

**Endogenous Coordination and Discoordination Games:†
Multiculturalism and Assimilation**

John P. Conley*

and

William S. Neilson**

† **Published as:** John P. Conley and William Neilson (2009) “Endogenous Games and Equilibrium Adaptation of Social Norms and Ethical Constraints” *Games and Economic Behavior*, Vol 6, pp. 761-774. We wish to thank Tom Gresik, Ehud Kalai, Herve Moulin, Simon Wilkie, and the participants of PET07, Vanderbilt for their comments and suggestions. We take responsibility for all remaining errors.

* Department of Economics, Vanderbilt University, j.p.conley@vanderbilt.edu

** Department of Economics, University of Tennessee, wneilson@utk.edu

Abstract

We argue that many strategic interactions are chosen by, rather than imposed upon, agents. The endogeneity of the implicit game substantially alters the incentive structure faced by agents. If agents can refuse to play a game and instead seek a different partner, then it becomes in their interest to be the sort of player with whom a high quality partner will agree to play. Formally, we explore games which are endogenous in the sense that players first choose to learn a specific set of strategies to play with future partners, and then enter a matching process to find a suitable partner. We find that in coordination games, the only equilibria involve all agents learning the exact same single strategy and playing this with the first partner they meet. Applying this to culture and language, our results imply that if there are benefits from interacting with similar agents then there is a strong bias towards assimilation and the emergence of a single culture (that is, all agents use the endogenous game structure to learn or adopt only a single common culture or language). This result can be overturned, however, by such factors as heritage, cultural affinities, caste, and government policies. Finally, we show that if agents benefit from interacting with agents dissimilar from themselves (we call this a “discoordination game”), not only will different cultures coexist, but also there must necessarily arise a class of “cultural market makers” who specialize in bridging the cultural divide.

1. Introduction

In a typical coordination game, two players must simultaneously choose from identical action sets without prior communication. If the players happen to choose the same action, they both get large payoffs. If they choose different actions, they both get low payoffs. Coordination games have multiple equilibria, and as such, are useful in exploring equilibrium selection issues. Since all of these equilibria require identical behavior by all agents, this, in turn, informs questions about how seemingly arbitrary social norms are established and persist. Equilibrium selection is commonly approached through experiments to study how actual subjects solve the problem, or through evolutionary games to study how a population evolves through time to arrive at one of the potential equilibria.

In this paper, we take a step back. We ask if the structure of the games used in these experimental and evolutionary treatments of coordination games agrees with the institutions agents experience in the real world. We use a new approach called “endogenous games” first developed in Conley and Neilson (2008) and explore the implications for equilibrium outcomes, and especially the possibility that multiple social norms might be able to coexist in an integrated society instead of one culture always assimilating into the other.

In a typical evolutionary approach, members of a large population are randomly matched to play the coordination game.¹ Each member has exactly one action available which may or may not be a best response to the action played by the partner with whom he is matched. Agents who are fortunate enough to match with a similar partner gain higher payoffs than those who do not, and types with higher expected payoffs gain larger population shares in the next period. The evolutionary equilibrium is the long-run, steady-state distribution of types that emerges from this process.

Much of the interest in the various evolutionary approaches to coordination and other games arises from the specific mechanisms that match partners and how those

¹ See, for example, the pioneering papers of Kandori, Mailath, and Rob (1993) and Young (1993).

mechanisms mimic real world societies.² Note that a maintained assumption in all of these matching protocols is that they are “forced” in that two randomly-matched agents must play the game against each other once the pairing occurs. None of the protocols allows agents to choose their own partners. We argue that this is very much at variance with many real world situations and that individuals often choose the set of agents with whom they interact. For example deciding to work with someone, joining a church or social organization or choosing to marry someone is typically a voluntary, strategic decision that logically precedes playing the resulting game between the agents in the group that eventually coalesces from this process. Bargaining games as discussed in Rubinstein and Wolinsky (1985) are an example of this type of interaction that has been well studied in the literature.

Since agents would have to be aware that groups form endogenously, we would expect that they would respond to implicit incentives that result. More specifically, since agents want to match with a partner who is “good” in the sense that the payoff received from playing the game together are high. However, good partners are in demand and can accept or reject an offered pairing. Thus, agents who know they will play an endogenously formed game have an incentive to make themselves attractive to good partners. Suppose, for example, I am seeking a business partner. In the abstract, I could choose to work honestly or try to cheat my partner as much as possible. Both are strategies that are physically available to all agents. However, if I could somehow exclude “cheat my partner” from my personal set of available strategies, then I could attract a similar partner in the choosing phase. Since two honest partners make more profit than two cheating partners, all would benefit from this exclusion. Thus, in a formal sense, we suggest that agents consider the whole set of strategies available to them in the game, but then learn to employ only a subset of these with a view to attracting a better partner.

In the real world, agents limit their strategy spaces and signal potential partners

² See, for example, Ellison (1993) and Morris (2000). Sugden (1995), Oechssler (1997), Ely (2002), and Bhaskar and Vega-Redondo (2004) show how manipulating the matching protocol can lead to the coexistence of different norms or conventions.

of these limits in a variety of ways.³ The most obvious is learning. Many types of strategies are unavailable to agents unless they have taken the time to learn them. One cannot speak French without learning it at some point. One cannot follow Japanese cultural norms without being immersed in Japanese culture for some time. It is difficult to embezzle without first gaining a knowledge of accounting, and it is difficult to cheat on one's wife without knowing how to charm women. Alternatively, it may be that agents have ethical codes that prevent them from employing certain strategies, and these ethical codes can be signaled through a religious affiliation or lifestyle. When an agent can credibly signal the boundaries of his strategy set, he is more likely to match with other agents who have the same boundaries.⁴

In this paper, we look at the coordination game in a new way that allows for these two types of endogenous choices on the part of participants. Specifically, in period 0 (which we refer to as the *pregame*) agents choose a subset of all physically available actions as the only ones that they wish to have available to use when they eventually are matched with a partner. We refer to the subset of the action space chosen by an agent as his *list*.⁵ In period 1, agents are randomly matched with a potential partner. Members of matched pairs do not know the full set of lists chosen by the population, but they do observe each others' lists and so decide on the basis of those observations whether to play or pass. If they both elect to play, they choose actions from their respective lists, collect their payoffs, and leave the game.⁶ They are then replaced in

³ Spier (1992) takes a traditional signaling approach in a matching market. In her model agents have two types, good and bad, and in the separating equilibrium good types select less-complete contracts than good types in order to signal their types to their partners. For example, if marriage partners come in stable and unstable types, the stable types would not ask for a prenuptial agreement in order to signal their stability, while an unstable type would ask for such an agreement. In Spier's approach the agents signal their types by taking specific actions. In contrast, in our approach agents restrict their action spaces to signal their resulting types.

⁴ These limited strategy sets may also be considered internal moral constraints as described by Stringham (2011).

⁵ Calabuig and Olcina (2009) also allow for a kind of "type" choice in the coordination game. In their model agents live for two overlapping generations, and parents choose what preferences to teach to their children.

⁶ Ghosh and Ray (1996) and Rob and Yang (2010) consider the opposite type of player choice, where players are first matched randomly but can choose to terminate relationships. In their papers the

the population by players of the same “type,” that is, with the same list (which is similar to the evolutionary approach). If, on the other hand, at least one member of the matched pair elects to pass, they both pay a waiting cost and are randomly rematched in the next period. This process continues forever. Since both agents’ types and coalitions of agents who eventually play together are endogenously determined, we refer this as an *endogenous game*.⁷

Endogenous games were introduced in Conley and Neilson (2008) to provide a rationale for observed cooperative behavior in the prisoner’s dilemma.⁸ We find all agents choosing to learn only to cooperate and all agents learning only how to cheat are both equilibria of the endogenous prisoner’s dilemma. If we add partially-cooperative actions, then (under mild conditions) all equilibria entail every agent choosing to learn the same least-cooperative strategy. That is, the equilibria look like “social norm” outcomes in which agents all know how to “cheat” but only to exactly the same degree. If an agent decides to defect by learning how to use a less-cooperative strategy than this social norm, he would never find a partner and so his greater skill at exploiting partners could never be used to any advantage. Thus, the set equilibria of the prisoner’s dilemma in the endogenous game setting is a superset of the equilibrium outcomes for the one-shot version but a subset of the equilibrium of the repeated version. The paper also looks at an anti-prisoner’s dilemma game (where the dominant strategy equilibrium is also Pareto efficient) and derives similar results. Any social norm (that

presence of myopic agents allows others to enforce cooperation through the threat of termination. Too few myopic agents causes this equilibrium to break down, similar to low unemployment rates causing efficiency wage equilibria to break down. Watson (1999) also allows players to exit relationships endogenously, and his model focuses on how players learn about each other over time. Our approach differs from theirs in that we allow agents to choose to interact in the first place, and agents play only once so there is no termination.

