ and $O_q$ such that if $\theta_\tau > P_q$ choosing 1 is strictly dominant for type-q agents; and if $\theta_\tau < O_q$ choosing 0 is strictly dominant for type-q agents.*

1 See Guimaraes et al. (2020) for a brief discussion on negative externalities.

# 3 Equilibrium

## 3.1 Multiple equilibria benchmark

Suppose $\theta_t$ is constant and equal to $\theta$. Because of our assumption of dominance regions, we can compute for any $q$ the upper dominance region boundary $P_q$, that is, a surface on the $\mathbb{R} \times [0,1]^Q$ space along which any type-$q$ agent is indifferent between playing 1 or 0 under her personal belief that every one else, of every other type, will choose to play 0 whenever the opportunity to change actions presents itself. More formally, we define the upper dominance region boundary for type $q$ as the surface $P_q(\boldsymbol{n}_\tau)$ satisfying

$$\int_\tau^\infty \mathrm{e}^{-(\rho+\delta^q)(t-\tau)} \mathbb{E}\left[\triangle u^q(\theta_t, \boldsymbol{n}_t^\downarrow) | \theta = P_q\right] \mathrm{d}t = 0 \tag{3.1}$$

where $\boldsymbol{n}_t^\downarrow = (n_{1_t}^\downarrow, ..., n_{Q_t}^\downarrow)$ and $n_{qt}^\downarrow := n_{q_\tau} \mathrm{e}^{-\delta^q(t-\tau)}$, for all $q \in \mathcal{Q}$ represents the dynamics of $n_{q_t}$ under the pessimistic belief described above.

Analogously, we can define the lower dominance region boundary for type $q$ as the surface $O_q(\boldsymbol{n}_\tau)$ satisfying

$$\int_\tau^\infty \mathrm{e}^{-(\rho+\delta^q)(t-\tau)} \mathbb{E}\left[\triangle u^q(\theta_t, \boldsymbol{n}_t^\uparrow) | \theta = O_q\right] \mathrm{d}t = 0 \tag{3.2}$$

where $\boldsymbol{n}_t^\uparrow = (n_{1_t}^\uparrow, ..., n_{Q_t}^\uparrow)$ and $n_{qt}^\uparrow := 1 - (1 - n_{q_\tau})\mathrm{e}^{-\delta^q(t-\tau)}, \forall q \in \mathcal{Q}$ represents the dynamics of $n_{q_t}$ under the agent's optimistic belief that every other type-$q$ agent will choose to play 1 whenever the opportunity to change actions presents itself.

For any $t > \tau$ and $\boldsymbol{n}_\tau$, $\boldsymbol{n}_t^\uparrow >> \boldsymbol{n}_t^\downarrow$. Hence, for any given $\theta$ and $\boldsymbol{n}_\tau$, the expression in (3.1) must be lower than (3.2). For both to equalize 0, it must be that $\theta$ satisfying (3.1) is larger than $\theta$ that solves (3.2), that is, $P_q > O_q$, for any initial value $\boldsymbol{n}_\tau$. Intuitively, this states that agents require a larger fundamental to be indifferent between actions 1 and 0 when they expect every one else to play 0 than when they expect every one else to play 1.

Moreover, because $n_{r_t}^\downarrow$ is increasing in $n_{r_\tau}$ for all $r \in \mathcal{Q}$, at least one element of the vector $\boldsymbol{n}_t^\downarrow$ increases as at least one element of $\boldsymbol{n}_\tau$ increases[1]. Then, since $\triangle u^q(\theta, \boldsymbol{n}_t)$ is (weakly) increasing in $n_{r_t}$ for all $r \in \mathcal{Q}$, the quantity $P_q$ is (weakly) decreasing in $n_{r_\tau}$ for all r $\in \mathcal{Q}$. The more sensitive $\triangle u^q(\theta, \boldsymbol{n}_t)$ is to a certain type $r$, the more $P_q$ decreases as $n_{r_\tau}$

[1] Abusing notation, one could state that $\boldsymbol{n}_t^\downarrow$ is increasing in $\boldsymbol{n}_\tau$.

increases for this specific type. Analogously, since $n^{\uparrow}_{r_t}$ also increases with $n_{r_\tau}$ for all $r \in \mathcal{Q}$, $\boldsymbol{n^{\uparrow}_t}$ increases with $\boldsymbol{n_\tau}$ and thus $O_q$ decreases with any increment in $n_{r_\tau}$ for all $r \in \mathcal{Q}$, the degree of such a decrease depending on the sensitiveness of $\triangle u^q(\theta, \boldsymbol{n_t})$ to each type.

The equilibrium uniqueness result in this heterogeneous environment is presented in the next section. It is based on an argument that depends only on properties of the upper and lower boundaries, built upon extreme beliefs about others' actions. As discussed below, when we introduce stochastic shocks to fundamentals such boundaries collapse to one single surface for each type. This, in turn, allows us to state and prove the unique equilibrium result.

It is useful representing both boundaries, $P_q$ and $O_q$, for a simplified two-type, linear version of the model. In what follows, $\mathcal{Q} = \{A, B\}$, $\boldsymbol{n_\tau} = (n_{A_\tau}, n_{B_\tau})$, $u^q_0(\theta^0, \boldsymbol{n_\tau}) = \theta^0 + \alpha^q(1 - n_{A_\tau}) + \beta^q(1 - n_{B_\tau})$ and $u^q_1(\theta^1, \boldsymbol{n_\tau}) = \theta^1 + \nu^q n_{A_\tau} + \phi^q n_{B_\tau}, \forall q \in \{A, B\}$. Thus,

$$\triangle u^q(\theta, n_{A_\tau}, n_{B_\tau}) = \theta + \gamma^q n_{A_\tau} + \lambda^q n_{B_\tau}, \quad \forall q \in \{A, B\} \tag{3.3}$$

where $\theta = \theta^1 - \theta^0 - \alpha^q - \beta^q$, $\gamma^q = \nu^q + \alpha^q$ and $\lambda^q = \phi^q + \beta^q$, for both $q \in \{A, B\}$. Also, types differ in their arrival rates, $\delta^A$ and $\delta^B$.

