

Local Interactions under Switching Costs

Ge Jiang* Simon Weidenholzer†

September 2016

Abstract

We study the impact of switching costs on the long run outcome in 2×2 coordination games played in the circular city model of local interactions. For low levels of switching costs the predictions are in line with the previous literature and the risk dominant convention is the unique long run equilibrium. For intermediate levels of switching costs the set of long run equilibria still contains the risk dominant convention but may also contain conventions that are not risk dominant. The set of long run equilibria may further be non-monotonic in the level of switching costs, i.e. as switching costs increase the prediction that the risk dominant convention is the unique long run equilibrium and the prediction that both conventions are long run equilibria alternate. Finally, for high levels of switching costs also non-monomorphic states will be included in the set of long run equilibria.

Keywords: Switching Costs, Local Interactions, Learning, Coordination Games.

JEL Classification Numbers: C72, D83.

*Nanjing University, School of Economics

†University of Essex, Department of Economics. Email: sweide@essex.ac.uk

1 Introduction

It is often costly to switch to a different technology or adopt a new social norm. For instance, switching from Windows to Apple requires not only getting familiarized with the new system but also moving files from one computer to the other. Further examples of switching costs include communicating one's new telephone number when switching providers in telecommunication or buying new tools when switching from inch screws to metric screws.

In these examples agents are better off when interacting with somebody who uses the same operating system, telecommunication provider, or industry standard, these situations typically give rise to coordination games. A wide range of models, starting with the seminal works of Kandori, Mailath, and Rob (1993) and Young (1993), have analyzed settings where a population of boundedly rational players decide on their actions in such coordination games using simple heuristics.¹ The message that emerges from these discussions is that, when players use best response learning, risk dominant strategies - that perform well against mixed strategy profiles - will emerge in the long run, even in the presence of payoff dominant strategies. In the context of the above examples, this implies that populations do not necessarily end up with technologies which maximize social welfare.

Norman (2009) has already analyzed the role of switching costs in a global interactions setting where everybody interacts with everybody else. In the global setting switching costs turned out to influence the speed at which the population approaches the long run equilibrium (LRE). The long run prediction remain unaffected, though. Quite frequently interactions are, however, local in nature, with interaction partners corresponding to family members, friends, or work colleagues. For instance, in the above examples on switching operating systems or telecommunication providers it is typically the case that this decision will to a larger degree be influenced by one's contacts or collaborators than by the overall distribution of technologies in the society.

The present paper aims to understand the role of switching costs in long run technology choice and the emergence of conventions under such local interactions. We capture local interactions by considering a model akin to the one proposed by Ellison (1993) where the agents are arranged around a circle and interact with their neighbors only. We focus on a setting where one strategy is risk dominant and the other strategy may or may not be

¹See Weidenholzer (2010) for a survey of the literature.

payoff dominant. This allows us to analyze circumstances under which strategies that are neither payoff- nor risk- dominant are selected. When determining which strategy to use the players play a best response to the distribution of play in their neighborhood in the previous period taking into account that switching strategies incurs a cost. In addition, choices are perturbed by occasional uniform (across agents and time) mistakes.

Our study of switching costs under local interactions reveals that: i) In the presence of sufficiently high switching costs risk dominant actions are no longer able to spread contagiously and profiles where different actions coexist are absorbing. ii) While for low switching costs the predictions are in line with the previous literature where risk dominant conventions are LRE, for higher switching costs the set of long run equilibria may additionally also contain conventions that are not risk dominant iii) The set of LRE may further be non-monotonic in the level of switching costs, i.e. as switching costs increase the prediction that the risk dominant convention is the unique long run equilibrium and the prediction that both conventions are long run equilibria alternate.

Let us now discuss our findings in more detail. When switching costs are relatively low, they do not affect agents' behavior. Thus, the risk dominant strategy is still able to spread contagiously, starting from a small cluster and eventually taking over the whole population. Consequently, low levels of switching costs do not change the predictions of the model as compared to the standard model without switching costs. However, for larger switching costs risk dominant strategies may no longer spread contagiously and non-monomorphic states, where different strategies coexist, become absorbing. The reason is that a player at the boundary of a risk dominant cluster will not switch under sufficiently high switching costs. Moreover, it is possible to move among all of these non-monomorphic absorbing states via a chain of single mutations. Transitions from different states to each others are, thus, characterized by step-by-step evolution as outlined in Ellison (2000).

The question which state will be LRE essentially boils down to how difficult the set of non-monomorphic states is to access from the two monomorphic states. Interestingly, if agents only interact with a few neighbors, there may exist a range of parameters where alongside the risk dominant convention also non-risk dominant conventions are LRE. Thus, switching costs may lead to the model's prediction no longer being unique. The reason behind this phenomenon is that -under the uniform noise approach- the number of mutations required to move from a convention to the set of non-monomorphic absorbing states is measured in

integers. Especially if agents only interact with a few neighbors, it may happen that the number of mistakes required to access the set of non-monomorphic absorbing states from the risk dominant convention equals the number of mistakes required to access this set from the non-risk dominant convention.

Perhaps even more interestingly, also owing to the fact the mutations are measured in integers, the prediction might be non-monotonic in the level of switching costs. That is, the prediction that the risk dominant convention is selected and the prediction that both conventions are selected alternate as switching costs increase. This curiosity is caused by i) the stepwise nature of rounding up and ii) by the fact that the number of mistakes required to leave the risk dominant convention and the number of mistakes required to leave the non-risk dominant convention only differ by a constant.

Finally, for very high levels of switching costs no player will switch in the absence of noise even if all neighbors choose the other strategy. Thus, all states are absorbing and can be connected via a chain of single mutations. Consequently, all absorbing states turn out to be LRE.

For large interaction neighborhoods the integer problem ceases to have impact and the risk dominant convention remains as unique LRE. In particular, this holds true if every agent interacts with every other agent and a sufficiently large population, thus, reconciling our results with those of Norman (2009).

We believe that two important insights arise from our analysis. The first implication stems from taking the model's predictions at face value. As our model predicts that under switching costs actions that are not risk dominant may arise as LRE, it may well contribute to our understanding of the emergence and survival of such actions. Note that the welfare implications of switching costs depend on whether the risk dominant action is also payoff dominant or not. If the risk dominant action is not payoff dominant, then the presence of switching costs implies that payoff dominant conventions will be observed with positive probability in the long run. Switching cost might, thus, be welfare improving. If, however, an action is both risk- and payoff- dominant the presence of switching costs may lead to (risk- and payoff-) dominated strategies surviving in the long run. Switching costs and local interactions may, thus, also explain why inefficient technology standards or norms survive in the long run.