⁷ Jackson and Wilkie (2005) also consider an endogenous game, but their approach is orthogonal to ours. We assume that the payoff function is fixed but that the players and action spaces are endogenous, while Jackson and Wilkie assume that the the players and action spaces are fixed but the payoff functions are endogenous. They achieve this endogeneity by allowing players to write binding contracts that change the payoff function before play begins. The focus of their paper is on the efficiency of outcomes and, as with Conley and Neilson (2008), they find that homogeneity need not lead to efficiency.

⁸ Mengel (2008) manipulates a matching protocol in an evolutionary version of the prisoner’s dilemma to discuss social norms.

is, outcomes where all agents have learned to play exactly the same strategies) is an equilibrium, regardless of whether the norm is Pareto dominated by the stage-game Nash equilibrium.

These differences between the set of equilibria in the one-shot game, the set of equilibria in the infinitely-repeated game, and the set of equilibria in the endogenous game suggest that an analysis of the endogenous coordination game is worthwhile.

One particularly interesting question is whether or not there exist equilibria in endogenous coordination games in which some agents choose to be one “type” and others choose to be another incompatible “type.” We find this is not possible, however. This contrasts with the one-shot game in which pairs of players must coordinate, but different pairs can coordinate on different actions. Instead, equilibria of the endogenous game involve every agents’ list containing the same common action, and having this common action played when agents are matched.⁹ We will say more about the implications of this below. We also look at discoordination games in which gains from trade occur when agents choose to be different. Again we find that the one-shot and endogenous treatments do not have the same equilibria. In the one-shot, two-player, two-action version of the game there are two pure-strategy Nash equilibria involving agents choosing different strategies with certainty, and one mixed strategy equilibria which is Pareto dominated. In contrast, we find that the only equilibrium in the endogenous game is for some agents to learn only the first strategy, some to learn only the second strategy and for some agents to learn both. The existence of these “social market makers” who learn both strategies is rather surprising. There is no role for market makers in the coordination game and they will not appear in non-endogenous game equilibria in general.

To make the coordination game results concrete and provide motivation, we consider language as an example. Agents can learn English, Spanish, or both. Two agents can observe the set of languages their potential partners are capable of speaking. They

⁹ The appropriate equilibrium concept is pregame perfect equilibrium, which is a refinement of subgame perfect equilibrium and is defined in Section 2.

both gain from trade only if they end up speaking the same language. If the waiting cost is small, someone who speaks only Spanish will choose to interact with a Spanish-only partner or a bilingual partner, but not with an English-only partner (and inversely). The question then is, given this ability to refuse interactions, can there be an equilibrium in which some agents speak only Spanish and some speak only English (as we see in reality), or does equilibrium dictate that all agents learn the same language? Our analysis of the symmetric game shows that everyone must learn the same language, although some may learn to be bilingual.

Since this is obviously counter-factual, we go on to explore how different social conventions might exist together in equilibrium.¹⁰ We note that the game may not be symmetric in the case of language. In particular, agents might prefer interactions made in their native language to those made in a non-native language. We check the robustness of our prior result that equilibrium prescribes a common language in this new setting. We find that the ability for different monolingual groups to coexist depends on the size of the waiting cost, the magnitude of the preference to interact in the native language, and the size of the native group within society. A society can consist of two monolingual groups that do not interact with each other only if the waiting cost is small, the groups are relatively equal in size, and the preference for speaking the native language is strong.

The paper precedes as follows. Section 2 gives a formal definition of the model and defines an equilibrium concept called *pregame perfect equilibrium*, which is an additional contribution beyond what is done in Conley and Neilson (2008). Section 3 looks at symmetric coordination games and establishes that one behavioral norm must dominate. Section 4 explores a discoordination game in which multiple conventions must coexist in equilibrium. Section 5 looks at issues of heritage and preferences for one mode of conduct over another. Section 6 offers some concluding remarks.

¹⁰ Despite the obvious coexistence of multiple languages in the world, Ku and Zussman (2010) provide evidence that a common language, in this case English, promotes trade.

2. The Model

Consider an economy with I agents indexed $i \in \{1, \dots, I\} \equiv \mathcal{I}$. Each agent has a set of strategies \mathcal{X} available for him to learn. Agents receive payoffs by finding a partner with whom to play a bilateral game. The payoff to each agent is given by the function $F : \mathcal{X} \times \mathcal{X} \rightarrow \mathfrak{R}$. Thus, the payoff to agent i choosing strategy x_i playing agent j choosing strategy x_j is $F(x_i, x_j)$.

We assume a two-part game. The first part is the *pregame*. Here, each agent chooses to learn a subset of strategies $\ell_i \subseteq \mathcal{X}$. We refer to ℓ_i as agent i 's *list*. Denote the set of all possible lists as \mathcal{L} (that is, all possible nonempty subsets of \mathcal{X}).

The second part of the game is a multistage matching game similar to the one described by Rubinstein and Wolinsky (1985). In each round of the stage game, agents are randomly matched with another agent and can choose one of two options: play or not play.

If both agents decide to Play (P), they choose strategies from their list, settle on a Nash equilibrium and retire from the game. We assume that when an agent retires, he is replaced by an agent with the same name who has chosen to play in the same way. While this is a strong assumption from a theoretical standpoint, it will become clear that in the applications we consider, it has very little bite.

If at least one agent decides to Not play (N), they receive no payoff but pay a *delay cost* of $z > 0$ and search for a new partner in the next round of play.

Formally, an agent's strategy in the second part of the game is a mapping from the strategy list of the agent with whom he is matched to his own list plus N. Since matching is random and anonymous and the population of agent types is stable across periods, we assume that history does not influence strategic choice.¹¹ Let agent i 's second stage strategy be denoted $g_i : \mathcal{L} \rightarrow \ell_i \cup N$. Let \mathcal{G} denote the set of all such

¹¹ The setup of the game precludes an history-dependent supergame strategies as would normally emerge in an infinitely-repeated game because agents only match a single time and then retire from the game after a single match. Agents who leave cannot be punished or rewarded for past behavior, which is why supergame strategies are not relevant. The setting also precludes a Markov-perfect equilibrium solution because each period begins in exactly the same state according to the assumption that agents who matched in the previous period were replaced by new agents otherwise identical to them.

mappings. We will refer to g_i as *agent i 's stage game strategy*.

An agent's strategy therefore consists of a pair $(\ell_i, g_i) \in (\mathcal{L}, \mathcal{G})$. Collectively, we will denote this as $s_i \in \mathcal{S}$. A *strategy profile* for the game is denoted: $S \equiv (s_1, \dots, s_i) \in \mathcal{S} \times \dots \times \mathcal{S}$. It will be useful to refer to the profile of lists and stage game strategies separately on occasion. We denote these $L = (\ell_1, \dots, \ell_i)$ and $G = (g_1, \dots, g_i)$. We also use the notation s_{-i} to denote the strategy profile for all agents excluding agent i .

The next step is to define a reasonable equilibrium concept for endogenous games. Rubinstein and Wolinsky (1984) employ subgame perfect equilibrium for a game which uses the same matching protocol as we do, but their game does not include a pregame. The existence of a pregame makes subgame perfection insufficient for analyzing our endogenous game. To see why, note that subgame perfection requires that agents play Nash equilibrium strategies in every proper subgame, where a proper subgame is one that begins at a decision node in a singleton information set. With the pregame, players have no singleton information sets. Consider, for example, a game with $I = 4$ agents and the strategy set $\{a, b\}$. Each agent can choose one of three possible lists: $\{a\}$, $\{b\}$, or the full list $\{a, b\}$. Suppose that in period 1 agent 1 is matched with agent 2, and agent 3 is matched with agent 4. Agent 1 observes agent 2's list, but not agent 3's or agent 4's. These two agents' lists are relevant to player 1, though, because if 1 and 2 choose not to play against each other, they are randomly rematched in period 2 in which case agent 1 is likely to face either 3's list or 4's. Because agent 1 does not observe these lists in period 1, though, agent 1's first-period information set contains 9 different decision nodes corresponding to the 9 different combinations of lists chosen by agents 3 and 4.¹² Thus, the game played by matched agents 1 and 2 in period 1 cannot be considered a proper subgame, and subgame perfection has no bite. Instead, an appropriate equilibrium concept must assign beliefs over the nodes in the information sets, and the essence of our equilibrium concept is that agents play sequentially rational strategies based on a reasonable set of consistent beliefs. While straightforward from

¹² In general, with I agents and n strategies in X , after being matched in period 1 an agent's information set contains $(n^2 - 1)^{I-2}$ nodes.

a conceptual standpoint, it is notationally dense to define formally. We therefore will follow the strategy of putting the detailed formal development of these concepts in the mathematical appendix and focus only on the points that are of real economic and theoretical interest.

One requirement of our equilibrium notion is that agents choose strategies in the pregame that maximize their expected payoffs. We will denote the *expected payoff* that agent i receives under strategy profile s as $EPO(s_i, s_{-i})$ (see the appendix for details).