Suppose that $\lambda^B = \gamma^A > \lambda^A = \gamma^B$, i.e. both are more sensitive to their own types. For this specific setting, $P_q$ and $O_q$ are planes in the $\mathbb{R} \times [0, 1] \times [0, 1]$ space parallel to one another, for both $q \in \{A, B\}$. Figure 1 represents both $P_A$ and $0_A$ for arbitrary parameter values. At any point "above"[2] $P_A$, a type-A agent strictly prefers to play 1: $\theta$ is sufficiently large to outweight the fact that every one else is prescribed to play 0 - in other words, choosing 1 is a dominant action. Likewise, choosing 0 is a dominant action for the region below $O_A$, where $\theta$ is smaller than what would make a type-A indifferent between 1 and 0 when every one else is expected to choose 1. A subset of the state space in between exhibits multiple equilibria: it supports both choosing 1 and 0, depending on the agent's beliefs about the actions of others. Although it is only possible to sketch the pessimistic and optimistic thresholds for a two-type case, the intuition extends to the general version of the model: for a given type $q$, the region where action 1 is dominant lies above the $P_q$ threshold, the action 0 dominant region lies below the threshold $O_q$ and always exists a subset of the state space with equilibrium multiplicity.

[2] We will use "above" to refer to regions where for any given contigency $(\theta, \boldsymbol{n})$, $\theta$ takes a larger value than the value $P_q(\boldsymbol{n})$, and we will use "below" whenever referring to regions where for all points $(\theta, \boldsymbol{n})$, $\theta < O_q(\boldsymbol{n})$.

## 3.2 Equilibrium Uniqueness

We now assume that the fundamental is subject to stochastic shocks. Agents can no longer perfectly anticipate the actions of others in the future, as there will always be a probability of $\theta$ falling in each one of each type's dominance regions. Assume that

$$d\theta_t = \mu dt + \sigma dB_t$$

where $dB_t$ is a standard Brownian motion, $\mu$ is a constant drift and $\sigma > 0$ is the volatility. The processes above is especially convenient as (i) it make $\theta$ subject to *frequent shocks* - for any interval of length $dt > 0$ we have that $\mathbb{P}(\theta_t = \theta_{t+dt}) = 0$; (ii) and those increments are *independent of history*, i.e. $\theta_{t+dt} - \theta_t$ is independent of $(\theta_s)_{s \le t}$. We can now state the equilibrium uniqueness result.

**Proposition 3.1.** *When $\theta$ is subject to shocks, there is a unique profile of strategies surviving iterated deletion of strictly dominated strategies. This profile of strategies is characterized by a vector of surfaces $\boldsymbol{\theta}^* = (\theta_1^*, ..., \theta_Q^*)$ in the $\mathbb{R} \times [0,1]^Q$ space such that, for each $q \in \mathcal{Q}$,*

$$a_{q,t} = \begin{cases} 1 & \text{if } \theta_t > \theta_q^*(\boldsymbol{n}_\tau) \\ 0 & \text{if } \theta_t < \theta_q^*(\boldsymbol{n}_\tau) \end{cases} \tag{3.4}$$

*Proof.* See Appendix A.1 □

Our proof builds on the proof of equilibrium uniqueness in Frankel and Pauzner (2000) and Frankel et al. (2003). It employs a strategy of iterative deletion of strictly dominated strategies, starting from the dominance regions. They key idea is that, no matter how remote these regions are, their existence triggers a iterative process of elimination of strategies until, for each type of agent, there is a single rationalizable strategy left. Pick a type-$q$ agent at any point on its upper dominance region boundary, $P_q$. She is indifferent between playing 0 and 1 under the most pessimistic believe that all other agents, of all other types, will choose to play 0 whenever they have the chance. As we introduce stochastic shocks to the fundamental, there is now a positive probability that, in some future period, $\theta$ falls within the upper dominance region of other type-$q$ agents, making them choose to play 1[3]. Thus, the type-$q$ agent can no longer hold the belief that all

[3] There is a positive probability that $\theta$ lies within the upper dominance regions of any other type, and the argument that follows would essentially be the same. For simplicity, we will present the argument in terms of updates of beliefs with respect to other type-$q$ players.

other players, of all other types, will choose 0 no matter the circumstances. The "new" most pessimistic belief is that agents will play 0 whenever it is not strictly dominated to do so, which is a bit more optimistic than the initial pessimistic belief under which we built $P_q$. Hence, there is a new level of fundamental, smaller than $P_q$, for which type-$q$ players are indifferent between the two actions. As the type-$q$'s upper dominance region boundary move downwards, playing 0 becomes strictly dominated for a larger portion of the state-space. The type-$q$ players, then, continue to update their beliefs, making their indifference surface (under the most possible pessimistic belief) even lower. Extending this reasoning to all following rounds and employing an analogous procedure starting from the lower dominance regions yields a unique rationalizable equilibrium.

As in Frankel and Pauzner (2000), this result can be achieved even with vanishing shocks to fundamentals (that is, when in the limit $\mu, \sigma \to 0$). Equilibria multiplicity in this environment is not robust to the introduction of the smallest amount of shocks.

