ESSAYS IN ECONOMIC THEORY

Charles Roddie

**A THESIS PRESENTED TO THE FACULTY
OF PRINCETON UNIVERSITY
IN CANDIDACY FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY**

**RECOMMENDED FOR ACCEPTANCE
BY THE DEPARTMENT OF ECONOMICS**

Adviser: Wolfgang Pesendorfer

November 2008

Acknowledgements

I am grateful to Wolfgang Pesendorfer for advice and support and to Stephen Morris, Dilip Abreu, Satoru Takahashi, Thomas Romer, Drew Fudenberg, and George Mailath for valuable discussions.

Dedication

This thesis is dedicated to my family, my friends and my church, with gratitude for their constant support.

Contents

Abstract **iii**

1 Repeated Signalling and Reputation **1**

1.1 Introduction . . . 1
1.2 Model . . . 5
1.3 Equilibrium selection . . . 10
1.4 The Iterated Riley solution . . . 11
1.5 The Iterated Riley equilibrium and the D_ω condition . . . 14
1.6 Limit properties of the Iterated Riley equilibrium . . . 17
1.7 Reputation . . . 18
1.8 The general case . . . 24
1.9 Further work and extensions . . . 25
1.10 Appendix . . . 27

2 A Repeated Signalling Model of Reciprocity **39**

2.1 Introduction . . . 39
2.2 Framework . . . 40
2.3 Riley equilibria . . . 41
2.4 Refinement and equilibrium selection . . . 45
2.5 Discussion . . . 47
2.6 Appendix . . . 47

3 Mobility and redistribution with non-linear taxes **49**

3.1 Synopsis 49
3.2 Model 50
3.3 The movement phase 51
3.4 Definitions 52
3.5 A weaker result 53
3.6 A stronger result 55
3.7 Discussion 56

Bibliography **57**

Abstract

In the first chapter, "Repeated signalling and reputation", a signaller repeatedly signals his type to an uninformed player. The paper adapts the divinity refinements of the static signalling game to the repeated signalling game, selecting a dynamic version of the Riley equilibrium, defined iteratively, in which types separate minimally in each period. The model provides an alternative framework for studying reputation, generating under appropriate limits a modified Stackelberg property: each type above the lowest takes the action that maximizes Stackelberg payoffs, subject to separating from the lowest type. In contrast to the usual approach to reputation there are no behavioural types. It can be solved under arbitrary discount factors of both players: if the signaller discounts, the result above holds with the signaller's Stackelberg payoffs replaced by simply defined "discounted Stackelberg" payoffs. If the respondent has preferences not only over the actions but also over the type of the signaller, a differential equation characterizes the limit, combining reputational and pure type-signalling motives. Applications include work incentives, reputations for product quality, and limit pricing.

The second chapter, "A two-way repeated signalling model of reciprocation", studies a model with two signallers who signal their good will to each other by acting generously. The incentive to signal good will is that preferences of each type are dependent not only on one's own good-will but also on the perceived good-will of the other. The model is found to be solvable as a finite game and as an infinite game under a Markov assumption and the solution is characterized.

The third chapter, "Mobility and redistribution under general taxes", studies taxation and redistribution in a multi-regional setting with free mobility. The situation is described by a two-stage game with policy setting by regional governments and then migration and work choice by the population. When taxes are allowed to be non-linear and can discriminate between residents and immigrants, competition for taxable workers implies severe restrictions on redistribution in a stable equilibrium. The restrictions are independent of how policies in each region are generated, as long as policies are locally Pareto efficient. The formulation allows the study of tax policy under both moral hazard and mobility. There is a clean division between these two issues, both of which reduce the ability of regions to redistribute wealth.

Chapter 1

Repeated Signalling and Reputation

1.1 Introduction

A signalling model of reputation

The economic idea of reputation is that a patient player by taking a certain action may cause others to expect him to do the same thing in the future, even if it will be against his immediate interests. By doing this he has effectively the ability to commit to any action, receiving (close to) Stackelberg payoffs[1]. In the standard model of reputation ([12], [13], [19]) this conclusion requires the possibility of the patient player being a *behavioural* or *commitment* type. Such a type uses a fixed exogenous strategy, independent of expectations about the other player's strategy; in the simplest case this strategy is a particular action that is always taken. The normal type(s) of the long-run player can then develop a reputation for one of these actions over time by repeatedly mimicing one of these behavioural types, if his discount factor is high enough.