The second insight is based on a more pessimistic reading of our results and concerns the

robustness of the local interaction model. As the local interaction model can no longer give a clear cut prediction, one may argue that it loses traction in the presence of switching costs. This is expressed by the non-uniqueness of the long run prediction but even more aggravated by the non-monotonicity of the prediction. While the risk dominant convention ceases to be unique LRE for high enough switching costs it might be again unique LRE for even higher switching costs. This is bad news since the circular city model of local interactions has some otherwise nice features as compared to the global model: i) It was observed by Ellison (1993) that in contrast to the global interaction model of Kandori, Mailath, and Rob (1993) it features a high speed of convergence. ii) Lee, Szeidl, and Valentinyi (2003) have shown that it is immune against the Bergin and Lipman (1996) critique. iii) (Weidenholzer 2012) has shown that it is robust to the addition and deletion of dominated strategies, a test which Kim and Wong (2010) have shown the global model fails.

The paper closest related to our work is Norman (2009) who studies switching costs in the context of a global interactions model. As already observed by Kandori, Mailath, and Rob (1993) a major drawback of the global interactions model lies in its low speed of convergence. Under global interactions the number of mistakes required to move from one convention to another turns out to depend on the population size. Thus, in large populations it is questionable whether the long run limit will be observed within any reasonable time horizon.² Norman (2009) shows how switching costs might speed up convergence. As in the present paper, the presence of switching costs implies that non-monomorphic states where agents use different actions become absorbing. This enables a transition from one convention to another by first accessing the class of non-monomorphic states and then moving through this class via a chain of single mutations to the other convention. Under switching costs the step from one convention to the set of non-monomorphic states is typically smaller than the direct step from that convention to the other. Consequently, switching costs may speed up the convergence to the long run prediction.

Norman's (2009) analysis focuses on the case of large populations, thus, ruling out the integer effects that drive our main results on the composition of the set of LRE. Our study reveals that under local interactions (but also under global interactions in small populations) the presence of switching costs may have implications that go beyond the speed of conver-

²Ellison (1993) pointed out that in the context of local context of local interactions where some strategies might spread contagiously the speed of convergence is independent of the population size and, thus, the LRE might be a reasonable predictor even in large populations.

gence. In this light, the predictions in populations characterized by local interactions (or in small populations with global interactions) may be fundamentally different to those in large populations with global interactions.

While the present paper adds to the ongoing discussion on learning in coordination games it also contributes to a wider discussion on how far received results in the literature on learning in games are robust to (minor) modifications. For instance, under imitation learning changing the interaction or information structures may result in different predictions in coordination games (see Robson and Vega-Redondo 1997 and Alós-Ferrer and Weidenholzer (2006, 2008)) or prisoner dilemma games (see Eshel, Samuelson, and Shaked 1998, Mengel 2009). Similarly, in Cournot games a number of contributions have analyzed conditions under which firms converge to the Walrasian state under imitation learning, as predicted by Vega-Redondo (1997). Alós-Ferrer (2004) shows that when agents have memory over the last two periods the Walrasian state is no longer uniquely stochastically stable.³ Apesteguia, Huck, Oechssler, and Weidenholzer (2010) find that, if there are differences in cost functions, all monomorphic states are absorbing. However, the Walrasian state remains unique LRE if no firm is the uniquely cheapest one, as shown by Tanaka (1999).

The rest of this paper is organized in the following way. Section 2 presents the model and discusses the main techniques used. Section 3 spells out our main results and Section 4 concludes.

2 The model

We consider a population of N agents who are located on a circle, as in Ellison (1993). A given agent i has agents $i - 1$ and $i + 1 \pmod{N}$ as immediate neighbors. Each agent interacts with her k closest neighbors to the left and to the right of her. We assume $k \leq \frac{N-1}{2}$ to ensure that no agent interacts with herself. Thus, agent i 's interactions are confined to the set of players $N(i) = \{i - k, i - k + 1, \dots, i - 1, i + 1, \dots, i + k - 1, i + k\}$. The agents in the set $N(i)$ are called neighbors of i .

We assume $|N|$ to be odd. This allows us to nest global interactions in our framework by setting $k = \frac{N-1}{2}$.⁴

³Alós-Ferrer and Shi (2012) consider asymmetric memory which turns out to affect equilibrium selection in coordination games but reinforces the stability of the Walrasian state in Cournot games.

⁴The results obtained for local interaction also hold for even populations.

Each agent i plays a 2×2 coordination game with strategy set $S = \{A, B\}$ against all agents in her neighborhood $N(i)$. We denote by $u(s_i, s_j)$ the payoff agent i with strategy s_i receives when playing against agent j with strategy s_j . We follow Eshel, Samuelson, and Shaked (1998) and consider the following (normalized) coordination game.

	A	B
A	α, α	$\beta, 0$
B	$0, \beta$	$1, 1$

We assume $\alpha > 0$ and $\beta < 1$, so that (A, A) and (B, B) are both strict Nash equilibria. Further, we assume $\alpha + \beta > 1$, so that the equilibrium (A, A) is risk dominant in the sense of Harsanyi and Selten (1988), i.e. A is the unique best response to a mixed strategy profile which puts equal probability on A and B . We denote by

$$q^* = \frac{1 - \beta}{1 + \alpha - \beta}$$

the critical mass put on A in a mixed strategy equilibrium. Risk dominance of the Nash equilibrium (A, A) translates into $q^* < \frac{1}{2}$. Note that if $\alpha > 1$, (A, A) is payoff dominant and if $\alpha < 1$, (B, B) is payoff dominant. However, no such assumption on α is made at this stage.

The number of A -players in the population is denoted by $m = \#\{i \in I | s_i = A\}$ and the number of A -players among agent i 's neighbors is denoted by $m_i = \#\{j \in N(i) | s_j = A\}$. Accordingly, the number of B -players in the population is given by $N - m$ and the number of B -players in i 's interaction set is given by $2k - m_i$.

We denote by $s_i(t)$ the strategy adopted by player i , by $s(t) = (s_1(t), \dots, s_N(t))$ the profile of strategies adopted by all players, and by

$$s_{-i}(t) = (s_{i-k}(t), \dots, s_{i-1}(t), s_{i+1}(t), \dots, s_{i+k}(t))$$

the strategies adopted by all of player i 's neighbors in period t . Further, the monomorphic states (s, s, \dots, s) where all agents adopt the same strategy s are denoted by \vec{s} .