The unique equilibrium is fully characterized by the indifference conditions

$$\int_{\tau}^{\infty} \mathrm{e}^{-(\rho+\delta^q)(t-\tau)} \mathbb{E}\Big[ \triangle\, u^q(\theta_t, \boldsymbol{n_t}) | \boldsymbol{\theta}^*, \theta_q^*(\boldsymbol{n_\tau}), \boldsymbol{n_\tau} \Big] \,\mathrm{d}t = 0 \quad \forall q \in \mathcal{Q} \tag{3.5}$$

where the operator $\mathbb{E}\left[\cdot|\tilde{\theta}, \theta_\tau, \boldsymbol{n}_\tau\right]$ denotes agents' expectation when the current state is $(\theta_\tau, \boldsymbol{n}_\tau)$ and she expects other to play according to the vector of strategies $\tilde{\theta}$.

Type-$q$ agents play 1 when $\theta$ is above $\theta_q^*$, and 0 if $\theta$ lies below $\theta_q^*$. Each threshold $\theta_q^*$ is a function of the distribution of arrival rates, of the own payoff function $\triangle u^q$ and, inasmuch as complementarities are relevant, of the payoff functions of other types. Differently from the homogeneous case and from the heterogeneous idiosyncratic preferences case in Guimaraes and Pereira (2017), the dynamics of the system does not depend only on $n$, the total mass of players playing 1, but on the vector $\boldsymbol{n}$ - that is, we must separately keep track of all types' different strategies in order to understand if a given initial condition $(\theta_\tau, \boldsymbol{n}_\tau)$ implies an upward or downward path for each $n_q$, and for the overall $n$.

# 4 The planner's problem

One important question is what kind of inefficiencies arise in this environment and how would a social planner optimally address them. We then turn to study the social planner's problem.

Instantaneous welfare is given by the sum of individual agents' payoff, weighted by the proportion of each type's agents locked at each action (0 or 1). Hence,

$$W(\theta_t, \boldsymbol{n}_t) = \sum_{q \in \mathcal{Q}} \alpha_q \left[ n_t^q u_1^q(\theta_t, \boldsymbol{n}_t) + (1 - n_t^q) u_0^q(\theta_t, \boldsymbol{n}_t) \right] \tag{4.1}$$

Total welfare is given by the discounted sum of instantaneous welfare:

$$\mathcal{W}_\tau = \mathbb{E}\left[ \int_{t=\tau}^{\infty} e^{-\rho(t-\tau)} W\left(\theta_t, \boldsymbol{n}_t\right) dt \right] \tag{4.2}$$

The social planner maximizes (4.2) choosing proportions $\phi_t^q \in [0,1]$ of each type among those that received a chance to switch actions that will pick action 1. Therefore, although the social planer can reevaluate the share of type-$q$ players choosing 1 at any point in time, he is subject to the same timing frictions as agents are. An optimum is characterized by a series of vectors $\{\phi_t^q\}_{q \in \{1,\ldots,Q\}}^{t \geq \tau}$ such that $d\mathcal{W}_t \leq 0$ for any possible deviation. Following Guimaraes et al. (2020), Proposition 2 characterizes the social planner's problem.

**Proposition 4.1.** *The planner's optimal strategy is characterized by a vector of surfaces* $\mathbf{Z} = (Z_1, .., Z_Q)$*, where each* $Z_q$ *belongs to the* $\mathbb{R} \times [0,1]^Q$ *space and for each* $q \in \mathcal{Q}$

$$\int_{t=\tau}^{\infty} e^{-(\rho+\delta^q)(t-\tau)} \frac{\partial W\left(Z_q, \boldsymbol{n}_t\right)}{\partial n_t^q} dt = 0 \tag{4.3}$$

*Proof.* See Appendix A.2 □

The social planner chooses $\phi^q = 1$ whenever $\theta > Z_q$ and $\phi^q = 0$ whenever $\theta < Z_q$. Importantly, the planner's problem is similar to the decentralized one: for each type $q$ he takes into account the exact same discount rate $(\rho + \delta^q)$ as the decentralized equilibrium and thus chooses according to a threshold $Z_q$ such that (4.3) holds with equality at $Z_q$ for all $\boldsymbol{n} \in [0,1]^Q$. The only difference is the planner's flow utility accrued over time.

The planner's decision rule incorporates the overall gains from inducing an additional type-$q$ agent to play 1. Due to strategic complementarities, this not only affects the agent who has been selected to switch, but also all other players. Therefore, the planner plays a modified game with its future selves in which the flow gain from marginally increasing a type's share of agents playing 1 is given by

$$\frac{\partial W(\theta_t, \boldsymbol{n}_t)}{\partial n_t^q} = \underbrace{[u_1^q - u_0^q]}_{\Delta u^q(\theta_t, \boldsymbol{n}_t)} + \underbrace{\sum_{r \in \mathcal{Q}} \alpha_r \left[ n_t^r \frac{\partial u_1^r}{\partial n_t^q} + (1 - n_t^r) \frac{\partial u_0^r}{\partial n_t^q} \right]}_{\text{Externality}} \tag{4.4}$$

The key difference between the planner's and the decentralized equilibrium, then, is the externality component. Whether the planner's threshold for a type-$q$ agent, $Z_q$, lies above or below the decentralized threshold $\theta_q^*$ at any given $\boldsymbol{n_t}$ depends on whether the positive externalities from marginally increasing a type-$q$'s share of agents playing 1 surpasses the negative externalities, that is, if $\sum_{r \in \mathcal{Q}} \alpha_r n_t^r \frac{\partial u_1^r}{\partial n_t^q} > - \sum_{r \in \mathcal{Q}} \alpha_r (1 - n_t^r) \frac{\partial u_0^r}{\partial n_t^q}$. This implies that the optimal policy boils down to making agents internalize the externalities they generate, which will make them decide according to the expression in (4.3).