We also require agents to follow optimal strategies in every information set, specifically the information sets that arise once an agent is matched with another. Thus, we will need to know the expected continuation payoff. That is, conditional on having been matched with an agent who has chosen list $\hat{\ell}$, we need to know the expected payoff from this point forward. Denote the set of lists that agent i might encounter during equilibrium play given S by

$$L_i \equiv \{\ell \in \mathcal{L} \mid \exists j \in \mathcal{I}, j \neq i \text{ s.t. } \ell_j = \ell\}$$

and the set of lists that should not be encountered by agent i given S by

$$\bar{L}_i \equiv \{\ell \in \mathcal{L} \mid \nexists j \in \mathcal{I}, j \neq i \text{ s.t. } \ell_j = \ell\}.$$

Similarly, define the set of lists that might be encountered by *any agent* in equilibrium play given S by

$$L \equiv \{\ell \in \mathcal{L} \mid \exists j \in \mathcal{I} \text{ s.t. } \ell_j = \ell\}$$

and the set of lists that should not be encountered by *any agent* in equilibrium given S by

$$\bar{L} \equiv \{\ell \in \mathcal{L} \mid \nexists j \in \mathcal{I} \text{ s.t. } \ell_j = \ell\}.$$

The expected payoff for an information set in which an agent faces an opponent with list $\hat{\ell} \in L_i$, that is, a list agent i expects to encounter with positive probability, is straightforward to compute and denoted by $EPO(s_i, s_{-i} \mid \hat{\ell})$. The expected payoff in an information set in which agent i faces a list that should *not* be encountered on

the equilibrium path, that is, a list $\ell \in \bar{L}_i$, is more problematic. Subgame perfection requires that if agent i and the agent with the deviant list $\ell \in \bar{L}_i$ agree to play, they must play a Nash equilibrium of the stage game. If the stage game has multiple equilibria agent i must act based on beliefs about which of those equilibria will ensue. If agent i chooses not to play against the agent with the deviant list, agent i must recompute the continuation payoffs based on the information that one agent has deviated from the equilibrium path. It could be that this is the only deviation, or it could be that there are multiple deviators. Agent i 's choice between playing or passing against the deviant list depends on his beliefs about what will happen if he chooses to play and what will happen if he chooses to pass.

We now turn to the task of specifying a set of consistent beliefs regarding the stage game strategies of both equilibrium and disequilibrium agents as well as beliefs about expected payoffs when a match is refused. Some of these beliefs are pinned down by subgame perfection, but others, especially regarding the expected payoffs relevant when paired with an agent with an out-of-equilibrium list, are not. The definition below extends the notion of credible beliefs to endogenous games.

We use the notation $p(g|S, \ell)$ to indicate agents' beliefs regarding the probability that a member of the population who has chosen list ℓ uses stage game strategy g given the strategy profile S . Given a strategy profile S , $p : \mathcal{G} \times S \times \dots \times S \times l \rightarrow [0, 1]$ is a *consistent belief system* held in common by all agents if:

- 1 For all $\ell \in \mathcal{L}$, $\sum_{g \in \mathcal{G}} p(g|S, \ell) = 1$ and $p(g|S, \ell) \geq 0$.
- 2 For all $\ell \in L$, $p(g|S, \ell) = \frac{|\{i \in \mathcal{I} \text{ s.t. } g_i = g \text{ and } \ell_i = \ell\}|}{|\{i \in \mathcal{I} \text{ s.t. } \ell_i = \ell\}|}$.
- 3 For all $\ell \in \bar{L}$, $p(g|S, \ell) > 0$ only if there does not exist $\bar{g} \in \mathcal{G}$ such that for all $\hat{\ell}$ such that there exists $i \in \mathcal{I}$ with $\ell_i = \hat{\ell}$.¹³

$$EPO(\bar{g}, \ell, S|\hat{\ell}) > EPO(g, \ell, S|\hat{\ell}).$$

¹³ This is a slight abuse of notation. Rather than define a new *EPO* for disequilibrium agents, we simply pretend the population expanded by 1 to $I + 1$ and use the old definition. Note in this case, S plays the role of s_i .

Condition 1 says p is a probability distribution. Condition 2 says the beliefs are consistent with the equilibrium choice of lists by agents. Condition 3 says for lists not seen in equilibrium, the common belief system about agents with such lists places positive probability weight only on stage game strategies that are payoff maximizing.

Given these beliefs, we will finally need to express the continuation payoff for an agent who encounters a list that he should not see in equilibrium. Note that there is a small complication since the out-of-equilibrium lists may be different for different agents. In particular, if *only* a single agent i chooses a given list ℓ then of course agent i would expect never to see this list in an opponent. The remainder of the population, on the other hand, would just assume it had encountered agent i . We have embedded one assumption in this regard in the definition in consistent beliefs: if an agent i who happens to be the only member of the population who chooses a list ℓ encounters this list anyway, he assumes that that this disequilibrium opponent is just like him and chooses the same stage game strategy, g_i , as he does. This keeps his beliefs consistent with the rest of the agents in the game. It would be more general to allow agent i 's beliefs to be unconstrained, but it would complicate notation and would not change any of the conclusions in the next section. We therefore opt for our simpler approach.

Given that strategies are complete contingent plans, every agent must specify a continuation game strategy for each possible list he could choose. The probability an agent with an out-of-equilibrium list plays stage game strategy g can be inferred from the equilibrium strategy profile S . The expected payoff from choosing to play against an out-of-equilibrium list can then be computed in a straightforward manner (see the appendix for details). As for the expected payoff from choosing to not play, we make a “small deviation” assumption that the out-of-equilibrium list disappears after the current round and is replaced by the appropriate equilibrium list. In essence, this means that agents never expect to face disequilibrium lists in the future, regardless of whether they currently face one or not. Given this is a zero probability event, this seems like a reasonable assumption. It also makes the equilibrium concept very simple to apply, as each agent's expected payoff from refusing to play remains constant through

time and is independent of who the agent is currently matched with.

We denote the expected payoff agents receive when encountering out of equilibrium lists ($\ell \in \bar{L}_i$) by $\overline{EPO}(s_i, s_{-i}|\ell)$. We use a “bar” to remind us that this continuation payoff refers only to out of equilibrium play.

We are now ready to define our equilibrium concept.

Pregame perfection: A strategy profile $s \in S$ and consistent belief system p is a pregame perfect equilibrium (PPE) if

1. For all agents $i \in \mathcal{I}$ and all $\hat{s}_i \in \mathcal{S}$,

$$EPO(s_i, s_{-i}) \geq EPO(\hat{s}_i, s_{-i}).$$

2. For all agents $i \in \mathcal{I}$, all $\ell \in L_i$, and all $\hat{g}_i \in \mathcal{G}$,

$$EPO(s_i, s_{-i}|\ell) \geq EPO(\hat{g}_i, \ell_i, s_{-i}|\ell).$$

3. For all agents $i \in \mathcal{I}$, all $\ell \in \bar{L}_i$, and all $\hat{g}_i \in \mathcal{G}$,

$$\overline{EPO}(s_i, s_{-i}|\ell) \geq \overline{EPO}(\hat{g}_i, \ell_i, s_{-i}|\ell).$$

4. For all agents $i \in \mathcal{I}$ and all $\ell \in \mathcal{L}$ if

$$\max_{x_i \in \ell_i} \min_{x \in \ell} F(x_i, x) \geq EPO(s_i, s_{-i}) - z,$$

then $g_i(\ell) \neq N$.

Condition 1 says that for all agents, the strategies chosen maximize the expected payoffs. Condition 2 says for all agents and for every list it is possible to encounter in equilibrium, the strategy chosen conditioned on this encounter maximizes the expected payoffs. Condition 3 is the same as 2 except that it applies to out-of-equilibrium encounters. Condition 4 says an agent will always agree to play with a potential partner when the worst that can happen is that he realizes a weakly higher payoff than he would expect from passing instead. This is necessary to prevent trivial outcomes

where all agents pass in response to encountering all lists. Passing, of course, is a weak best response to passing, so this forms a degenerate equilibrium.

In the next two sections, we maintain the two assumptions that population size, I , is even and $I \geq 4$. The first is needed because this is a matching game, and if the population were odd, we would have to include the possibility of not finding a potential match in any given period. This would needlessly complicate our model. The second is needed because if $I = 2$ agents would always be matched with the same player each period which would make the endogeneity of the game degenerate.

3. The coexistence of conventions in a simple coordination game

Let Γ be a simple coordination game with $X = \{A, B\}$ and $F(A, A) = \alpha$, $F(B, B) = \beta$, and $F(A, B) = F(B, A) = 0$, as shown below.

Coordination Game		
	A	B
A	α, α	$0, 0$
B	$0, 0$	β, β

with $\alpha, \beta > 0$.

This game has the feature that the two agents receive gains from trade if they choose the same actions, and no gains from trade if they do not choose the same actions. The possible lists are $\{A\}$, $\{B\}$, and $\{A, B\}$. We will refer to an agent with list $\{A\}$ as an A-type individual, one with list $\{B\}$ as a B-type individual, and one with list $\{A, B\}$ as an AB-type individual. The issue we wish to address is the existence of a pregame perfect equilibrium in which some agents choose list $\{A\}$ and some choose list $\{B\}$, and A-types and B-types do not interact with each other. That is, when do agents choose not to interact with people who are different from them, and when can two conventions or social norms coexist?