The payoff for player i is given by the average payoff received when interacting with all

neighbors.

$$U_i(s_i(t), s_{-i}(t)) = \frac{1}{2k} \sum_{j \in N(i)} u(s_i(t), s_j(t)).$$

We consider a myopic best response process with *switching costs*. In each period $t = 1, 2, \dots$ each agent receives the opportunity to revise her strategy with exogenous probability $\eta \in (0, 1)$.⁵ Changing strategies is assumed to be costly. Whenever an agent changes her strategy she is subject to a switching cost. We follow Norman (2009) and consider switching costs c which are independent of the current action choice and enter the payoff function in an additive way.⁶ The following function formalizes this idea

$$c(s_i(t), s_i(t+1)) = \begin{cases} c & \text{if } s_i(t) \neq s_i(t+1) \\ 0 & \text{if } s_i(t) = s_i(t+1) \end{cases}.$$

When a revision opportunity arises an agent switches to a myopic best response, i.e. she plays a best response to the distribution of play in her neighborhood in the previous period, taking into account the switching costs. More formally, at time $t+1$, when given revision opportunity, player i chooses

$$s_i(t+1) \in \arg \max_{s_i(t+1) \in S} \left[U(s_i(t+1), s_{-i}(t)) - c(s_i(t), s_i(t+1)) \right].$$

If a player has multiple best replies, it is assumed that she randomly chooses one of them with exogenously given probability. If she does not receive an opportunity to revise her strategy, she chooses $s_i(t+1) = s_i(t)$. Further, with fixed probability $\epsilon > 0$, independent across agents and across time, the agent ignores her prescription and chooses a strategy at random, i.e. she makes a mistake or mutates.

We denote the state space by Ω and a state of the process by ω . The process with mistakes is called *perturbed process*. Under the perturbed process any two states can be reached from each other. Thus, the only absorbing set is the entire state space, implying that the process is ergodic. The unique invariant distribution of this process is denoted by $\mu(\epsilon)$. We are interested in the *limit invariant distribution* (as the rate of experimentation tends to zero),

⁵Thus, we are considering a model of positive inertia where agents may not adjust their strategy every period.

⁶Alternative formulations of switching costs encompass situations where the level of switching costs depends on the current strategy used or on the current level of payoffs.

$\mu^* = \lim_{\epsilon \rightarrow 0} \mu(\epsilon)$. Such a distribution exists (see Foster and Young (1990), Young (1993), or Ellison (2000)) and is an invariant distribution of the process without mistakes (the so called *unperturbed process*). It gives a stable prediction for the original process, in the sense that for ϵ small enough the play approximates that described by μ^* in the long run. The states in the support of μ^* , are called *Long Run Equilibria (LRE)* or stochastically stable states. The set of LRE is denoted by $\mathcal{S} = \{\omega \in \Omega \mid \mu^*(\omega) > 0\}$. We use a characterization of the set of LRE due to Freidlin and Wentzell (1988).⁷ Consider two absorbing sets of states X and Y and let $C(X, Y) > 0$ (referred to as a *transition cost*) denote the minimal number of mutations for a transition from the X to Y . An X -tree is a directed tree such that the set of nodes is the set of all absorbing sets, and the tree is directed into the root X . For a given tree one can calculate the cost as the sum of the costs of transition for each edge. According to Freidlin and Wentzell (1988), a set X is a LRE if and only if it is the root of a minimum cost tree.

3 The role of switching costs

In a first step we will study how switching costs influence the agent's decision to switch strategies. Consider an A -player. She will switch strategies with probability one if her payoff from playing B minus the switching cost strictly exceeds her payoff from remaining an A -player, i.e.

$$\frac{1}{2k} \left(m_i \alpha + (2k - m_i) \beta \right) < \frac{1}{2k} \left(2k - m_i \right) - c.$$

Rearranging terms yields

$$m_i < 2kq^* - \frac{2kc}{1 + \alpha - \beta} := m^A(c, k)$$

An A -player will remain an A -player with certainty whenever $m_i > m^A(c, k)$ and will choose A and B with positive probability if $m_i = m^A(c, k)$. As $m^A(c, k)$ is the minimum number of A -playing neighbors such that keeping A is a unique best response, it cannot be negative.

Likewise, consider a B -player. She will switch strategies with probability one if the payoff

⁷See Fudenberg and Levine (1998) or Samuelson (1997) for textbook treatments. Ellison (2000) provides an enhanced (and sometimes easier to apply) algorithm for identifying the set of LRE. We chose to work with the original formulation as it allows for a characterization in case of multiple LRE.

from playing A minus the switching cost exceeds her current payoff, which yields

$$m_i > 2kq^* + \frac{2kc}{1 + \alpha - \beta} := m^B(c, k).$$

A B -player will remain a B -player if $m_i < m^B(c, k)$, and will randomize between the two strategies if $m_i = m^B(c, k)$. Note that $m^B(c, k)$ is defined as the number of A -players such that a player with less than $m^B(c, k)$ A -neighbors chooses to stay at B with certainty and, thus, cannot exceed $2k$.

Note that $m^A(0, k) = m^B(0, k) = 2kq^*$, i.e. in the absence of switching costs the thresholds are the same as in Ellison's (1993) model. For $c > 0$, we have $m^A(c, k) < m^A(0, k) = m^B(0, k) < m^B(c, k)$. Hence, in the presence of switching costs, it takes more players of the other type to induce a switch than in the absence of switching costs. Further, a B -player will require more A -opponents to switch strategies than an A -player requires to stay at her strategy. Likewise, an A -player will switch to B at a lower number of A -opponents than it takes a B -player to remain at her strategy. Thus, switching costs create regions where players with the same distribution of play in their neighborhood but with a different current strategy may behave differently. This may lead to the emergence of non-monomorphic absorbing states where clusters of players with different strategies coexist. In such states all players want to remain at their current strategies, i.e. $m_i > m^A(c, k)$ for all A -players and $m_j < m^B(c, k)$ for all B -players.

In the following, G denotes the set non-monomorphic absorbing states, i.e.

$$G = \{s \in S \mid s \neq \vec{A}, \vec{B}, m_i > m^A(c, k) \forall i \text{ with } s_i = A, \text{ and } m_j < m^B(c, k) \forall j \text{ with } s_j = B\}.$$

and an element of this set is denoted by AB . Further, G_ℓ denotes the set of non-monomorphic absorbing states with ℓ A -players (and $N - \ell$ B -players), i.e.

$$G_\ell = \{s \in G \mid m = \ell\}.$$

3.1 Two-neighbor interaction

In order to build intuition and to highlight the main mechanisms at work, our analysis starts with an informal discussion of the special case where each agent only interacts with her two

most immediate neighbors, i.e. $k = 1$. A comprehensive analysis of the case $k \geq 1$ is provided in Section 3.2.