**Proposition 4.2.** *A subsidy schedule that compensates each agent for the amount of externalities it generates at every point in time implements the first best.*

*Proof.* See Appendix A.3 □

In this sense, the optimal (flow) subsidy scheme abstracts from dynamical issues and instead focuses on resolving current externalities at each period, every period. This does not imply that the *distribution* of arrival rates is irrelevant to the overall amount of subsidies disbursed. Types' equilibrium strategies differ inasmuch as their parameter of revision differ. Hence, the distribution of arrival rates affects the distribution of firms investing in each state of the economy, which is crucial in determining optimal subsidies. However, Proposition 4.2 does imply that a social planner might as well be uninformed about the different arrival rates in order to design the optimal subsidy: it suffices to observe the vector $\boldsymbol{n_t}$ and the network externalities. This result generalizes the results found for the homogeneous case in Guimaraes et al. (2020) and the idiosyncratic tastes case in Guimaraes and Pereira (2017).

Even though we present a dynamic coordination model, our optimal subsidy scheme bears resemblance with that of the static model in Sákovics and Steiner (2012). Their

planner should "target agents who impose high externalities on others and on whom others impose low externalities". This intuition is partly extended to our case: for a given $\boldsymbol{n_t}$, flow subsidies are increasing in the amount of positive externalities each type generates. Unlike Sákovics and Steiner (2012), however, we also allow for negative externalities. Those could be originated in positive spillovers type-$r$ players currently playing 0 might derive from type-$q$ players playing 0, as in the case of two-sided externalities. Then, a marginal increase in the share of type-$q$ agents playing 1 would negatively affect the flow utility of those type-$r$ players. The planner would have to take that into account and possibly impose a *tax* on certain types to make them less likely to choose action 1. This result holds even when abstracting from budgetary issues from the social planner problem[1].

On top of that, our model also features *history dependence*. At any given time the distribution of subsidies depends not only on how each agent's decision affects the others but also on the stock of agents of each type playing 1 or 0. The initial condition $\boldsymbol{n_0}$, then, is crucial. To better discuss the implications of this let us revisit the business cycle example from the introduction. It would be appealing to target subsidies to firms that most generate externalities. Suppose, however, that this big influential firm largely affects firms who are currently not investing, whereas there is another firm which does not generates as much externalities but is placed at a network largely comprised of firms which are already investing. It could be optimal to allocate more resources to the former firm rather than to the latter.

A subsidiary case would be to consider firms whose performance affects very large sectors. The amount of total externalities generated by a given type $q$ is bigger the larger the mass of firms affected by that given type, $\alpha_r$. Thus, the planner must also bear in mind the the size of a firm's network.

We should then call to mind all those features when answering who matters in dynamic coordination problems. In accordance with common intuition, types whose performance has large spillovers on the rest of the economy are natural target of stimulus. This does not tell the entire story, though. The initial condition of the economy and the composition and size of each types' networks is also relevant. Our model shows that when

[1] Sákovics and Steiner (2012) discuss imposing taxes to certain groups when solving their planner's problem under a balanced budget. The reason for imposing taxes on their model is the underlying trade-off between subsidizing certain groups at the expense of others under a fiscal neutral scheme. This is different from the reasons why even a fiscally unrestricted planner would impose taxes on certain types in our model.

flow subsidies are available the planner should allot more resources to types who impose big *net* externalities on a *large* fraction of the population.

# 5 Conclusion

This paper extends the homogeneous dynamic coordination model of Frankel and Pauzner (2000) to study how heterogeneous network externalities and timing frictions impact subsidy schemes. We show that as long as flow subsidies are available, the optimal policy is designed in such a fashion as to embrace network externalities, but is insensitive to heterogeneous timing frictions.

There are three different dimensions in each network relevant for allocating subsidies: (i) how influential is each type to the decision of others on his network; (ii) the size of each type's network; and (iii) the fraction of agents among those affected by a given type who currently choose the same action as the targeted type. The planner should at each time allocate more resources to types who have larger spillovers on the economy, after taking into account those three dimensions.

We reach those conclusions in a rather general environment, allowing for a broad range of payoffs and shocks structures. However, we restrict our attention to a highly stylized version of the planner's problem, where no budgetary or commitment issues are at stake. This work would greatly benefit from future research relaxing those two important assumptions.

Most papers building on the framework of Frankel and Pauzner (2000) consider a homogeneous population. Our extension, generalizing Guimaraes and Pereira (2017), could be useful for a variety of settings where heterogeneity matters. Indeed, applied work on static coordination games employing the global games methodology has often allowed for heterogeneous populations [1].

[1] Sákovics and Steiner (2012) is one example, but this literature has also considered heterogeneity in wealth (Goldstein and Pauzner, 2004); roles (Goldstein, 2005); risk aversion and consumption profile (Guimaraes and Morris, 2007); financial health (Choi, 2014); and ideological convictions (Morris and Shadmehr, 2018).

# Bibliography

Argenziano, R. (2008). Differentiated networks: equilibrium and efficiency. *The RAND Journal of Economics 39*(3), 747–769.

Choi, D. B. (2014). Heterogeneity and stability: Bolster the strong, not the weak. *The Review of Financial Studies 27*(6), 1830–1867.

Frankel, D., S. Morris, and A. Pauzner (2003). Equilibrium selection in global games with strategic complementarities. *Journal of Economic Theory 108*(1), 1–44.

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

Goldstein, I. (2005). Strategic complementarities and the twin crises. *The Economic Journal 115*(503), 368–390.

Goldstein, I. and A. Pauzner (2004). Contagion of self-fulfilling financial crises due to diversification of investment portfolios. *Journal of Economic Theory 119*(1), 151–183.

Groote, O. D. and F. Verboven (2019). Subsidies and time discounting in new technology adoption: Evidence from solar photovoltaic systems. *American Economic Review 109*(6), 2137–72.

Guimaraes, B. and C. Machado (2018). Dynamic coordination and the optimal stimulus policies. *The Economic Journal 128*(615), 2785–2811.

Guimaraes, B., C. Machado, and A. E. Pereira (2020). Dynamic coordination with timing frictions: Theory and applications. *Journal of Public Economic Theory*, forthcoming.