The signalling model of reputation proposed here has a different logic. Here there are only normal types of the signaller, giving a range of preferences. Instead of pooling with a behavioural type, each normal type separates from "worse" types. Types are correlated over time and this separation occurs at every stage. The signaller wants to be seen as a *higher* type and this give an incentive to take *higher* actions than are myopically optimal: by taking a *higher* action today, the signaller will be seen as a *higher* type today, and so will be expected expected to be a *higher* type tomorrow, and so expected to take a *higher* action tomorrow, leading to more favourable treatment by the other player. This is the reputational incentive to take higher actions than would be myopically optimal.

Suppose that types are unlikely to change from one period to the next (an assumption) and that each type's action does not change much from a given period to the next (a limit property as the length of the game tends to infinity). What generates a Stackelberg

[1] The Stackelberg payoff is the payoff in the stage game to a player who can commit to any action, while the other player best-responds. A Stackelberg action maximizes this payoff.

property is that by taking an action, the signaller signals that he is the type that preferred to take this action, and will be expected to be the same type in the following stage, and so to do (approximately) the same action in the next period. This holds only when the signaller chooses an action that is taken with positive probability by some type; the set of these such actions in the limit determines how the Stackelberg result is qualified.

Assume that the signaller is patient and that the receiver's preferences are over actions only and do not involve the signaller's type - as in standard reputation models. Taking the limit as the number of periods from the end of the game tends to infinity, as types become dense in some interval, and when type change becomes infinitely unlikely, we get the following reputation result: The lowest type takes his myopically optimal action. All other types take the actions that give them the greatest Stackelberg payoff, subject to separating (by the lowest type's Stackelberg payoff) from the lowest type[2]. So low types take the minimal action that separates them from the lowest type, while high types take their Stackelberg actions. The limit result is a combination of a separation property, such as is often seen in signalling models, with a Stackelberg property, often seen in reputation models.

The signalling model is more tractable than standard reputation models tend to be, generating a unique, simple and calculable solution under the refinement. In particular, the model remains solvable under general discount factors of both players. Standard reputational models require the reputation-builder to have discount factor 1 or tending to 1, and often myopic play by the other player, or at least a level of patience that becomes infinitely less than the reputation-builder's. In the repeated signalling model, the respondent's discount factor has no effect on the solution, and when the signaller has discount factor other than 1, the result given above only requires Stackelberg payoffs to be replaced with simply defined "discounted Stackelberg" payoffs.

The type of the signaller and actions of both players lie in intervals of real numbers, with monotonicity properties that will be spelled out presently. This is a loss of generality from the standard reputation model, which can work with very general stage games. But it is a natural specification for a large class of applied models. Examples include developing reputations for product quality, with the signaller's type being firm quality, monetary policy of a central bank, with type being toughness on inflation, and work incentives, with type being ability. As a model of work incentives it could be seen as a development of Holmstrom's model [15]. In that model the worker signals his ability, but without knowing his own ability - which makes strategies simpler. Fudenberg and Tirole [14] refer to this type of model as a "signal-jamming" model. A combination of signal, the signaller's action, and random noise is observed at each stage. In a particular specification of incentives, with work and ability being perfect substitutes, and with normal distributions of noise, there is a Stackelberg result. A repeated signalling model does the same type of thing with a more standard approach to signalling, and allowing for more general specifications of payoffs and information.

Mailath and Samuelson ([21]) have also studied reputational issues in terms of dy-

[2] That is to say, they take actions that make the lowest type not want to mimic them.

namic signalling. In a two-type model they find that reputation effects can be supported with the high type separating from the low type. In common with the repeated signalling model, the effect is generated by separation from bad type(s) rather than pooling with a "good" type. The importance of types being changeable over time is also emphasized. There is no clear relation to Stackelberg actions and payoffs: their interest is in qualitatively supporting reputational concerns and also seeing how reputation can be built up and lost gradually.