In a first step, let us consider under which circumstances non-monomorphic states are absorbing. To this end, consider states where clusters of A -players and B -players, each of at least size two, alternate, e.g.

$$\dots BBAABBBAAAA \dots$$

Players in the middle of such a cluster only interact with players of their own kind and, hence, will never switch. Thus, let us focus on the boundary between two such strings. Note that whenever $m^A(c, 1) < 1$ holds the boundary A -player will keep her strategy. This translates into $2c > 1 - \alpha - \beta$, which is implied by risk dominance of A . Thus, the boundary A -player will remain. Now consider the B -player. Note that if $m^B(c, 1) \leq 1$ holds, the boundary B -player will switch to A with positive probability. This translates into $c \leq \frac{\alpha + \beta - 1}{2}$. Thus, provided switching costs are low, the A -cluster will grow contagiously, even in the absence of mistakes. If this condition is violated, $c > \frac{\alpha + \beta - 1}{2}$, the boundary B -player will stay a B -player with certainty. This, in turn, implies that for sufficiently high switching costs non-monomorphic states are absorbing.

Surprisingly, switching costs may not only alter the set of absorbing states but may also change the set of LRE. To see this, first note that one can move among the set of non-monomorphic states via a chain of single mutations. More precisely, it is possible to move from a state in G_ℓ to either a state in G_a or in G_b , with $a < \ell < b$ at the cost of one mutation. While, it is clear that one mutation to A (or B) increases (decreases) the number of A -players by one, this initial mutation might also trigger additional changes.

Further, note that in the presence of non-monomorphic states the transition from one monomorphic state to the other can occur via a series of intermediate steps. Which state will be LRE depends on how difficult it is to move from the two monomorphic states, \vec{A} and \vec{B} into the set of non-monomorphic states.⁸ First, consider states where there is only one A -player.

$$\dots BBABB \dots$$

⁸As the non-monomorphic states can be connected to each other and to the monomorphic states via a chain of single mutations which tree will be of minimum cost will be determined by how difficult it is to escape the monomorphic states. The next section elaborates on this in more detail.

As the A -player has no A -neighbors she will switch to B with positive probability if $m^A(c, 1) \geq 0$ which translates into $c \leq 1 - \beta$. In this case lonesome A -players will disappear. However, states with two adjacent A -players are absorbing. Conversely, if $c > 1 - \beta$ holds, the A -player will keep her strategy and states with lonesome A -players are absorbing. Likewise, consider the case when there is a lonesome B -player.

...AABAA...

The B -player has two A -neighbors and will switch strategies with positive probability provided that $m^B(c, 1) \leq 2$, which can be rewritten as $c \leq \alpha$. However, whenever $c > \alpha$ a lonesome B -player will remain. Note, by risk dominance of A , $\alpha > 1 - \beta$. This implies that whenever lonesome B -players will keep their strategy, lonesome A -player will do the same.

Summarizing, if $c \leq \frac{\alpha + \beta - 1}{2}$, only the monomorphic states are absorbing and A can spread out contagiously. Thus, \vec{A} is unique LRE. If $\frac{\alpha + \beta - 1}{2} < c \leq \alpha$ and $c \leq 1 - \beta$, non-monomorphic states are absorbing and it is possible to move among the non-monomorphic states and from these states to the monomorphic ones via a single mutation chain. It is further possible to move from the two monomorphic states to the set of non-monomorphic states at the cost of two mutations. Thus, one can exhibit A - and B - trees which are of cost smaller than any AB -tree. Hence, \vec{A} and \vec{B} are LRE. If, however, $c > 1 - \beta$, moving from \vec{A} to the set of non-monomorphic states takes two mutations, whereas escaping \vec{B} is possible at the cost of one mutation. Thus, in this case one can exhibit A -trees which are of minimum cost, implying that \vec{A} is unique LRE. If $c > \alpha$, all absorbing states are accessible from each other via a chain of single mutations, implying that all of them are LRE.

Thus, the presence of switching costs may imply that the prediction for the long run is altered. The essential mechanism that underlies this result is that switching costs may stop the contagious spread of the risk dominant strategy. Without contagion the question of which equilibrium will emerge in the long run boils down to how difficult it is to access the set of non-monomorphic states from the two conventions. Risk dominance only implies that it can never be easier to move out of the risk dominant convention than it is to move out of the non-risk dominant convention. This, in turn, implies that the risk dominant convention is always contained in the set of LRE. However, it might not be the unique prediction. In particular, there may exist a parameter range where both conventions can be left with two

mutations and, thus, are both LRE.

Whether it is actually possible that a non-risk dominant convention is LRE does not only depend on the level of switching costs but also on the parameters of the underlying game. To see this point note that both monomorphic states are LRE if $c > \frac{\alpha+\beta-1}{2}$ and $c \leq 1-\beta$. It, thus, has to be the case that $\frac{\alpha+\beta-1}{2} < 1-\beta$. This translates into $\alpha+3\beta < 3$. This condition is fulfilled if the advantage of strategy A over B is not too large, but per se is not related to payoff dominance or risk dominance.⁹ Importantly, it may hold if $\alpha > 1$. Thus, even if action A is risk- and payoff- dominant, it might not be unique LRE. We illustrate the set of LRE depending on the level of switching cost in this case in Figure 1. It is interesting to note that the prediction is “non-monotonic” in the level of switching costs. With increasing switching costs the prediction switches from \vec{A} to $\vec{A} \cup \vec{B}$ back to \vec{A} and finally to $\vec{A} \cup \vec{B} \cup G$ in games with $\alpha+3\beta < 3$.




Figure 1: LRE under two player interaction with switching costs and $\alpha+3\beta < 3$.

3.2 $2k$ -neighbor interaction and global interactions

We will now generalize the insights of the two player interaction model to $2k$ -neighbor interaction. We show that we can expect similar phenomena as in the simple two-neighbor model for small interaction neighborhoods. However, as the the size of the interaction neighborhood, k , increases switching costs do no longer influence the prediction, with the exception of very high levels of switching costs, where in the absence of noise no player would switch regardless of the distribution of strategies in her neighborhood. The following lemma provides a characterization of the set of absorbing states.

Lemma 1. *For positive switching costs, $c > 0$,*

i) there are no non-singleton absorbing sets.

ii) the only absorbing states are \vec{A} , \vec{B} , and G .

⁹If strategy A is sufficiently advantageous compared to B , $\alpha+3\beta > 3$, it will be uniquely selected up to the point where $c > \alpha$ (where all absorbing states are LRE.)

Proof. To prove the first part consider an absorbing set W . Consider a state $\tilde{s} \in W$ where the number of A -players is maximal. Let \tilde{m} be the number of A -players at this state. It follows that at this state there does not exist a B -player who, when given revision opportunity, switches to A with positive probability. Thus, $m_i < m^B(c, k)$ for all i with $s_i = B$. If it is the case that $m_j > m^A(c, k)$ for all j with $s_j = A$, then \tilde{s} is the only state in W . If $m_j \leq m^A(c, k)$ for some players j with $s_j = A$, we proceed in the following manner. With positive probability, one of these agents receives revision opportunity and switches to B . We reach a new state s' . At this new state there are strictly fewer A -players. Provided that $c > 0$ for the new B -player we have $m_j \leq m^A(c, k) < m^B(c, k)$, implying that she will not switch back. For all old B -players it is still true that $m_i < m^B(c, k)$, implying that none of them will switch. If there is no A -player with $m_j \geq m^A(c, k)$ left, the state s' is absorbing (contradicting that $\tilde{s} \in W$). If there are still such A -players left, we iterate the procedure until we reach an absorbing state, eventually contradicting the assumption $\tilde{s} \in W$.