Guimaraes, B. and S. Morris (2007). Risk and wealth in a model of self-fulfilling currency attacks. *Journal of Monetary Economics 54*(8), 2205–2230.

Guimaraes, B. and A. E. Pereira (2017). Dynamic coordination among heterogeneous agents. *Journal of Mathematical Economics 73*, 13–33.

Hall, R. E. (2004). Measuring factor adjustment costs. *The Quarterly Journal of Economics 119*(3), 899–927.

Langer, A. and D. Lemoine (2018). Designing dynamic subsidies to spur adoption of new technologies. Working Paper 24310, National Bureau of Economic Research.

Matsuyama, K. (1991). Increasing returns, industrialization, and indeterminacy of equilibrium. *The Quarterly Journal of Economics 106*(2), 617–650.

Morris, S. and M. Shadmehr (2018). Inspiring regime change. Working paper.

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

Schaal, E. and M. Taschereau-Dumouchel (2018). Coordinating business cycles. Working paper.

# Appendix

# APPENDIX A – Proofs

## A.1 Proof of Proposition 1

The proof of equilibrium uniqueness is similar to the proof of equilibrium uniqueness in Frankel et al. (2003), even though that is a static global-game model. 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 all strategy 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). All exemplifying figures depict the linear two-type case, with $\mathcal{Q} = \{A, B\}$ and $\boldsymbol{n} = (n_A, n_B)$.

### A.1.1 Iterations from above.

Fix a type $q$ and a contingency $(\theta, \boldsymbol{n})$ such that $\theta = P_q$. A typical type-$q$ agent is indifferent between choosing 0 and 1 under the pessimistic belief that everyone, whenever called upon to choose, will pick 0 under any circumstances. But when $\theta$ moves stochastically, there will always be a positive possibility that it will spend some time above other players' dominance region boundaries. Notice that even if $q$ is such that $P_q > P_r$ for all $r \in \mathcal{Q}$ when on contingency $\boldsymbol{n}$ (Point X in Figure 2c, from the perspective of type-$A$ agents) she cannot expect every other player to choose 0 under any circumstances. If $\theta$ moves slightly upwards it will be strictly dominant at least for other type-$q$ players to pick 1, and thus a fraction $\alpha_q$ of the agents that get the chance will not choose 0. Hence, the new most pessimistic believe that agents can hold consistent with the dominance regions is that each type-$q$ agent plays 1 when above $P_q$, and 0 when below it. In other words, agents do not play strictly dominated strategies. Under this not-so-pessimistic new belief, the agent on $P_q$ is not indifferent anymore, but strictly prefers to play 1. To make her indifferent, $\theta$ must be lower. We can then construct for each type $q$ a new boundary $P_q^2$, with $P_q^2 < P_q$, above which a type-$q$ player chooses 1 when she expects all other agents to play according

to $(P_q)_{q\in\mathcal{Q}}$. In Figure 3a we depict the new threshold $P_A^2$ for type-A players.

For any $\boldsymbol{n}$, $P_q^2 < P_q$. This implies that when on a point between $P_q^2$ and $P_q$, strategies that prescribe playing 0 for type-$q$ are now dominated. Hence, the not-so-pessimistic belief we had constructed does not hold in equilibrium anymore: when on $P_q^2$ under the belief that everyone else is playing according to $(P_q)_{q\in\mathcal{Q}}$, type-$q$ players cannot be indifferent between action 0 and 1 as there is a positive probability of $\theta$ falling in the region between $P_q^2$ and $P_q$. We must again lower $\theta$ of indifference to $P_q^3$, which is entirely below $P_q^2$ (Figure 3b). This procedure can be repeated *ad infinitum*. At each round, we look for the curve $P_q^k$ on which a type-$q$ player has zero discounted payoff when assuming that other agents play according to $(P_q^{k-1})_{q\in\mathcal{Q}}$. Denote the limit of this sequence by $(P_q^\infty)_{q\in\mathcal{Q}}$. Notice that each agent $i$ playing according to $P_q^\infty(i)$ is, in fact, an equilibrium: if she expects others to play according to $(P_q^\infty)_{q\in\mathcal{Q}}$, her best response is to play according to $P_q^\infty(i)$.

### A.1.2 Iterations from below.

We now turn to a iterative process starting from $O_q$. Again, consider a type-$q$ player on $O_q$. She is indifferent between 0 and 1 under the optimistic belief that everyone when called upon choosing will pick 1, always. When $\theta$ follows a Brownian motion, there is a positive probability of it lying in some players action 0 dominance region. As before, the most optimistic belief consistent with dominance regions must be updated and we must move $O_q$ upwards, to $O_q^2$. Again, we repeat the procedure *ad infinitum*. We reach the limit $(O_q^\infty)_{q\in\mathcal{Q}}$, which is a Nash Equilibrium.

### A.1.3 Limit of iterations coincide.

Finally, we argue that it must be that $O_q^\infty = P_q^\infty$ for all $q \in \mathcal{Q}$. As a result, there is a single surface $\theta_q^*$ for each type $q$ such that type-$q$ players choose to play 1 whenever $\theta > \theta_q^*$, and to play 0 whenever $\theta < \theta_q^*$.