We first address a narrower question: Can there be a pregame perfect equilibrium with both A-types and B-types but no AB-types, and in which A-types and B-types do not interact? We show below that the answer is “no” .

Proposition 1. *There is no pregame perfect equilibrium with $0 < I_A < I$ A-type agents who do not play against B-type agents, $I - I_A$ B-type agents who do not play against A-type agents, and no AB-types at all.*

Proof/

See appendix.

Next, we consider whether this nonexistence result is robust to the introduction of AB-types as intermediaries. When an A-type is matched with an AB-type, the only Nash equilibrium of the restricted base game is for both agents to play A , and so an A-type is indifferent between interacting with another A-type and interacting with an AB-type. B-types are also indifferent between interacting with other B-types and interacting with AB-types. Can these intermediaries allow the different conventions to coexist? As the next proposition shows, the answer is also negative.

Proposition 2. *There is no pregame perfect equilibrium with $I_A > 0$ A-type agents who do not play against B-type agents and $I_B > 0$ B-type agents who do not play against A-type agents.*

Proof/

See appendix.

The intuition behind Propositions 1 and 2 is the same. If an A-type’s expected payoff is at least as high as a B-type’s, the B-type has an incentive to defect and choose the list $\{A\}$. Switching enlarges the population of A-types, increasing the A-type’s payoff advantage. So, switching from list $\{B\}$ to list $\{A\}$ improves the agent’s payoff in two ways, by moving him to a type that earns higher expected payoff in the first place, and by increasing the expected payoff of that type by making them more likely to

match successfully. Intermediaries who have list $\{A, B\}$ play no role in this argument, and cannot help two conventions coexist.

Having established that some classes of equilibria do not exist, it remains to establish the classes that do. The next proposition characterized the entire set of PPE equilibria.

Proposition 3. *There exist exactly two classes of pregame perfect equilibria: equilibria with no A-types and equilibria with no B-types.*

Proof/

See appendix.

If the coordination game had n actions x_1, \dots, x_n with n distinct coordination payoffs ($F(x_j, x_j) \neq F(x_k, x_k)$ for $j \neq k$), then Proposition 3 could be extended. In such a case there would be n classes of pregame perfect equilibria, and in class k every agent would have strategy x_k in his list.

We delay applying these results to language and culture until section 6 in order to discuss the “discoordination game” first.

4. Discoordination Games

In this section, we turn our attention to the “discoordination game” which is in a sense the opposite of the coordination games above. Here, agents want to find a partner who will play the opposite strategy as he does. For example, someone who knows how to cook meat would prefer to partner with someone who knows how to cook side dishes for a picnic. We will argue that not only will it necessarily be the case that multiple conventions coexist, but also that there must exist “social market makers” who can interact with either convention.

Consider the symmetric *discoordination game* in which $X = \{A, B\}$ and $F(A, A) = F(B, B) = 0$, and $F(A, B) = F(B, A) = \alpha$, as shown below.

Discoordination Game		
	A	B
A	0, 0	α, α
B	α, α	0, 0

Note the following:

1. It is not a PPE for all agents to learn only one strategy. Clearly, any single agent who deviated would always find a partner willing and able to play the opposite strategy and would therefore receive the maximum payoff. This would certainly be a higher payoff than what is received by the remaining agents who would either have to wait to find this player or accept a zero payoff by playing with a coordinating agent.
2. It is not a PPE for all agents to learn both strategies. When two such players meet, they would not know which of them should play A and which should play B . The best they could do is flip a coin resulting in an expected payoff of $\frac{1}{2}\alpha$. If any agent deviated and learned only A then any two-strategy agent would know to play B in the resulting game. Thus, the deviant agent would always find a partner willing and able to play the opposite strategy and would also receive the maximum payoff.
3. A PPE does exist for this game when I is divisible by 3. In the equilibrium $I/3$ agents choose list $\{A\}$, $I/3$ choose list $\{B\}$, and $I/3$ choose the full list $\{A, B\}$. All agents use the same continuation game strategy $g(\ell'|\ell)$ given by

$$\begin{aligned}
g_A(\{A\}) &= N & g_A(\{B\}) &= A & g_A(\{AB\}) &= A. \\
g_B(\{A\}) &= A & g_B(\{B\}) &= N & g_B(\{AB\}) &= B. \\
g_{AB}(\{A\}) &= A & g_{AB}(\{B\}) &= B & g_{AB}(\{AB\}) &= N.
\end{aligned}$$

It is straightforward to verify that the continuation strategies satisfy the conditions for pregame perfect equilibrium.¹⁴ To see that the choices of lists are consistent

¹⁴ Because all possible lists occur in the equilibrium, the consistent belief requirement that governs beliefs when encountering an off-the-equilibrium-path list plays no part in the equilibrium.

with equilibrium, note that $EPO_A(S) = EPO_B(S) = EPO_{AB}(S) = \alpha - \frac{I-3}{2i}z$, and if any agent switches to a different list, his expected payoff changes to $\alpha = \frac{I}{2I-3}z$, which is worse for all $z > 0$.

In a coordination game, agents benefit from playing against others who behave the same as themselves, and pregame perfect equilibrium allows only one mode of behavior to exist in equilibrium. In a discoordination game, in contrast, two modes of behavior *must* coexist. Pregame perfect equilibrium requires the existence of intermediaries to serve as “social market makers,” that is, agents with full lists who can interact with both single-list types. Thus, pregame perfect equilibrium in discoordination games precludes full specialization by type.

To understand the discoordination game more fully, consider the following classical interpretation. (We provide two interpretations that relate more directly to culture below.) Two agents are matched and, if they agree to play, receive positive payoffs if they find each other and zero payoffs if they do not. The two available actions are to “go looking” or “remain in one place.” If they both remain in one place they fail to find each other, and if they both go looking they also fail to find each other. Successful outcomes occur only when one agent goes looking and the other remains in one place.¹⁵ Parents realize this and teach their children to remain in one place when they become separated. However, that rule only works for asymmetric populations where parents are always matched with children. What would happen in a population with only adults? In the endogenous game the equilibrium specifies a group that only goes looking (finders), a group that only remains where they are (stayers), and a group that can do either (intermediaries), all of equal size. Finders refuse to play against other finders, stayers refuse to play against other stayers, and intermediaries refuse to play against other intermediaries. If a stayer, say, switches to become a finder, then he joins a now-larger group and has a higher probability of matching with another finder, which is unsatisfactory. The same thing occurs if he switches to become an

¹⁵ Burdett et al. (1995) investigate a seeker-finder game in a monetary model. The key difference between their game and ours is that theirs has a fundamental asymmetry between agents, with sellers being agents who already own the good and buyers being those who do not.

intermediary.

The difference between the discoordination game and the coordination game, then, and what drives the existence of equilibria with distinct groups in the former but not in the latter, is the property that in a coordination game a player wants to match with someone like himself but in a discoordination game a player wants to match with someone who is different. This, then, is the key to the failure of the coexistence of conventions. Conventions cannot coexist when individuals with one convention benefit from interacting with others who share the same convention but not with agents who follow different conventions.

Analysis of the discoordination game also highlights a subtle difference between the endogenous and standard game theoretic approaches. In the classical approach the labels for the actions do not matter, and so the discoordination would be as easy to achieve as coordination would be in a standard coordination game. Under the endogenous approach, though, the equilibria of coordination game are very different in nature from the equilibria of the discoordination game, and the labels for the actions become an important part of the structure of the game. One way to think about this is that in the classical approach a player has an identity, either row or column, and can condition the play on that identity. Because of this, reversing the action labels for the column player has no impact on the equilibria of the game. In the endogenous approach, though, when the game is symmetric the only identity a player has is the list chosen in the pre-game, and reversing the action labels for the column player requires a way to distinguish column from row players in the matching phase of the endogenous game. Without this extra means of identification, reversing the action labels has an impact on the outcomes.

Discoordination games have interesting implications in the context of culture and behavior we are concerned with in this paper.