The second part follows from the definition of \vec{A} , \vec{B} , and G . □

With the help of this lemma we are able to provide the following result.¹⁰

Proposition 2. *In the $2k$ -neighbor interaction model,*

- a) if $c \leq \frac{\alpha+\beta-1}{2}$ and $N > k(k+1)$, then $\mathcal{S} = \{\vec{A}\}$,
- b) if $\frac{\alpha+\beta-1}{2} < c \leq \alpha$ and
 - i) if $\lfloor m^A(c, k) \rfloor = \lfloor 2k - m^B(c, k) \rfloor$, then $\mathcal{S} = \{\vec{A}, \vec{B}\}$
 - ii) if $\lfloor m^A(c, k) \rfloor < \lfloor 2k - m^B(c, k) \rfloor$, then $\mathcal{S} = \{\vec{A}\}$, and
- c) if $c > \alpha$, then $\mathcal{S} = \{\vec{A}, \vec{B}\} \cup G$.

Proof. For part a) note if $c \leq \frac{\alpha+\beta-1}{2}$, one has $m^B(c, k) \leq k$, implying that a B -player switches to A with positive probability whenever half (or more) of her $2k$ -neighbors choose A . Thus, A may spread contagiously and we are back in the model outlined by Ellison (1993), where $\mathcal{S} = \{\vec{A}\}$ if $N > k(k+1)$.¹¹

¹⁰In the following we denote by $\lfloor x \rfloor$ the largest integer not greater than x and by $\lceil x \rceil$ the smallest integer not less than x .

¹¹Note that we have a model with positive inertia whereas Ellisons model features strategy adjustment in each round. See Weidenholzer (2010) for a discussion of the model with inertia.

We now consider the case where $c > \frac{\alpha + \beta - 1}{2}$. Here we have $m^B(c, k) > k$. Thus, B -players will no longer switch if they have half of their neighbors playing A . This implies A can no longer spread out contagiously. Further, non-monomorphic states are now absorbing, meaning that the set G is non-empty.

We next show that it is possible to move from an absorbing state $AB \in G_\ell$ to either a state in G_a or in G_b , with $a < \ell < b$ at the cost of one mutation. We will show that there exists an A - (and a B -player) such that if she mutates to B (to A), she will not switch back and no other player will switch to A (to B). By the definition of G_ℓ we have $m_i > m^A(c, k)$ for all i with $s_i = A$ and $m_j < m^B(c, k)$ for all j with $s_j = B$. Consider now an A -player i whose adjacent neighbor j is playing B . As they are direct neighbors they have only one player who is not a joint neighbor. Call i 's disjoint neighbor \tilde{i} and j 's disjoint neighbor \tilde{j} . Further j also faces i who is an A -player. It follows that j faces either the same number of A -neighbors as i (if $s_{\tilde{i}} = A$ and $s_{\tilde{j}} = B$), has one more A -neighbors than i (if $s_{\tilde{i}} = s_{\tilde{j}}$), or two more A -neighbors (if $s_{\tilde{i}} = B$ and $s_{\tilde{j}} = A$). Thus, $m_j \in \{m_i, m_i + 1, m_i + 2\}$. Assume that j mutates to A . Since $m_j \geq m_i > m^A(c, k)$ she will not switch back. Further, as there are now more A -players, none of the old A -players will switch, showing that we will reach a state G_b with $b > \ell$. An analogous argument can be used to show that it is also possible with one mutation to move to a state G_a with $a < \ell$.

Now consider \vec{B} . We want to find the minimum number of mutations required for a transition from \vec{B} to a state in the set G . Let $C(\vec{B}, AB)$ denote this number. Recall that $m^A(c, k)$ is defined such that if a player has strictly more than $m^A(c, k)$ A -neighbors, she will strictly prefer to stay at A . If $m^A(c, k) < 0$, we have that an A -player remains even if she does not have an A neighbor. Thus, if $m^A(c, k) < 0$, one mutation is enough to move from \vec{B} to a state in G_1 . Now consider $m^A(c, k) \geq 0$. First, consider the case where $m^A(c, k) \notin \mathbb{Z}$ (where \mathbb{Z} denotes the integers). In this case, if $\lceil m^A(c, k) \rceil + 1$ adjacent players mutate to A each of them will have $\lceil m^A(c, k) \rceil > m^A(c, k)$ neighbors choosing B . Thus, none of them will switch and we have reached an absorbing state in the set $G_{\lceil m^A(c, k) \rceil + 1}$. Note that if less than $\lceil m^A(c, k) \rceil + 1$ players switch to A , all of them will switch back when given revision opportunity. It follows that $C(\vec{B}, AB) = \max\{\lceil m^A(c, k) \rceil, 0\} + 1$ for $m^A(c, k) \notin \mathbb{Z}$. Now consider $m^A(c, k) \in \mathbb{Z}$. In this case for all A players to stay with probability one each of them needs strictly more than $m^A(c, k)$ A -neighbors. Thus, if $m^A(c, k) + 2$ players switch to A , each of them will have $m^A(c, k) + 1$ neighbors playing A and will not switch back

with positive probability. Thus, $C(\vec{B}, AB) = \max\{m^A(c, k) + 1, 0\} + 1$ for $m^A(c, k) \in \mathbb{Z}$. Summing up, we have

$$C(\vec{B}, AB) = \begin{cases} \max\{\lceil m^A(c, k) \rceil, 0\} + 1, & \text{if } m^A(c, k) \notin \mathbb{Z} \\ \max\{m^A(c, k) + 1, 0\} + 1, & \text{if } m^A(c, k) \in \mathbb{Z} \end{cases}.$$

This can be written as $C(\vec{B}, AB) = \max\{\lfloor m^A(c, k) \rfloor + 1, 0\} + 1$.