From static to repeated signalling

Signalling games have been a fruitful area in pure and applied work, beginning with Spence's model of education and job-market signalling [28]. See Sobel [27] for a survey of applications. In the canonical monotonic signalling model the signaller has a type, which is private information, and takes an action in a space embedded in $\mathbb{R}$ or $\mathbb{R}^n$. Types are a subset (finite or continuum) of some real interval, and higher types have more of a preference for higher actions, a single-crossing condition. The respondent observes the action, forms a belief about the type, and replies accordingly, treating higher types more favourably - a preference given by external considerations. In Spence's example, types who find education easier are given a higher wage because they are expected to do better work: this is not modelled but is the reason for the respondent's preference. In the repeated signalling model both players move simultaneously in the stage game[3] and as described above, which allows the study of reputation, and there are reputational reasons for higher types being treated more favourably. Beyond the usual assumptions of monotonic signalling games there is currently an additive separability assumption: the signaller's stage-game payoff is separable between the respondent's action and the signaller's type and action. This provides uniqueness of equilibrium in the stage game and allows simple characterization of the solution and allows discounted Stackelberg payoffs to be defined simply.[4]

There are many perfect Bayesian equilibria of the one-shot signalling game, some separating and some pooling, and the most used equilibrium of the signalling game is the Riley equilibrium, which is minimally separating. That is to say no two types take the same action with positive probability, and each type takes his most preferred action given the requirement of separating from lower types (so that no lower type would want to mimic him). In a finitely repeated signalling game, applying this property at every stage, starting with the last, gives what I call the iterated Riley equilibrium, in which types separate minimally at each stage. This means that at every stage the current type of the signaller is revealed. (Type is changeable but correlated over time, following a Markov process). This equilibrium has a particularly simple form, with the signaller's actions depending only on his current type and how many periods he is from the end of the game, and the respondent responding myopically (regardless of his actual discount

[3] Although this has not been formally shown, there will be no change to any results if the respondent moves first.

[4] Work is in progress to relax this assumption.

factor) to current expectations of play. To calculate the signaller's strategy in a given period, we only need to find the minimally separating equilibrium between his current action and the respondent's response in the next period. This response depends on the signaller's strategy in the next period and it is an inductive calculation.

From a theoretical angle, attempts were made to cut down the number of equilibria with restrictions on beliefs off the equilibrium path, some beliefs being considered more reasonable than others. Particularly successful are the set of related refinements that go by the name "divinity", including D1, defined in Cho, Kreps [5].[5] In this paper I consider a refinement, labelled D_ω, which is an extension of divinity to the repeated signalling game. The spirit of the refinement D_ω is this: sub-optimal actions by the signaller are interpreted as over-confidence, over-confidence about the respondent's response to these actions. I define the justifying beliefs of an action to be those beliefs about the respondent's immediate response that would justify the action over the equilibrium action. And the criterion D_ω is that if one type has a smaller set of justifying beliefs for a particular action than another type, then the first type is ruled out, assuming that the second type was given positive probability before the action was observed. Suppose, informally, that larger belief-mistakes are infinitely less likely than smaller ones, uniformly across types. Then any type requiring a larger belief-mistake (overconfidence) to justify an action than another type must be assigned infinitely less probability than this other type. The criterion D_ω is a weakening of this condition.

Cho, Sobel [6] show that in monotonic signalling games the D1 criterion selects the Riley equilibrium uniquely (assuming pooling at the highest action is ruled out). A similar logic is used here to show that D_ω uniquely selects the iterated Riley equilibrium. There are two steps. First pooling is eliminated. At any point in the game if two types pool on the same action, it is shown that by taking a slightly higher action, the signaller is considered to be at least the higher type, so payoffs increase discontinuously on raising the action from this point. Second, separating is shown to be minimal. If a type does more takes a more costly action than necessary to separate from the preceding type then changing the action slightly is shown not to affect beliefs, so the original action cannot be optimal.