Alpha, Beta, and Gamma Males: Consider a situation in which men come together for a partnership to complete a project and share the profits. In each partnership, a leader and a follower are needed. Men develop skill at taking either role,

or sometimes, both roles. Examples include things like hunting expeditions, preparing a presentation for the boss, or trying to meet women at a fraternity party. Two alpha-males means two leaders and a failed venture. Two beta-males means no leadership, and a failed venture. Two gamma-males leads to a coordination problem and only half the partnerships consisting of a leader and a follower. Thus, we find that all three types of males will exist in equilibrium and agents will spend time finding the compatible partner.¹⁶

Caste systems: A society needs people to specialize in and perform all kinds of work. The analysis above offers a potential explanation for the apparent stability of caste-based societies. Suppose successful partnerships consist of a high-caste and a low-caste member (or even several low-caste members). The major difference here is that in this context search is not an entirely random process. Caste is usually a clearly identifiable characteristic, even from a distance. If a high-caste agent needs a low-caste partner, he can readily identify one and partner without search (and the same is true for a low-caste agent.) Thus, market makers do not arise in this context. More importantly, a low-caste agent would not have any incentive to learn high-caste skills. He would never be asked for a partnership by another low-caste agent. Thus, the endogeneity of the game serves to reinforce the caste-system. If an agent can “pass” for being a member of another type, however, the story is different. A low-caste person who wants to provide a high-caste skill must appear outwardly to be high-caste as well in order to have success in finding partners. This might explain the relative persistence of caste distinctions based on race, gender, family connections, and even such characteristics as height and beauty. On the other hand, caste distinctions based on behaviors or appearances that *any* agent can learn and adopt are more difficult to

¹⁶ Of course, in some (but not all) games like this, it might be possible for gamma-males to meet and then agree who will be the alpha and who will be the beta. In this case all agents would choose to be gamma-males. Essentially, equilibrium strategies are asymmetric with some otherwise-identical players choosing one strategy when they are paired and others choosing a different strategy. One way around this caveat is to require symmetry, as we do above. However, we can also recover the result by adding a cost of learning a strategy (or mode of behavior). In this case, the higher the cost the fewer gamma-males we see in equilibrium, but we would still see all types existing together (at least until the cost of learning a second strategy got prohibitive.)

maintain. Social mobility results.

5. An Application of Endogenous Coordination Games to Culture and Language

Our motivation to study endogenous games is that we believe that they are common in real life and therefore are an institution that should be studied. In this section, we consider the a game in which agents can choose to learn English, Spanish or both. Since it is more productive to play any game with a partner you can understand, this is an example of coordination game. Suppose that A-types from section 3 speak only English, B-types speak only Spanish, and AB-types are bilingual. Propositions 1 and 2 state there is no equilibrium with both English-only speakers and Spanish-only speakers, regardless of whether or not the population contains bilingual people. Proposition 3 establishes the existence of equilibria with a combination of English-only and bilingual agents and equilibria with a combination of Spanish-only and bilingual agents. Moreover, these are the only equilibria, and in equilibrium a single language is sufficient for all interactions. Thus, it appears that cultural assimilation is inevitable.

This result suggests that integration in schools or neighborhoods should be enough to generate the benefits of peer effects. However, casual inspection attests to the fact that we do see cultures coexisting within integrated societies. Clearly, the benchmark case considered in Section 3 must be missing something. The reasoning from that section suggests that when agents match randomly with the population at large, they benefit from sharing the same behavioral norm as the majority group. Several factors might make it more likely that minority cultures are able to continue to exist in light of the pressures towards assimilation.

- A group that wants to preserve its culture can improve its odds by increasing the likelihood that its members will interact with each other rather than the population at large. For example, members may go to the same church, engage in the same line of work or live in a ghetto.

- Policies may make it cheap or expensive to learn certain strategies. For example, teaching Welsh, Celtic or Catalan in schools, or forbidding Native American children to use anything but English in school will affect assimilation.
- Agents might have a strong preference for one of the coordinated outcomes over the other. For example, assimilation is more likely if the “affinity” benefits are outweighed by the “large network” benefits of joining the larger community; conversely, cultural segregation might arise when the affinity benefits are large. If the size of a group falls below a critical mass, however, separation is hard to sustain. Given the search cost, it will not benefit agents to hold out in hopes of finding another agent of their affinity type, and agents will instead learn the strategies used by the larger society.

We consider in more detail the case in which there are two types of players: a -natives and b -natives. The first step is to adapt the theory of endogenous games to allow for these “heritage-based” preferences.

To construct an endogenous version of the coordination game with heritage based preferences, suppose that the numbers of a -natives and b -natives in the population are H^a and H^b , respectively, and these numbers are common knowledge. Set $I = H = H^a + H^b$. In the pregame agents choose lists, as before. In the continuation game, agents are randomly matched, and observe their potential opponents’ lists *along with their heritage* (i.e. their opponents’ payoff functions).¹⁷ The game proceeds as before, with agents deciding whether to play or pay a delay cost of z and be rematched. When a pair elects to play, they play actions from their lists, receive their payoffs, and exit the game to be replaced by agents with the same heritage and the same strategies. The definitions of consistent beliefs and pregame perfect equilibrium can be extended to this setting in a straightforward manner by replacing lists in the earlier definitions by list-heritage pairs.

¹⁷ Up to this point in the paper, an agent’s “type” was determined by his chosen list, so “type” was endogenous. For consideration of heritage, a “type” becomes a list-heritage pair, which has an exogenous component.

First, consider the base game in which

$$F^a(A, A) = F^b(B, B) = \alpha + \eta$$

$$F^a(B, B) = F^b(A, A) = \alpha$$

$$F^h(A, B) = F^h(B, A) = 0$$

for $h = a, b$ where $\eta > 0$. The parameter η captures agents' preferences for their native actions, with increases in η corresponding to stronger preferences for the native action. Since we assume that both heritage and lists are observable, the strategies must be conditioned on both. Formally $g_\ell^h(h', \ell')$ is the strategy employed by an agent of affinity type h with list ℓ who encounters an agent with affinity type h' with list ℓ' . We are now able to address the question of whether there exists a pregame perfect equilibrium in which agents refuse to interact with agents of the other heritage.

Proposition 4. *If the delay cost z is small enough, there exists a pregame perfect equilibrium in which agents' lists do not contain the non-heritage action and agents play only against others from the same heritage.*

Proof/

See appendix.

Proposition 4 establishes that, if the delay cost z is small relative to the preference for the native action, η , there exists an equilibrium of cultural separation in which a -natives never learn action B and only interact with other a -natives, while b -natives never learn action A and only interact with other b -natives. When the delay cost is small, a -natives find it more beneficial to wait to meet another a -native in order to gain the benefit η than to learn action B and play the first opponent, and b -natives share the same incentives.

The actual condition for “small enough” z is worth a closer look. It reduces to

$$z < \frac{H^j - 1}{H - 1} \eta$$

for $j = a, b$. The restriction on z becomes looser, and therefore cultural separation is more likely to occur, when the population is more evenly split among the two types and the affinity for native behavior η is large. Conversely, when the population contains a large majority and a small minority, assimilation is more likely to occur.¹⁸

This analysis highlights the reason why cultural separation can persist in integrated schools, for example. Even though agents with different preferences are intermingled and are sometimes matched with members of the other heritage group, if the native preference is strong, the heritage group is relatively large, or the cost of waiting for a match with someone from the same heritage group is small, cultural separation can persist. In the presence of these conditions, integrating two cultural groups is not enough, because while integration can lead to face-to-face encounters, individuals still can refuse to turn these into mutually beneficial relationships.¹⁹

6. Conclusions

This paper uses the endogenous game framework introduced in Conley and Neilson (2008) to explore coordination and discoordination games. We argue that agents really do choose the partners with whom they play in many types of games and thus may choose to do such things as failing to learn unattractive actions in hopes of obtaining a better partner. Thus, endogenous games approximate many real world situations and it is therefore important to study their implications on equilibrium outcomes.

¹⁸ This result can be compared to the results derived from the empirical homophily literature, which is interested in groups integrating rather than coexisting. In other words, these studies look for situations in which z is sufficiently large that the premise of Proposition 4 fails to hold. In a recent study, Weinberg (2013) finds that integrated relationships between Hispanics and non-Hispanics are maximized when they are in a 50:50 split, which suggests that for such interactions the affinity parameter η is small. In contrast, relationships between blacks and non-blacks are maximized at a much more skewed split, which would occur if η is comparatively large.

¹⁹ Of course, the converse is also true. That is, if delay cost is large, native preference weak, or the heritage group relative small, then cultural segregation becomes more difficult to sustain. At some point, all PPE involve cultural assimilation as in section 3.

For endogenous discoordination games we find that the equilibria involve some agents learning only the first strategy, some learning only the second strategy and some agents learning both. The presence of these social market makers who learn both strategies is rather surprising and is not seen in the one-shot discoordination game. The equilibria of the endogenous and one-shot coordination games are also different, but there is no role for market makers and they will not appear in endogenous game equilibria in general.

Since one of our main motivations is to explore the positive implications of endogenous games, this equilibrium result for endogenous coordination games is disturbing. It implies that two non-segregated groups with different modes of behavior or different cultural conventions cannot coexist. Provided everyone has the ability to make the same choices, (that is, if there are no inherent differences between agents before the game starts) *de facto* cultural segregation is inconsistent with equilibrium and so one mode of behavior or cultural convention that everyone in the society shares should emerge.

Since many cultures and languages do in fact manage to coexist in the real world, we explore what modification to the underlying coordination game are needed to be consistent with this. We propose adding an asymmetry called “heritage” . Formally, “ heritage” means that agents have a preference for using one or the other strategies (but still get more benefits from coordinating strategies with the agents with whom they end up playing the game). With this modification, we find that two cultures can coexist if the delay cost is small, the size of a heritage group is relatively large, or the heritage preference is relatively large.