Conversely, consider the convention \vec{A} . We aim to understand how many mutations to B we need so that the new B -players will keep their strategy with certainty. If $m^B(c, k) > 2k$, this would be the case even if all neighbors choose A . Thus, one mutation is enough to move from \vec{A} to a state in G_{N-1} whenever $m^B(c, k) > 2k$. Assume $m^B(c, k) \leq 2k$. Now a B -player will keep her strategy whenever $m_i < m^B(c, k)$. Initially the B -players had $2k$ A -neighbors. Thus, each of them needs strictly more than $2k - m^B(c, k)$ of their neighbors to play B to keep their strategy with probability 1. Again, let us distinguish the cases $2k - m^B(c, k) \in \mathbb{Z}$ and $2k - m^B(c, k) \notin \mathbb{Z}$. In the latter case with $\lceil 2k - m^B(c, k) \rceil + 1$ mutations one can move from \vec{A} to a state in the set $G_{N - \lceil 2k - m^B(c, k) \rceil - 1}$. Thus, $C(\vec{A}, AB) = \max\{\lceil 2k - m^B(c, k) \rceil, 0\} + 1$. If $2k - m^B(c, k) \in \mathbb{Z}$, we need $2k - m^B(c, k) + 2$ mutations to ensure that each B player has more than $2k - m^B(c, k)$ neighbors playing B . As above, the cases $2k - m^B(c, k) \in \mathbb{Z}$ and $2k - m^B(c, k) \notin \mathbb{Z}$ can be unified by using $C(\vec{A}, AB) = \max\{\lfloor 2k - m^B(c, k) \rfloor + 1, 0\} + 1$.

Finally, let us determine the set of LRE. Let L denote the number of non-monomorphic absorbing states. Thus, together with the states \vec{A} and \vec{B} there are $L + 2$ absorbing states. We can connect all L AB states to each other and to \vec{A} and \vec{B} via a chain of single mutations. Further, we can move from \vec{B} into the class of AB states at the cost of $C(\vec{B}, AB)$. Thus, we can exhibit minimum A -trees of cost $L + C(\vec{B}, AB)$. Likewise, the minimum B -trees have cost $L + C(\vec{A}, AB)$. Further, for each state $AB \in G$ we can exhibit a minimum cost tree of cost $L - 1 + C(\vec{A}, AB) + C(\vec{B}, AB)$.

First note that if $c > \alpha$, we have $m^A(c, k) < 0$ and $m^B(c, k) > 2k$. It follows $C(\vec{A}, AB) = C(\vec{B}, AB) = 1$. Thus, the minimum cost \vec{A} -, the \vec{B} -, and all minimum cost AB -trees have cost $L + 1$. Thus, $\mathcal{S} = \{\vec{A}, \vec{B}\} \cup G$.

Now, consider $\frac{\alpha + \beta - 1}{2} < c \leq \alpha$. Observe that $\lfloor 2k - m^B(c, k) \rfloor = \lfloor 2k(1 - 2q^*) + m^A(c, k) \rfloor \geq \lfloor m^A(c, k) \rfloor$. Thus, $C(\vec{A}, AB) \geq C(\vec{B}, AB)$. So, we either have $C(\vec{A}, AB) > C(\vec{B}, AB)$ in which case $\mathcal{S} = \vec{A}$ or $C(\vec{A}, AB) = C(\vec{B}, AB)$ in which case $\mathcal{S} = \vec{A} \cup \vec{B}$. \square

Thus, the presence of switching costs may imply that under local interactions the risk dominant convention is no longer unique LRE. Let us provide some technical intuition for this result. First, if $c \leq \frac{\alpha+\beta-1}{2}$, the risk dominant strategy may still spread contagiously and, thus, remains unique LRE. For $\frac{\alpha+\beta-1}{2} < c \leq \alpha$ there exist absorbing AB states. Whether the risk dominant or the payoff dominant convention is LRE boils down to the question from which of the two conventions it is more difficult to move to the set of AB -states. This is measured by the numbers $C(\vec{A}, AB)$ and $C(\vec{B}, AB)$ which are in turn rounded down values of the functions $2k - m^B(c, k) + 2$ and $m^A(c, k) + 2$. Risk dominance implies that $C(\vec{B}, AB) \leq C(\vec{A}, AB)$. Thus, the risk dominant convention is always contained in the set of LRE. The functions $2k - m^B(c, k) + 2$ and $m^A(c, k) + 2$ only differ by a constant and are linearly decreasing in the switching costs. It may very well be the case that the rounded down values are the same, $C(\vec{A}, AB) = C(\vec{B}, AB)$. In this case both conventions turn out to be LRE. Finally, for $c > \alpha$ we have that agents will not switch strategies, no matter what the distribution of strategies among their neighbors is and all absorbing states turn out to be LRE.

In Figure 2 we plot the transition costs from either convention to the set of non-monomorphic states as a function of the switching costs. Whenever $C(\vec{A}, AB)$ lies above $C(\vec{B}, AB)$ the convention \vec{A} is unique LRE. When $C(\vec{A}, AB)$ and $C(\vec{B}, AB)$ coincide both conventions, \vec{A} and \vec{B} , are LRE. When the two functions are equal to one, both conventions, \vec{A} and \vec{B} , and the set of non-monomorphic states G are LRE. Note that as in the two player interaction case the prediction is non-monotonic in the level of switching costs. In particular, the prediction that the risk dominant convention is unique LRE and the prediction that both of them are LRE alternate k -times.

The following corollary explores the circumstances under which switching costs may influence the set of LRE. In case switching cost may change the prediction, it shows that the prediction will be non-monotonic as switching costs vary.

Corollary 3. *If $\frac{\alpha+\beta-1}{2} < c \leq \alpha$ and*

a) *if $2k(1 - 2q^*) \geq 1$, then $\mathcal{S} = \vec{A}$*

b) *if $2k(1 - 2q^*) < 1$, then there exist thresholds $\bar{c}^{k+1} < \underline{c}^k < \bar{c}^k < \underline{c}^{k-1} < \bar{c}^{k-1} < \dots < \underline{c}^1 < \bar{c}^1$ (with $\bar{c}^{k+1} = \frac{\alpha+\beta-1}{2}$ and $\bar{c}^1 = \alpha$) such that if $c \in (\bar{c}^{\ell+1}, \underline{c}^\ell]$ for $\ell = 1, 2, \dots, k$, then $\mathcal{S} = \vec{A} \cup \vec{B}$ and $\mathcal{S} = \vec{A}$ otherwise.*




Figure 2: LRE in the game $[\alpha, \beta] = [1.1, 0.1]$ with interaction radius $k = 3$. The solid line plots the transition costs $C(\vec{A}, AB)$ and the dashed line plots the transition costs $C(\vec{B}, AB)$. Whenever $C(\vec{A}, AB)$ lies above $C(\vec{B}, AB)$ the convention \vec{A} is unique LRE. When $C(\vec{A}, AB)$ and $C(\vec{B}, AB)$ coincide both conventions, \vec{A} and \vec{B} , are LRE. When the two functions are equal to 1, both conventions, \vec{A} and \vec{B} , and the set of non-monomorphic states G are LRE.

Proof. Consider case bii) in the previous Proposition. First, note that $2k - m^B(c, k) = m^A(c, k) + 2k(1 - 2q^*)$. Thus, the functions $m^A(c, k)$ and $2k - m^B(c, k)$ only differ by the constant $2k(1 - 2q^*)$. Risk dominance implies $2k(1 - 2q^*) > 0$. Further, note that $m^A(c, k)$ (and thus also $2k - m^B(c, k)$) is linearly decreasing in c .