First notice that it cannot be the case that $O_q^\infty > P_q^\infty$, for any type $q$ on any contigency. Otherwise there would be no strategy that survives IESDS and we know that this cannot be true, since a Nash Equilibrium exists and every Nash Equilibrium survives IESDS. Suppose by contradiction that for some $q \in \mathcal{Q}$, $O_q^\infty$ and $P_q^\infty$ do not coincide, as exemplified in Figure 4. Then, we can always make an upward translation of $(O_q^\lambda)_{q\in\mathcal{Q}}$ by $\lambda$

where

$$\lambda := \sup_{q \in \mathcal{Q}, \boldsymbol{n} \in [0,1]^Q} \{P_q^\infty - O_q^\infty\}$$

That is, $\lambda$ is the biggest distance between any two points among all possible pairs $P_q^\infty$ and $O_q^\infty$ of all $q \in \mathcal{Q}$, and all contingencies $\boldsymbol{n}$. For all agents of each type $q \in \mathcal{Q}$, we move $O_q^\infty$ upwards by $\lambda$ and arrive at $O_q^\lambda$. That is, for any given $\boldsymbol{n}$

$$O_q^\lambda = O_q^\infty + \lambda \quad \forall q \in \mathcal{Q}$$

This implies that for every $q$, $O_q^\lambda$ lies entirely above $P_q^\infty$ and for at least some type $r$, $O_r^\lambda$ touches $P_r^\infty$ in at least one point.

Without loss of generality, suppose that there is a single type $r$ and a single point at which $O_r^\lambda = P_r^\infty$. We shall call this point K, as exemplified in Figure 4. For all types other than $r$, $O_q^\lambda > P_q^\infty$ at all points. Consider a type-$r$ player choosing at K, under the belief that every other agent will play according to $(O_q^\lambda)_{q \in \mathcal{Q}}$. Her pay-off of choosing 1 must be strictly negative, since it would be zero if she believed every one would play according to $(P_q^\infty)_{q \in \mathcal{Q}}$, and therefore choose 1 more often for any Brownian path.

Consider now the payoff of type-$r$ choosing at point J, where $J = K - \lambda$. For her, the increments $(\theta_{\tau+s} - \theta_\tau)$ follow the same distribution as the increments for player positioned at point K, since $(O_q^\lambda)_{q \in \mathcal{Q}}$ is a translation of $(O_q^\infty)_{q \in \mathcal{Q}}$. Thus, for any Brownian path, player positioned at K will have $\theta_t$ above (below) each types' threshold if, and only if, player positioned at J also has observes $\theta_t$ above (below) of each types' threshold. Thus, they will always experience the same dynamics of $\boldsymbol{n}_t$, but player J always experiences a smaller fundamental. Thus, the relative pay-off of player J of choosing action 1 (which we know to be zero, since $(O_q^\infty)_{q \in \mathcal{Q}}$ is an equilibrium) must be smaller than that of player K. Thus we get the contradiction:

$$0 > (\text{Player K's gain of choosing } 1) > (\text{Player J's gain of choosing } 1) = 0$$

□

## A.2 Proof of Proposition 2

Fix $r \in \{1, ...Q\}$ and suppose that at a given date $\tau$ it is optimal for the planner to choose $\phi_\tau^r < 1$ and consider the following deviation: the planner increases $\phi_\tau^r$ in

$\Delta\phi^r > 0$ units today, but keeps the future values of $\phi_t^q \quad \forall q \in \{1, ..., Q\}$ unchanged, for any realization of the Brownian path. Increasing $\phi_\tau^r$ by $\Delta\phi^r$ today raises $n_\tau^r$ by $\delta^r \Delta\phi^r dt := d\phi^r$, but leaves $n_\tau^q$ unchanged for all other types $q \neq r$. For any given period $T > \tau$, the change in $n_T^r$ due to deviations in $\phi_\tau^r$ is "depreciated" by rate $e^{-\delta^r(T-\tau)}$. Because $\phi_t^q$ is unchanged, we have no deviations in $n_t^q$, $\forall q, t > \tau$. It then follows that

$$dn_T^r = d\phi^r e^{-\delta^r(T-\tau)}$$

This deviation is not profitable if

$$\frac{\partial \mathcal{W}_\tau}{\partial n_\tau^r} = \mathbb{E}\left[\int_{t=\tau}^{\infty} e^{-\rho(t-\tau)} \frac{\partial W(\theta_t, \mathbf{n}_t)}{\partial n_t^r} \frac{dn_t^r}{dn_\tau^r} dt\right] \leq 0$$

which becomes

$$\mathbb{E}\left[\int_{t=\tau}^{\infty} e^{-(\rho+\delta^r)(t-\tau)} \frac{\partial W(\theta_t, \mathbf{n}_t)}{\partial n_t^r} dt\right] \leq 0 \tag{A.1}$$

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

$$\mathbb{E}\left[\int_{t=\tau}^{\infty} e^{-(\rho+\delta^r)(t-\tau)} \frac{\partial W(\theta_t, \mathbf{n}_t)}{\partial n_t^r} dt\right] \geq 0 \tag{A.2}$$

This implies the following necessary conditions for optimality: $\forall q \in \{1, ..., Q\}$, if $\phi_\tau^q = 0$ then (A.1) holds; if $\phi^q = 1$, then (A.2) holds; if $\phi_\tau^q \in (0,1)$ then (A.1) holds with equality. Those are the necessary conditions for a Nash Equilibrium in the game where the type-specific relative payoff $\Delta u^q(\theta, n)$ is replaced by $\frac{\partial W(\theta_t, \mathbf{n}_t)}{\partial n_t^q}$. 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 $\frac{\partial W(\theta_t, \mathbf{n}_t)}{\partial n_t^q}$ satisfies the same conditions we imposed on $\Delta u^q(\theta, n)$ and we have shocks, the equilibrium is unique and therefore these necessary conditions are also sufficient for optimality.

□

## A.3 Proof of Proposition 3

Fix a type q. The solution to the planner's problem prescribes action 1 if

$$\int_{t=\tau}^{\infty} e^{-(\rho+\delta^q)(t-\tau)} \frac{\partial W(Z_q, \mathbf{n}_t)}{\partial n_t^q} \mathrm{d}t > 0 \tag{A.3}$$

and action 0 if this inequality is reversed. For the decentralized problem, choosing 1 is optimal if

$$\int_{\tau}^{\infty} \mathrm{e}^{-(\rho+\delta^{q})(t-\tau)} \mathbb{E}[\triangle u^{q}(\theta_t, \boldsymbol{n_t})] \,\mathrm{d}t > 0 \tag{A.4}$$

Adding the externality generated by a typical type-$q$ player to (A.4) yields the condition for the planner to set type-$q$ agents to choose 1, as in (A.3).