To relate these results to the real world, consider the case of Welsh speakers in Wales. The theory predicts that either everyone speaks English and some also speak Welsh, or everyone speaks Welsh and some also speak English. Nearly everyone in Wales can communicate in English. Furthermore, according to British census data, 54% of the population spoke Welsh in 1954, but by 1981 that number had fallen to 19%. This raises the issue of how a group can go about preserving its culture, given that

equilibrium requires that everyone adopt the same dominant culture. Wales addressed the disappearance of its language with the Education Act of 1988, which required the teaching of Welsh to schoolchildren up to age 16. In part because of this effort, the fraction of bilinguals in the population rose to 21% in the 2001 census.

The Welsh institutional response essentially encourages the expansion of the bi-cultural group, but makes no attempt to create or preserve a mono-cultural, Welsh-only group. In fact, the theory implies that there is no way for “cultural intermediaries” who are able and willing to act with two different mono-cultural groups to preserve the multiculturalism. The reason is that members of the two mono-cultural groups are sometimes put in a position in which they can benefit from interacting with each other but prevented from doing so by cultural differences. If passing up interactions is costly, equilibrium prescribes one dominant culture, and cultural struggles are winner-take-all. The only way, then, to preserve mono-cultural societies is through institutional means that prevent opportunities for members of different cultures to meet in the first place.

Our model shows that self-segregation in a geographically-intermingled population is not enough. Rather, physical separations that keep members of different cultures from meeting are the only way to prevent such cultural struggles. Thus, separatist movements, such as those among the Basque population of Spain, are consistent with the model, while efforts to force cultural unification, such as English-only legislation in some U.S. states, are not. The model predicts that one culture will become universal throughout the population anyway, so attempts to force this are unnecessary, while efforts to keep different cultural groups separate are needed to preserve mono-cultural identities.

Appendix A - Justifying a new equilibrium concept

While Pregame Perfect Equilibrium is formally different from Sequential Equilib-

rium, we see PPE as simply the natural extension of subgame perfection to endogenous games. To see why we need to make this modification, consider the following simple example. Suppose there are four players, each of whom chooses a list from $\{a, b\}$. We will pay particular attention to player 1 who is going to be matched with player 2. Each individual has three possible lists: $\{a\}$, $\{b\}$, and $\{a, b\}$. For any matched pair, then, there are 9 possible pairs of lists:

$$\begin{aligned}
 &\{a\}, \{a\}, \\
 &\{a\}, \{b\}, \\
 &\{a\}, \{a, b\}, \\
 &\{b\}, \{a\}, \\
 &\{b\}, \{b\}, \\
 &\{b\}, \{a, b\}, \\
 &\{a, b\}, \{a\}, \\
 &\{a, b\}, \{b\}, \\
 &\{a, b\}, \{a, b\}.
 \end{aligned}$$

Thus, for each matched pair there are 9 possible nodes corresponding to the 9 possible pairs of lists.

Now suppose that player 1 is matched with player 2 and observes player 2's list. Player 1 knows her own list, and she also observes player 2's list. So, player 1 knows which node pertains from the set of 9 generated by the (player 1, player 2) pairing. However, player 1 does not know which of the 9 nodes from the (player 3, player 4) pairing is the actual decision set. Consequently, player 1's information set is not a singleton, and contains all of the possible nodes from the (player 3, player 4) pairing.

The definition of subgame perfection requires that play be a Nash equilibrium in every proper subgame, and a proper subgame is one that extends from a singleton information set. As described above, player 1's information set is not a singleton, and therefore there are no proper subgames extending from the first match. Consequently, subgame perfection has no bite. In essence, player 1 is missing the information about the lists of the players with whom he could be matched if he does not play against 2.

If player 1 and 2 agree to play, the information about the lists of players 3 and 4 becomes irrelevant. The information set is not a singleton, but every node in the information set yields the same set of branches with the same payoffs. We could just choose one of the nodes and work out the Nash equilibrium from there. But, if players 1 and 2 do not agree to play, the lists of players 3 and 4 become relevant. This is why subgame perfection will not work.

Alternative solution concepts, notably sequential equilibrium, govern behavior from non-singleton information sets. However, sequential equilibrium is constructed so that player 1 can infer things about player 2's type after observing player 2's action. This does not help in our game, because player 2's action, his choice of list, completely reveals his "type." We need an equilibrium concept that allows player 1, after observing player 2's action, to infer things about players 3 and 4. In essence, we want player 1 to observe player 2's list and then infer something about which node pertains for the (player 3, player 4) pairing so that player 1 can form expectations about what will happen if 1 and 2 decide not to play against each other.

Because neither subgame perfection nor sequential equilibrium tie down beliefs about the nodes reached by other pairings, we identify a new equilibrium concept, pregame perfect equilibrium. Like other refinements of perfect Bayesian equilibrium, it operates by tying down beliefs following off-equilibrium-path moves. In essence, a combination of lists and subsequent actions is a pregame perfect equilibrium if every player believes that all other players choose lists according to the equilibrium profile, so that no player ever expects to see an out-of-equilibrium list.

To see how this works with the above example, suppose that the candidate equilibrium has all four players choosing list $\{a\}$, choosing to play against anyone with list $\{a\}$ and choosing not to play otherwise, and playing a when they play. When players 1 and 2 are matched, player 1 believes that players 3 and 4 have both chosen the list $\{a\}$, and so player 1 places zero probability on the other 8 nodes. This, in essence, makes player 1's information set a singleton and the structure of subgame perfection suffices for solving the game.

Appendix B - Formal definition of pregame perfection

In this appendix, we give formal definitions to some of the concepts used in the body of the paper. We begin by defining the various notions of expected payoffs.

The structure of the game is for agents to match randomly and decide either to play or continue searching. Given a pair of strategies, the following binary function $PL : \mathcal{S} \times \mathcal{S} \rightarrow \{0, 1\}$ indicates the outcome of any given random pairing:

$$PL(s_i, s_j) \equiv \begin{cases} 0 & \text{if } g_j(\ell_i) = N \text{ or } g_i(\ell_j) = N \\ 1 & \text{if } g_j(\ell_i) \neq N \text{ and } g_i(\ell_j) \neq N \end{cases}$$

Given a particular agent i 's choice of strategy s_i and a strategy profile for the remaining agents s_{-i} , we can divide the set of potential partners two sets: those who will accept a match with i and with whom i will also agree to play, and matches that will result in at least one of the two agents choosing not to play. Respectively, these may be defined formally as follows:

$$\mathcal{I}^P(s_i, s_{-i}) \equiv \{j \in \mathcal{I} \setminus i \mid PL(s_i, s_j) = 1\},$$

$$\mathcal{I}^N(s_i, s_{-i}) \equiv \{j \in \mathcal{I} \setminus i \mid PL(s_i, s_j) = 0\}.$$

Thus, given any strategy choice s_i for agent i and strategy profile s_{-i} for the remaining agents, the probability that an agent does not match in any given period can be calculated as follows:

$$P^N(s_i, s_{-i}) \equiv \frac{|\mathcal{I}^N(s_i, s_{-i})|}{(I - 1)}$$

where $|\cdot|$ denotes the cardinality of a set. Given this, the expected payoff that agent i receives under strategy profile s is

$$EPO(s_i, s_{-i}) = \sum_{t=0}^{\infty} (P^N(s_i, s_{-i}))^t \left(\frac{1}{I-1} \sum_{j \in \mathcal{I}^P(s_i, s_j)} F(g_i(\ell_j), g_j(\ell_i)) - P_i^N(s_i, s_{-i})z \right)$$

In order to test that agents are following optimal strategies in every subgame, we will also need to know the expected continuation payoff from every subgame. That is, conditional on having been matched with an agent who has chosen list $\hat{\ell}$, we need to know the expected payoff from this point forward. To calculate this we need to restate the objects above to take account of the current matching realization. Thus:

$$\mathcal{I}^P(s_i, s_{-i} \mid \hat{\ell}) \equiv \{j \in \mathcal{I} \setminus i \mid PL(s_i, s_j) = 1 \text{ and } \ell_j = \hat{\ell}\},$$

$$\mathcal{I}^N(s_i, s_{-i} \mid \hat{\ell}) \equiv \{j \in \mathcal{I} \setminus i \mid PL(s_i, s_j) = 0 \text{ and } \ell_j = \hat{\ell}\}.$$

$$P^N(s_i, s_{-i} \mid \hat{\ell}) \equiv \frac{|\mathcal{I}^N(s_i, s_{-i} \mid \hat{\ell})|}{|\mathcal{I}^N(s_i, s_{-i} \mid \hat{\ell}) \cup \mathcal{I}^P(s_i, s_{-i} \mid \hat{\ell})|}$$

The *expected continuation payoff* can then be expressed as follows:

$$EPO(s_i, s_{-i} \mid \hat{\ell}) = \frac{1}{|\{j \in \mathcal{I} \setminus i \mid \ell_j = \hat{\ell}\}|} \sum_{j \in \mathcal{I}^P(s_i, s_j \mid \hat{\ell})} F(g_i(\ell_j)g_j(\ell_i)) + P^N(s_i, s_{-i} \mid \hat{\ell}) (EPO(s_i, s_{-i}) - z)$$

Next, we define the consistent belief system for a particular agent as derived from a commonly held consistent belief system $p(g|S, \ell)$. Informally, the beliefs held by agent i , $p_i(g|S, \ell)$, regarding the probability distribution of strategies he himself will encounter is a modification of p that removes agent i from the population probabilities:

1. For all $g \in \mathcal{G}$ and $\ell \neq \ell_i$, $p_i(g|S, \ell) = p(g|S, \ell)$.
2. For all $g \in \mathcal{G}$ and $\ell \in L_i$, $p_i(g|S, \ell) = \frac{|\{j \in \mathcal{I} \setminus i \text{ s.t. } g_j = g \text{ and } \ell_j = \ell\}|}{|\{j \in \mathcal{I} \setminus i \text{ s.t. } \ell_j = \ell\}|}$.
3. For all $g \in \mathcal{G}$ and $\ell \in \bar{L}_i$, $p_i(g|S, \ell) = p(g|S, \ell)$.