Consider part a). Note if $2k(1 - 2q^*) \geq 1$, then $\lfloor 2k - m^B(c, k) \rfloor > \lfloor m^A(c, k) \rfloor$.

Now consider part b). Let \underline{c}^ℓ be the value of switching costs c that solves $m^A(c, k) + 1 = \ell$. Note that $\lfloor m^A(c, k) + 1 \rfloor = \ell$ for $\underline{c}^{\ell+1} < c \leq \underline{c}^\ell$. Likewise, define \bar{c}^ℓ to be the value of switching costs c for which $2k - m^B(c, k) + 1 = \ell$. We have $\lfloor 2k - m^B(c, k) + 1 \rfloor = \ell$ for $\bar{c}^{\ell+1} < c \leq \bar{c}^\ell$.

As $2k - m^B(c, k) = m^A(c, k) + 2k(1 - 2q^*) > m^A(c, k)$ and $m^A(c, k)$ is decreasing in c it follows that $\underline{c}^\ell < \bar{c}^\ell$. Further, note that for $2k(1 - 2q^*) < 1$ one has $m^A(c, k) + 1 < m^A(c, k) + 2k(1 - 2q^*) + 1 < m^A(c, k) + 2$. Thus, $\underline{c}^\ell < \bar{c}^\ell < \underline{c}^{\ell-1}$. The last two observations imply $\bar{c}^{k+1} < \underline{c}^k < \bar{c}^k < \underline{c}^{k-1} < \bar{c}^{k-1} < \dots < \underline{c}^1 < \bar{c}^1$. Now note that $\lfloor m^A(c, k) + 1 \rfloor = \lfloor 2k - m^B(c, k) + 1 \rfloor = \ell$ if $c \in (\underline{c}^{\ell+1}, \underline{c}^\ell]$ and $c \in (\bar{c}^{\ell+1}, \bar{c}^\ell]$. This is the case for $c \in (\bar{c}^{\ell+1}, \underline{c}^\ell]$. On the contrary, if $c \in (\underline{c}^\ell, \bar{c}^{\ell-1}]$, then $\lfloor m^A(c, k) + 1 \rfloor = \ell - 1 < \ell = \lfloor 2k - m^B(c, k) + 1 \rfloor$. \square

The main idea behind this corollary is that the functions $2k - m^B(c, k) + 2$ and $m^A(c, k) + 2$ only differ by the constant $2k(1 - 2q^*)$. If this constant is greater than or equal to one, we

will have $C(\vec{B}, AB) = \lfloor m^A(c, k) + 2 \rfloor < C(\vec{A}, AB) = \lfloor 2k - m^B(c, k) + 2 \rfloor$, regardless the level of switching costs. If, however, this constant is smaller than one, there exist levels of switching costs for which $\lfloor 2k - m^B(c, k) \rfloor = \lfloor m^A(c, k) \rfloor$. Further, note that if there exists a range of switching costs for which, e.g. $\lfloor 2k - m^B(c, k) \rfloor = \lfloor m^A(c, k) \rfloor = 1$, then due to the stepwise nature of the floor function, there also exists a range of switching costs for which $\lfloor 2k - m^B(c, k) + 2 \rfloor = \lfloor m^A(c, k) + 2 \rfloor = r$ for every $r \in \mathbb{Z}$. Thus, the prediction that the convention \vec{A} is unique LRE and the prediction that both conventions, \vec{A} and \vec{B} , are LRE alternate k times as c increases (and $C(\vec{B}, AB)$ and $C(\vec{A}, AB)$ decrease from $k + 1$ to 1). Finally, it is interesting to note that since $\underline{c}^1 = 1 - \beta$ and $\bar{c}^1 = \alpha$, for $c \in (1 - \beta, \alpha]$ one has $\mathcal{S} = \vec{A}$. Further, if $c > \alpha$, one has $\mathcal{S} = \{\vec{A}, \vec{B}\} \cup G$. Hence, just before the model's prediction includes all absorbing states it uniquely selects the risk dominant convention.

A straightforward implication of the first part of the corollary is that if agents interact with sufficiently many other agents (k large) or if the risk dominant action has a relatively large basin of attraction (q small), switching costs do not influence the prediction. The second part of the corollary implies that if agents interact only with a few other agents (k small) and/or the risk dominant action's basin of attraction is relatively small (q close to $\frac{1}{2}$), then the prediction may not be unique and moreover is non-monotonic in the level of switching costs.

Finally, note that it is possible to reconcile our findings with the results of Norman (2009) by simply setting $k = \frac{N-1}{2}$, thus, obtaining a model of global interactions. For small populations switching costs may very well have an impact on the set of LRE. However, in large populations, as considered by Norman (2009), the prediction is robust to switching costs. In this case, switching costs speed up convergence but do not alter the long run behavior of the population.

4 Conclusion

We have established that under local interactions the set of LRE may be altered by the presence of switching costs. In particular, the risk dominant convention may no longer be unique LRE. If, however, agents interact with sufficiently many other agents our critique does not apply and risk dominant conventions are still uniquely selected.

One question that immediately comes to mind is whether our findings hold in a more

general context. In the context of this paper switching costs played the following role: i) Under switching costs non-monomorphic states may become absorbing. ii) Switching costs may change the transition costs, measured in the number of required mistakes, with which these non-monomorphic states can be accessed from the monomorphic ones. Rounding up, when calculating switching costs, may lead to the effects outlined in this paper. If the number of required mistakes is relatively small, the effect of rounding up will be most pronounced. However, for a large number of required mistakes these effects will be most likely negligible. We, thus, conjecture that switching costs will play a similar role in models where only a relatively small number of mutations is needed to move from one convention to another. There are two natural dimensions along which our results might be generalized. First, we expect switching costs to impact the long run prediction in the circular city model of local interactions if we move beyond the class of 2×2 -coordination games (as in e.g. Ellison 2000, Alós-Ferrer and Weidenholzer 2007). Secondly, switching costs will also influence the prediction in models where the way in which agents interact with each other implies that only few mistakes are necessary to move among conventions. Examples include the torus model outlined in Ellison (2000), multiple location models (as in Anwar 2002, Ely 2002, Blume and Temzelides 2003, Shi 2014, Alós-Ferrer and Kirchsteiger 2010), network formation models under asynchronous adjustments of links and actions (see Jackson and Watts 2002) or under constrained interactions (as in Staudigl and Weidenholzer 2014).