□

# APPENDIX B – Figures

Figure 1 – Dominance regions for type-A agent

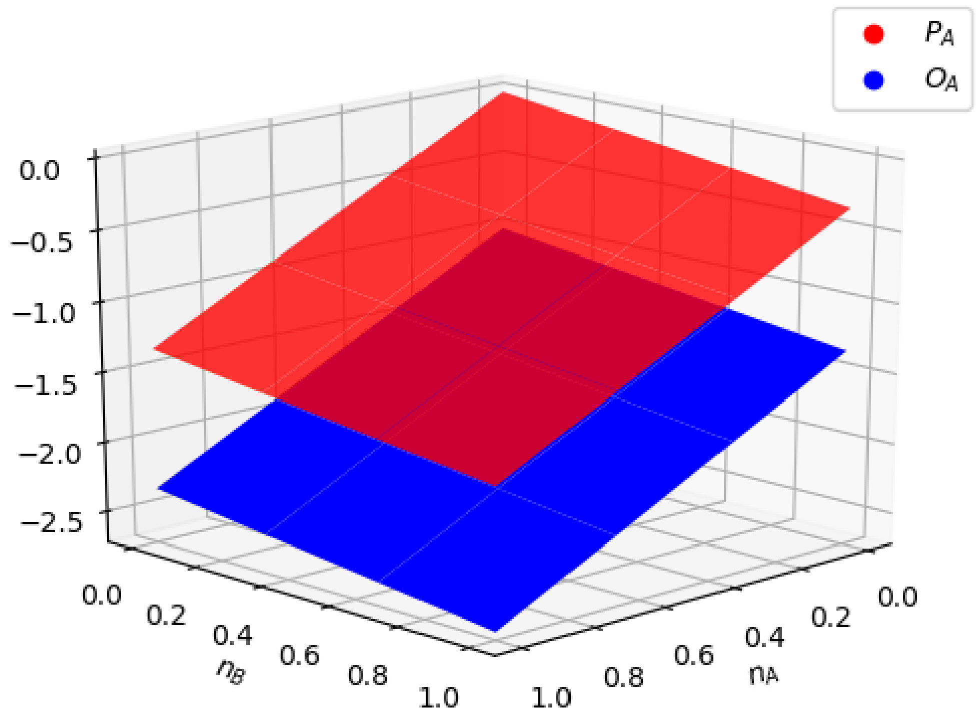


*Note:* Figure 1 presents upper and lower boundaries, $P_A$ and $O_A$, for a type-A agent who faces a linear utility function of the form of (3.3) with parameter values $\delta^A = \rho = 1/2$, $\gamma^A = 2$ and $\lambda^A = 1$ and is called to play at period $\tau = 0$. In this example, $P_A$ is such that $\mathbb{E}\left[\int_0^\infty \mathrm{e}^{-t}(\theta + 2n_A^{\downarrow} + n_B^{\downarrow})\,\mathrm{d}t \middle| \theta = P_A\right] = 0$, where $n_q^{\downarrow} = n_{q_0}e^{-t/2}$, for both $q \in \{A, B\}$.