We will place a bar over the next several objects we define to remind us that these refer to out of equilibrium play. Note that the following are only well defined for these out of equilibrium encounters. We begin by defining the analog of \mathcal{I}^P and \mathcal{I}^N for disequilibrium situations. In this case, we need to know the set of stage game strategies used by disequilibrium players that will or will not result in a match taking place.

$$\bar{\mathcal{G}}^P(s_i, s_{-i} \mid \hat{\ell}) \equiv \{g \in \mathcal{G} \mid PL(s_i, g, \hat{\ell}) = 1\},$$

$$\bar{\mathcal{G}}^N(s_i, s_{-i} | \hat{\ell}) \equiv \{g \in \mathcal{G} \mid PL(s_i, g, \hat{\ell}) = 0\}.$$

Given this and a set of consistent beliefs p , the probability of no match taking place in a given round is easily calculated:

$$\bar{P}^N(s_i, s_{-i} | \hat{\ell}) \equiv \sum_{g \in \bar{\mathcal{G}}^N} p_i(g | S, \ell)$$

Using this, we can finally state expected payoff in the event of an out of equilibrium encounter as follows:

$$\begin{aligned} \overline{EPO}(s_i, s_{-i} | \hat{\ell}) = \\ \sum_{g \in \bar{\mathcal{G}}^P(s_i, s_j | \hat{\ell})} p_i(g | S, \hat{\ell}) F(g_i(\hat{\ell}), g(\ell_i)) + P^N(s_i, s_{-i} | \hat{\ell}) (EPO(s_i, s_{-i}) - z). \end{aligned}$$

Appendix C - Proofs

Proposition 1. *There is no pregame perfect equilibrium with $0 < I_A < I$ A-type agents who do not play against B-type agents, $I - I_A$ B-type agents who do not play against A-type agents, and no AB-types at all.*

Proof/

Let I_A , I_B , and I_{AB} denote the numbers of A-types, B-types, and AB-types, respectively. We want to establish the existence or nonexistence of a pregame perfect equilibrium in which $I_A > 0$, $I_B > 0$, and $I_{AB} = 0$. Note that in any PPE, agents will always agree to play with their own types since this give the highest payoff possible. Thus, the expected payoff to A-types in this case is

$$EPO_A(S) = \frac{I_A - 1}{I - 1} \alpha + \frac{I_B}{I - 1} (-z + EPO_A(S)) = \alpha - \frac{I_B}{I_A - 1} z$$

and similarly the expected payoff to a B-type is

$$EPO_B(S) = \beta - \frac{I_A}{I_B - 1} z.$$

The key point to note is that $EPO_i(S)$ increases in I_i and decreases in I_j for $i, j \in \{A, B\}$ with $i \neq j$. Thus, if the payoff from being a B-type, for example, is higher than the payoff from being an A-type (either because β is higher or $I_B > I_A$), then A-types are better off becoming B-types. In addition, by making the switch, they further increase the payoff advantage of being an B-type. Even if the payoffs to each type are initially equal, an A-type increases his payoff by becoming a B-type. Consequently, for any possible combination of $I_A > 0$ and $I_B > 0$ with $I_A + I_B = I$ some type has an incentive to choose the other type's list, and so there is no PPE that satisfies the hypothesis.

■

Proposition 2. *There is no pregame perfect equilibrium with $I_A > 0$ A-type agents who do not play against B-type agents and $I_B > 0$ B-type agents who do not play against A-type agents.*

Proof/

Suppose there is such an equilibrium. The expected payoff to an A-type is

$$EPO_A(S) = \frac{I_A + I_{AB} - 1}{I - 1} \alpha + \frac{I_B}{I - 1} (-z + EPO_A(S)) = \alpha - \frac{I_B}{I_A + I_{AB} - 1} z$$

and similarly the expected payoff to a B-type is

$$EPO_B(S) = \beta - \frac{I_A}{I_B + I_{AB} - 1} z.$$

We can compute

$$EPO_A(S) - EPO_B(S) = \alpha - \beta + \frac{I_A - I_B}{(I_A + 1)(I_B + 1)} (I - 1) z,$$

which again is increasing in I_A and decreasing in I_B . The argument is therefore the same as for Proposition 1. In all cases either an A-type, a B-type, or both have an incentive to defect, contradicting the hypothesis that an equilibrium with both types exists.

■

Proposition 3. *There exist exactly two classes of pregame perfect equilibria: equilibria with no A-types and equilibria with no B-types.*

Proof/

First consider an economy in which I_A agents choose list $\{A\}$, and the other $I - I_A$ agents choose list $\{A, B\}$. Two cases arise based on the value of the delay cost z .

Case 1: $z < \alpha$. Suppose that A-types choose the following strategy in the continuation game:

$$g_A(\{A\}) = A \quad g_A(\{B\}) = N \quad g_A(\{AB\}) = A$$

and AB-type agents choose the strategy given by:

$$g_{AB}(\{A\}) = A \quad g_{AB}(\{B\}) = B \quad g_{AB}(\{AB\}) = A.$$

In addition, suppose that both types of agents think that if they happen to encounter a B-type agent, his continuation strategy is:

$$\bar{g}_B(\{A\}) = N \quad \bar{g}_B(\{B\}) = B \quad \bar{g}_B(\{AB\}) = B.$$

We will leave aside for the moment the question of whether this generates a consistent belief system.

Observe that given these strategies, all agents match in the first round and get the maximum payoff. Thus, no strategy could yield a higher payoff from the standpoint of the pregame and so condition (1) of PPE is satisfied. This also means that in any matching that occurs after the pregame, the strategies for each type of player still yield the maximum payoff. Thus, in equilibrium encounters, the conditional expected payoff is maximized and condition (2) of the PPE is satisfied. Now suppose that an A or AB type happens to be matched with a B type. This is an out of equilibrium encounter. For AB-types, this presents no problem as both agents play “B” and get maximum

payoffs. For A-types, the expected payoff to this strategy is to suffer a cost of $z < \alpha$ this period and then get the maximal payoff of 1 with certainty next period. Since $z < \alpha$, condition (3) of PPE is satisfied. Finally, from the above arguments, it is clear that when $z < \alpha$, it is optimal for A-types to pass in this case, and so condition (4) is satisfied.

To complete the proof for this case, we have to show that this belief system is consistent. Parts (1) and (2) are trivially satisfied, so only part (3) is in question. In particular, we have to check that the strategy ascribed to the (non-existent) B-types is, in fact, an optimal strategy. Clearly, B-types should play B against other B-types and against AB-types. The question is playing N against A-types credible? Since A-types play N when meeting B-types, however, it does not matter what B-types do. Thus, no other strategy would yield a higher payoff, and so part (3) of the definition of a consistent belief system is satisfied as well.

Case 2: $z \geq \alpha$. Under this condition an A-type's best response when facing a B-type is to play, since the penalty for delay is too steep. Accordingly, suppose that A-types choose the following strategy in the continuation game:

$$g_A(\{A\}) = A \quad g_A(\{B\}) = A \quad g_A(\{A, B\}) = A$$

and AB-type agents choose the strategy given by

$$g_{AB}(\{A\}) = A \quad g_{AB}(\{B\}) = B \quad g_{AB}(\{A, B\}) = A.$$

In addition, suppose that both types of agents think that if they happen to encounter a B-type agent, his continuation strategy is

$$\bar{g}_B(\{A\}) = B \quad \bar{g}_B(\{B\}) = B \quad \bar{g}_B(\{A, B\}) = B.$$

As with Case 1, this generates a consistent belief system, and the continuation game strategies are consistent with PPE.

It remains to show for this case that no agent wants to change his list in the pregame. An A-type has no incentive to change to a B-type or an AB-type because he

receives the maximal payoff in the first period. For the same reason, an AB-type has no incentive to change to being an A-type or a B-type. So, even though the high delay cost makes it optimal for an A-type to play against a B-type, there is no incentive to become a B-type.

We conclude that a PPE exists with no B-types. A similar argument establishes the existence of a PPE with no A-types for $z \geq \beta$ and $z < \beta$. Proposition 2 then implies that the *only* pregame PPE have either no B-types or no A-types when AB-types can solve the coordination problem.