Admittedly, the integer problem that is driving our results is an artefact of the uniform noise approach. While other learning models such as the logit dynamics as advocated by e.g. Blume (1993, 1995) do not face this problem, their predictions sometimes may depend on other specifics such as the timing of revision opportunities or tie breaking assumptions (see Alós-Ferrer and Netzer 2014 for a discussion).¹²

Acknowledgments

We are indebted to an associate editor, two anonymous referees, Carlos Alós Ferrer, Roy Bailey, Zhiwei Cui, Andrea Galeotti, David Gill, Erik Mohlin, Thomas Norman, Peyton Young, and to audiences at the University of Oxford, the 16th International Symposium

¹²Maes and Nax (2014) experimentally assess whether uniform- or logit- noise provides a better fit to actual deviation from best response learning. Noise levels turn out to decrease with the payoff loss implied by a deviation, providing support for logit noise. However, for low implied deviation costs (as in the interesting parameter range of our paper) a constant error model may offer a better fit than logit noise.

on Dynamic Games and Applications in Amsterdam and the 2014 annual meeting of the Austrian Economic Association in Vienna for inspiring discussions and for helpful comments and suggestions. Financial support from the Vienna Science and Technology Fund (WWTF) under project fund MA 09-017 is gratefully acknowledged.

References

- ALÓS-FERRER, C. (2004): “Cournot vs. Walras in Oligopoly Models with Memory,” *International Journal of Industrial Organization*, 22, 193–217.
- ALÓS-FERRER, C., AND G. KIRCHSTEIGER (2010): “General Equilibrium and the Emergence of (Non) Market Clearing Trading Institutions,” *Economic Theory*, 44, 339–360.
- ALÓS-FERRER, C., AND N. NETZER (2014): “Robust stochastic stability,” *Economic Theory*, pp. 1–27.
- ALÓS-FERRER, C., AND F. SHI (2012): “Imitation with asymmetric memory,” *Economic Theory*, 49(1), 193–215.
- ALÓS-FERRER, C., AND S. WEIDENHOLZER (2006): “Imitation, Local Interactions, and Efficiency,” *Economics Letters*, 93, 163–168.
- (2007): “Partial Bandwagon Effects and Local Interactions,” *Games and Economic Behavior*, 61, 1–19.
- (2008): “Contagion and Efficiency,” *Journal of Economic Theory*, Volume 143, 251–274.
- ANWAR, A. W. (2002): “On the Co-existence of Conventions,” *Journal of Economic Theory*, 107, 145–155.
- APESTEGUIA, J., S. HUCK, J. OECHSSLER, AND S. WEIDENHOLZER (2010): “Imitation and the Evolution of Walrasian behavior: Theoretically Fragile but Behaviorally Robust,” *Journal of Economic Theory*, 145(5), 1603–1617.
- BERGIN, J., AND B. L. LIPMAN (1996): “Evolution with State-Dependent Mutations,” *Econometrica*, 64, 943–956.

- BLUME, A., AND T. TEMZELIDES (2003): “On the Geography of Conventions,” *Economic Theory*, 22, 863–873.
- BLUME, L. (1993): “The Statistical Mechanics of Strategic Interaction,” *Games and Economic Behavior*, 5, 387–424.
- (1995): “The Statistical Mechanics of Best-Response Strategy Revision,” *Games and Economic Behavior*, 11, 111–145.
- ELLISON, G. (1993): “Learning, Local Interaction, and Coordination,” *Econometrica*, 61, 1047–1071.
- (2000): “Basins of Attraction, Long-Run Stochastic Stability, and the Speed of Step-by-Step Evolution,” *Review of Economic Studies*, 67, 17–45.
- ELY, J. C. (2002): “Local Conventions,” *Advances in Theoretical Economics*, 2, 1–30.
- ESHEL, I., L. SAMUELSON, AND A. SHAKED (1998): “Altruists, Egoists, and Hooligans in a Local Interaction Model,” *The American Economic Review*, 88, 157–179.
- FOSTER, D., AND P. YOUNG (1990): “Stochastic evolutionary game dynamics,” *Theoretical population biology*, 38(2), 219–232.
- FREIDLIN, M., AND A. WENTZELL (1988): *Random Perturbations of Dynamical Systems*, 2nd Ed. Springer Verlag, New York.
- FUDENBERG, D., AND D. LEVINE (1998): *The Theory of Learning in Games*. The MIT Press, Cambridge, Massachusetts.
- HARSANYI, J., AND R. SELTEN (1988): *A General Theory of Equilibrium Selection in Games*. The MIT Press, Cambridge, MA.
- JACKSON, M. O., AND A. WATTS (2002): “On the Formation of Interaction Networks in Social Coordination Games,” *Games and Economic Behavior*, 41, 265–291.
- KANDORI, M., G. J. MAILATH, AND R. ROB (1993): “Learning, Mutation, and Long Run Equilibria in Games,” *Econometrica*, 61, 29–56.
- KIM, C., AND K. WONG (2010): “Long-run equilibria with dominated strategies,” *Games and Economic Behavior*, 68, 242–254.

- LEE, I. H., A. SZEIDL, AND A. VALENTINYI (2003): “Contagion and State Dependent Mutations,” *The B.E. Journal of Theoretical Economics*, 0(1).
- MAES, M., AND H. H. NAX (2014): “A Behavioral Study of ‘Noise’ in Coordination Games,” *Available at SSRN 2521119*.
- MENGEL, F. (2009): “Conformism and cooperation in a local interaction model,” *Journal of Evolutionary Economics*, 19(3), 397–415.
- NORMAN, T. W. (2009): “Rapid evolution under inertia,” *Games and Economic Behavior*, 66(2), 865–879.
- ROBSON, A. J., AND F. VEGA-REDONDO (1996): “Efficient Equilibrium Selection in Evolutionary Games with Random Matching,” *Journal of Economic Theory*, 70, 65–92.
- SAMUELSON, L. (1997): *Evolutionary Games and Equilibrium Selection*. The MIT Press, Cambridge, Massachusetts.
- SHI, F. (2014): “Long-run technology choice with endogenous local capacity,” *Economic Theory*, pp. 1–23.
- STAUDIGL, M., AND S. WEIDENHOLZER (2014): “Constrained interactions and social coordination,” *Journal of Economic Theory*, 152, 41–63.
- TANAKA, Y. (1999): “Long Run Equilibria in an Asymmetric Oligopoly,” *Economic Theory*, 14, 705–715.
- VEGA-REDONDO, F. (1997): “The Evolution of Walrasian Behavior,” *Econometrica*, 65, 375–384.
- WEIDENHOLZER, S. (2010): “Coordination Games and Local Interactions: A Survey of the Game Theoretic Literature,” *Games*, 1(4), 551–585.
- (2012): “Long-run equilibria, dominated strategies, and local interactions,” *Games and Economic Behavior*, 75(2), 1014–1024.
- YOUNG, P. (1993): “The Evolution of Conventions,” *Econometrica*, 61, 57–84.