When type-change from period to period becomes very unlikely I find limit properties as the number of periods from the end of the game tends to infinity, calculating a limit map from types of the signaller to actions. When types become dense in an interval we get the modified Stackelberg result given above. This happens when the respondent has preferences over actions of both players, and not over the type of the signaller, so that signalling incentives derive entirely from reputational concerns. When the respondent has preferences over actions and the signaller's type, rewarding both high expected actions and high types, I find that the limit map from types to actions is characterized by a differential equation. This is the same differential equation as in the Riley equilibrium of

[5] Other criteria less connected to the current work include the *intuitive criterion* of Cho, Kreps [5], which provides a unique solution when there are only two types, strategic stability [18], which is defined on finite action spaces, the weak condition of "*undefeated equilibium*" (Mailath, Okuno-Fujiwara, and Postlewaite [20]), and evolutionary stability ([23]).

the single-stage game when the signaller moves first, so is already a Stackelberg leader, with a different starting point. Thus the limit combines commitment and pure type-signalling motives, commitment motives arising from reputational considerations.

Contents

Chapter 2 defines the model and states its main assumptions. Chapter 3 defines perfect-Bayesian equilibrium and the D_ω refinement. I define the iterated Riley solution in chapter 4, and show how it results uniquely from the D_ω refinement in chapter 5. In chapter 6 limit properties are found as type change becomes very unlikely and as the number of periods tends to infinity. Limit properties of the reputation case are derived and discussed in chapter 7, and limits of the general case in chapter 8. Chapter 9 proposes further work.

1.2 Model

Actions, types and utilities

There are two players and k periods. In each period both players take actions simultaneously; actions are observable. The signaller takes actions from the set $A \subseteq \mathbb{R}$, $A = [a_{\min}, a_{\max}]$; the respondent takes actions from $R \subseteq \mathbb{R}$, $R = [r_{\min}, r_{\max}]$.

The respondent has no private information; he has a type in each period (a "period-type") which determines both players' payoff functions in the stage game in that period. The signaller in each period knows his current and previous period-types. Each period-type lies in a finite set $T \subseteq \mathbb{R}$, $T = \{\tau_0, ...\tau_h\}$. Let the global type, the vector of period-types, of the signaller be $t^k \in T^k$; the signaller's period-type in period i is then t_i. A sub-vector of types $t^n \in T^n$ describes period-types in periods 1 to n.

The signaller has the discounted utility function $U_1 = \sum_i \delta_1^i u_1 (t_i, a_i, r_i)$ from outcomes $O = (T \times A \times R)^k$ to $\mathbb{R}$, with u_1 a continuous function $T \times A \times R \to \mathbb{R}$ and with $0 < \delta_1 \leq 1$.

The respondent has utility function $U_2 = \sum_i \delta_2^i u_2 (t_i, a_i, r_i) : O \to \mathbb{R}$, with u_2 a continuous function $T \times A \times R \to \mathbb{R}$ and with $0 < d_2$.

Assumptions on u_2

Assumption 1 *$\int u_2(., r) d\mu_{ta}$ is strictly quasi-concave in r for any probability measure μ_{ta} on $T \times A$.*

The above integral is continuous by continuity of u_2 and so has a maximum in r for each probability measure μ_{ta}. The quasi-concavity assumption ensures that there is a unique maximum. Call this maximum $r^*(\mu_{ta})$, the myopic best response of the respondent to the belief μ_{ta}.

Definition 1 *For any measure μ_{ta} on $T \times A$, let* $r^*(\mu_{ta}) = \arg\max \int u_2(.,r)d\mu_{ta}$

Assumption 2 $\mathrm{Im}(r^*) \subseteq (r_{\min}, r_{\max})$

Assumption 3 *Increasing response to types or actions:*

u_2 is differentiable in the third argument and $(\partial/\partial r)u_2(t,a,r)$ is strictly increasing in (t,a).

Here $(t_1,a_1) < (t_2,a_2)$ iff $t_1 \leq t_2$ and $a_1 \leq a_2$ with at least one inequality strict.

Assumptions 2 and 3 imply the following fact, which is their only role in this paper: fixing a map between types and actions, if the distribution of types increases in the sense of first order stochastic dominance, then the myopic best response of the respondent will increase.