■

Proposition 4. *If the delay cost z is small enough, there exists a pregame perfect equilibrium in which agents' lists do not contain the non-heritage action and agents play only against others from the same heritage.*

Proof/

We suppose that a -affinity agents learn list $\{A\}$, b -affinity agents learn list $\{B\}$, and these agents play the following strategies for $h = a, b$:

$$\begin{aligned} g^a(h, \{A\}) &= A & g^a(h, \{B\}) &= N & g^a(h, \{AB\}) &= A \\ g^b(h, \{A\}) &= N & g^b(h, \{B\}) &= B & g^b(h, \{AB\}) &= B. \end{aligned}$$

We must also state a belief system for heritage/list combinations that do not appear in equilibrium. Here the “bar” reminds us that these are not equilibrium agents and the notation $\bar{g}_{\ell'}^{h'}(h, \ell)$ denotes the continuation game strategy of an h' -affinity player with list ℓ' matched against an h -affinity player with list ℓ . Thus, for $h = a, b$:

$$\begin{aligned} \bar{g}_B^a(h, \{A\}) &= N & \bar{g}_B^a(h, \{B\}) &= B & \bar{g}_B^a(h, \{AB\}) &= B \\ \bar{g}_{AB}^a(h, \{A\}) &= A & \bar{g}_{AB}^a(h, \{B\}) &= N & \bar{g}_{AB}^a(a, \{AB\}) &= A & \bar{g}_{AB}^a(b, \{AB\}) &= N \\ \bar{g}_A^b(h, \{A\}) &= A & \bar{g}_A^b(h, \{B\}) &= N & \bar{g}_A^b(h, \{AB\}) &= A \end{aligned}$$

$$\bar{g}_{AB}^b(h, \{A\}) = N \quad \bar{g}_{AB}^b(h, \{B\}) = B \quad \bar{g}_{AB}^b(a, \{AB\}) = N. \quad \bar{g}_{AB}^b(b, \{AB\}) = B.$$

We begin by showing the definition of PPE is satisfied for a -natives. From the above, we can calculate an a -native's expected equilibrium payoff as:

$$\begin{aligned} EPO^a(S) &= \\ \frac{H^a - 1}{H - 1}(\alpha + \eta) + \frac{H^b}{H - 1}(-z + EPO^a(S)) &= \\ (\alpha + \eta) - \frac{H^b}{H^a - 1}z & \end{aligned}$$

There are a variety of alternative strategies that an a -native might employ. In principle, showing that the strategy we suggest produces the highest expected payoff starting from the pregame requires that we test them all. Fortunately, most of these are obviously suboptimal, so we will begin by focusing on the one that is non-trivial and potentially optimal. Suppose that an a -native considers an alternative strategy s' in which learns B as well as A , and plays B when encountering a b -native. In this case, his expected payoff given the new strategy profile S' is:

$$\begin{aligned} EPO^a(S') &= \\ \frac{H^a - 1}{H - 1}(\alpha + \eta) + \frac{H^b}{H - 1}\alpha &= \\ \frac{H^a - 1}{H - 1}(\alpha + \eta) + \frac{H - H^a}{H - 1}\alpha &= \\ \frac{\alpha H^a + \eta H^a - \alpha - \eta + \alpha H - \alpha H^a}{H - 1} &= \\ \frac{(H - 1)\alpha + (H^a - 1)\eta}{H - 1}. & \end{aligned}$$

For this to be a best response for the deviating a -native, it must be the case that:

$$\frac{(H - 1)\alpha + (H^a - 1)\eta}{H - 1} \geq (\alpha + \eta) - \frac{H^b}{H^a - 1}z$$

The other alternative strategies are:

1. Learn only B and play this against all partners: clearly this is dominated by learning A and B and playing A if the agent happens to be partnered with an a -native.
2. Learn A or A and B but play only against b -natives: obviously, this is a dominated strategy.
3. Learn anything but always pass: obviously dominated.
4. Learn A and B but play only against a -natives: this has the same payoff as s and results in the same play.

We conclude that s maximizes the expected payoff a -natives when the delay cost is small and so condition (1) of the definition of PPE is satisfied.

Now we need to show that these strategies continue to be optimal in any equilibrium subgame. Suppose an a -native meets an a -native. Clearly playing A maximizes the expected payoff. Suppose instead he meets a b -native. Since b -natives do not play in this case, no stage game strategy can give anything more than $-z$ plus the continuation payoff. It follows that not playing is expected payoff maximizing. Thus, condition (2) of the definition of PPE is satisfied.

Next, consider out of equilibrium encounters.

1. If an a -native meets an a -native who only can play B , not playing is optimal for small z .
2. If an a -native meets an a -native who can play both A and B , playing A is optimal.
3. If an a -native meets a b -native who only can play A , playing A is optimal.
4. If an a -native meets a b -native who can play both A and B , not playing is optimal for small z .

Thus, condition (3) of the definition of PPE is satisfied. Finally, the only time a -natives choose to not play is when they encounter a b -native who can only play B . We show above this is optimal for small z . Thus, condition (4) of the definition of PPE is satisfied.

Obviously, since the strategies are symmetrically inverted for b -natives, the arguments above also show that the strategy S with this belief system also satisfies the

definition of PPE for b -natives as well.

It remains only to show that the strategies assigned to players not seen in equilibrium form a consistent belief system (that is, condition (3) of the definition of consistent beliefs is satisfied). Observe that the strategies, \bar{g} , given for these players fall into two categories: (1) For agents who have not learned their native strategies (a -natives who only know B , for example) \bar{g} requires them to play the non-native strategy when matched with someone who also has that strategy in their list, but not against someone without that strategy in their list. They play in order to earn a payoff of α , and for sufficiently small z waiting for the payoff of α is optimal. (2) For agents who know both strategies, \bar{g} requires them to play their native strategies against any equilibrium player who will also play this strategy and to not play otherwise. Again, for small z , this is expected payoff maximizing. We conclude that this PPE is supported by a consistent belief system.

■

References

- Bhaskar, V. and Vega-Redondo, R.** (2004): "Migration and the Evolution of Conventions," *Journal of Economic Behavior and Organization*, **55** pp. 397-418.
- Burdett, K., Coles, M., Kiyotaki, N., Wright, R.** (1995): "Buyers and sellers: Should I stay or should I go?," *American Economic Association Papers and Proceedings*, **85**, pp. 281-286.
- Calabuig, V., Olcina, G.** (2009): "Cooperation and cultural transmission in a coordination game," *Journal of Economic Behavior and Organization*, **72**, pp. 188-201.
- Conley, J.P., Neilson, W.S.** (2008): "Endogenous games and equilibrium adoption of social norms and ethical constraints," *Games and Economic Behavior*, **66**, pp. 761-774.
- Ellison, G.** (1993): "Learning, Local Interaction, and Coordination," *Econometrica*, **61** pp.1047-1071.
- Ely, J. C.** (2002): "Local Conventions," *Advances in Theoretical Economics*, Vol. 2, No. 1, Article 1.

- Ghosh, P., Ray, D.** (1996): "Cooperation in community interaction without information flows," *Review of Economic Studies*, **63**, pp. 491-519.
- Jackson, M.O., Wilkie, S.** (2005): "Endogenous games and mechanisms: Side payments among players," *Review of Economic Studies*, **72**, pp. 543-566.
- Kandori, M, Mailath, G., and Rob, R.** (1993): "Learning, Mutation, and Long Run Equilibria in Games," *Econometrica*, **61** pp. 29-56.
- Ku, H., Zussman, A.** (2010): "Lingua Franca: The role of English in international trade," *Journal of Economic Behavior and Organization*, **75**, pp. 250-260.
- Mengel, F.** (2008): "Matching structure and the cultural transmission of social norms," *Journal of Economic Behavior and Organization*, **67**, pp. 608-623.
- Morris, S.** (2000): "Contagion," *Review of Economic Studies*, **67** pp. 57-78.
- Oechssler, J.** (1997): "Decentralization and the Coordination Problem," *Journal of Economic Behavior and Organization*, **32** pp. 119-135.
- Rob, R., Yang, H.** (2010): "Long-term relationships as safeguards," *Economic Theory*, **72**, pp. 43, 143-166.
- Rubinstein, A. and Wolinsky, A.** (1985): "Equilibrium in a Market with Sequential Bargaining," *Econometrica*, **53** pp. 1133-1150.
- Spier, K.E.** (1992): "Incomplete contracts and signaling," *RAND Journal of Economics*, **72**, pp. 23, 432-443.
- Stringham, E.P.** (2011): "Embracing morals in economics: The role of internal moral constraints in a market economy," *Journal of Economic Behavior and Organization*, **78**, pp. 98-109.
- Sugden, R.** (2004): *The Economics of Rights, Co-operation and Welfare*. Palgrave Macmillan.
- Watson, J.** (1999): "Starting small and renegotiation," *Journal of Economic Theory*, **85**, pp. 52-90.
- Weinberg, B.A.** (2013): "Group design with endogenous associations," *Regional Science and Urban Economics*, **43**, pp. 411-421.
- Young, H. P.** (1993): "The Evolution of Conventions," *Econometrica*, **81** pp. 57-84.