Figure 2 – Action 1 dominance region for types $A$ and $B$

(a) Action 1 dominance region for type-$B$

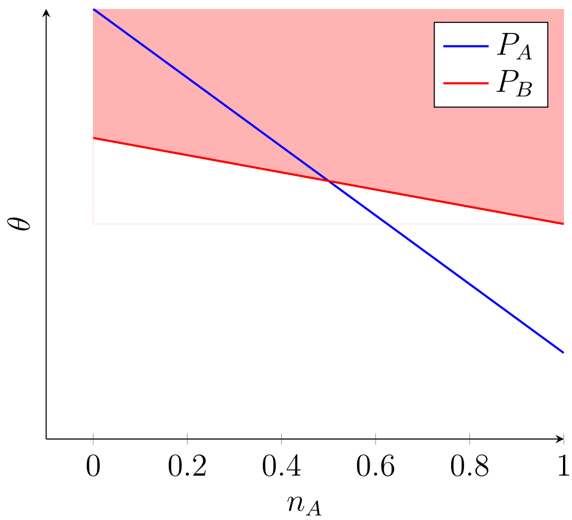


(b) Action 1 dominance region for type-$A$

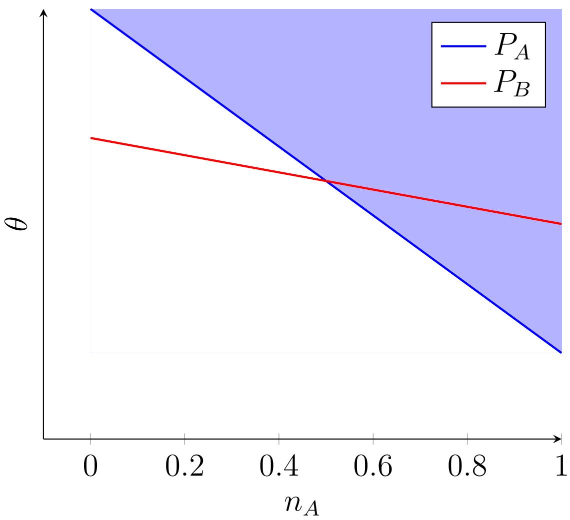


(c) Intersection of regions

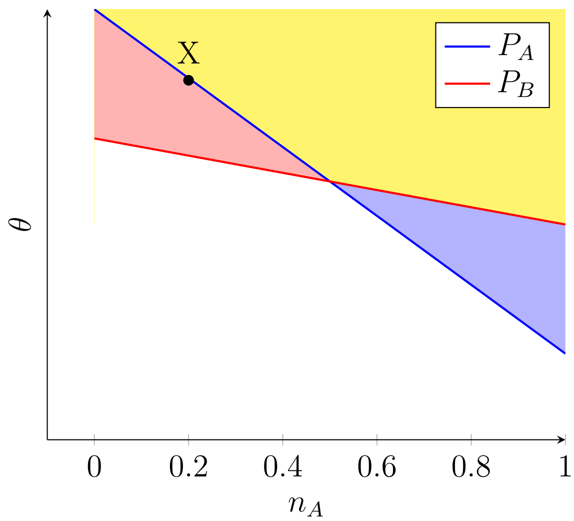


*Note:* We slice the 3-D figure at some arbitrary $n_B$. Figures above depict thresholds $P_A$ and $P_B$ for the two-type linear case. The red region shows the points for which action 1 is dominant for all type-$B$ agents and the blue region, for type-$A$ agents. The yellow area represents the intersection of both regions.

Figure 3 – Iterations from above

(a) Type-$A$ first round of iterations

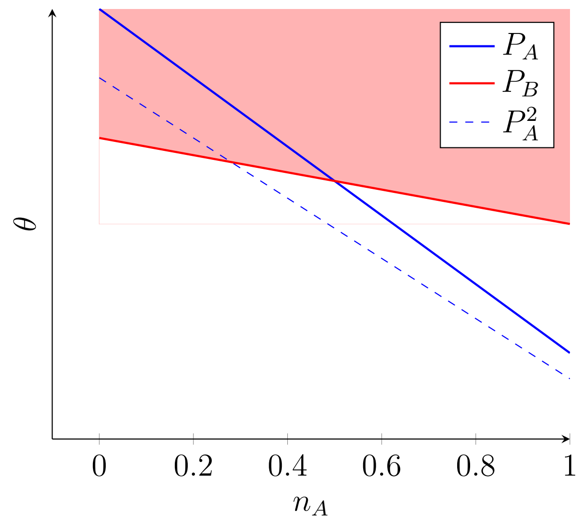


(b) Type-$A$ second round of iterations

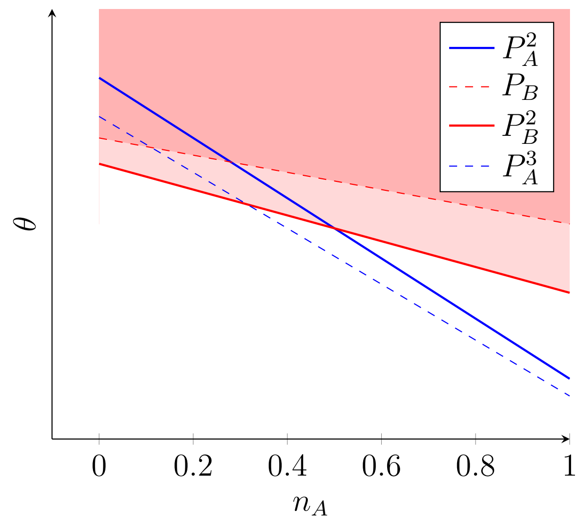


*Note:* We slice the 3-D figure at some arbitrary $n_B$. Figures above depict thresholds $P_A$ and $P_B$ for the two-type linear case. The red region shows the points for which action 1 is dominant for all type-$B$ agents.

Figure 4 – Translations

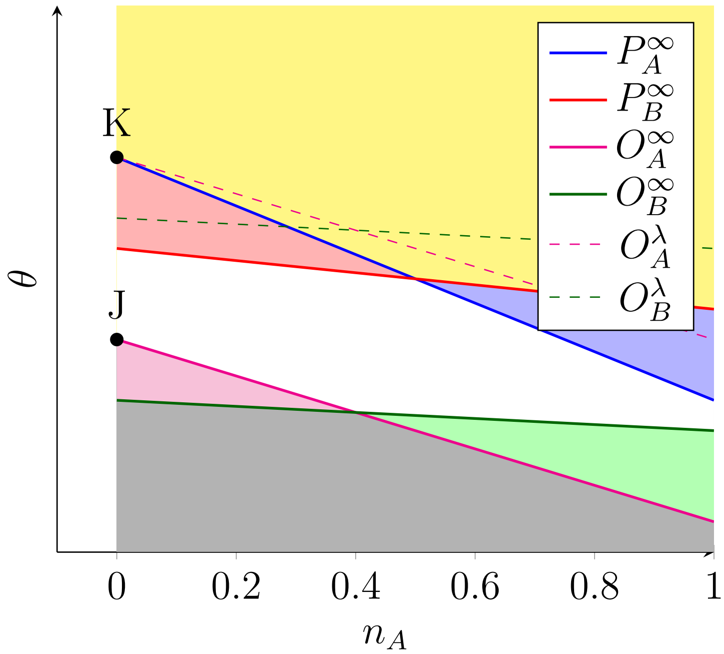


*Note:* We slice the 3-D figure at some arbitrary $n_B$. The figure depicts thresholds for the two-type linear case. When under the belief that all other agents play according to $\{P_q^\infty\}_{q\in\mathcal{Q}}$, type-$A$ agents choose 1 whenever $\theta$ lies in the red or yellow regions. Analogously, type-B players choose to play 1 whenever $\theta$ lies in the blue or yellow regions and they expect all others to play according to $\{P_q^\infty\}_{q\in\mathcal{Q}}$. Type-$A$ players choose to play 0 whenever $\theta$ is on the pink or gray regions and they expect all others to play according to $\{O_q^\infty\}_{q\in\mathcal{Q}}$; and type-$B$ players always choose 0 if $\theta$ lies in the green or gray regions and they expect all others to play according to $\{O_q^\infty\}_{q\in\mathcal{Q}}$.