Fact 1 *If $\alpha : T \to A$ is a strictly increasing function, and $f : T \to T \times A$ by $f(\tau) = (\tau, \alpha(\tau))$, and if $\mu_t < \mu'_t$ in the sense of first order stochastic dominance, then $r^*(\mu_{ta}) < r^*(\mu'_{ta})$, where $\mu_{ta} = \mu_t \circ f^{-1}$ and $\mu'_{ta} = \mu'_t \circ f^{-1}$.*

Proof. $r^*(\mu_{ta})$ maximizes the quasi-concave function of r, $\int u_2(t,a,r)\,d(\mu_t \circ f^{-1})(t,a)$, so $(\partial/\partial r)\int u_2(t,a,r)\,d(\mu_t \circ f^{-1})(t,a) = \int(\partial/\partial r)u_2(t,a,r)\,d(\mu_t \circ f^{-1})(t,a) = 0$ at $r^*(\mu_{ta})$, so $\int(\partial/\partial r)u_2(t,a,r)\,d(\mu'_t \circ f^{-1})(t,a) > 0$ at $r^*(\mu_{ta})$, but $= 0$ at $r^*(\mu_{ta})$, so we must have $r^*(\mu_{ta}) < r^*(\mu'_{ta})$ by quasi-concavity of $\int u_2(t,a,r)\,d(\mu'_t \circ f^{-1})(t,a)$. ■

Say that any function $\widehat{r} : \Delta(T \times A) \to R$ for which the above property holds satisfies *increasing response to types or actions.*

Assumptions on u_1

Assumption 4 *Additive separability:* $u_1(t,a,r) \equiv v_a(t,a) + v_r(r)$

Assumption 5 *v_r is strictly increasing*

The first monotonicity assumption for the signaller, assumption 5 requires that higher actions by the respondent are preferred by the signaller. It is equivalent to $u_1(t,a,.)$ being a strictly increasing function for each t,a.

Assumption 6 *Single crossing: If $a_1 < a_2$ and $t_1 < t_2$ then $v_a(t_2,a_2) - v_a(t_2,a_1) > v_a(t_1,a_2) - v_a(t_1,a_1)$*

The second monotonicity assumption, what assumption 6 expresses is that higher types of the signaller are more disposed to taking higher actions.

It follows from this assumption that if $a_1 < a_2$ and $t_1 < t_2$ then $u_1(t_1, a_1, r_1) \leq u_1(t_1, a_2, r_2)$ implies $u_1(t_2, a_1, r_1) < u_1(t_2, a_2, r_2)$; and this is condition that will be used in this paper

Assumption 7 *$v_a(t, a)$ is strictly quasi-concave in a for all t.*

This is equivalent to $u_1(t, a, r)$ being strictly quasi-concave in a for all t and r. This assumption will be used to give unique solutions to optimization problems by the signaller. In particular, fixing r, there is a be a unique action for any type which maximizes $v_a(t, a)$: this is implied by quasi-concavity and continuity of v_a. We will call this action $a^*(t)$.

Definition 2 $a^*(t) = \arg\max_a v_a(t, a)$

Assumption 8 *Undesirable $a_{\min}$ and $a_{\max}$:*

For each $t \in T$, $t > \tau_0$, $v_a(t, a)$ is not maximized at $a_{\min}$.

For any $r_1, r_2 \in \mathrm{Im}(r^*)$, $t \in T$, $v_a(t, a^*(t)) + v_r(r_1) > v_a(t, a_{\max}) + v_r(r_2)$

The undesirable $a_{\max}$ assumption is that no change in the respondent's action (within the myopic best-response set $\mathrm{Im}(r^*)$) will compensate for taking the action $a_{\max}$ over $a^*(t)$.

It is important to eliminate the possibility of pooling at the highest action because then my game structure in which types are revealed each period breaks down. Also, without pooling there will be a simple map from types to actions in a given period, independent of the current type distribution, while the type-action correspondence for a pooled equilibrium depends on the type distribution.

The undesirable $a_{\min}$ assumption is needed to ensure that the map from types to actions is always strictly increasing, required to generate the limit reputation result. If the respondent has a concern for type as well as actions then we do not need this assumption as the map will be strictly increasing from the penultimate stage back. (It is possible that a reworking of the limit results will eliminate the need for either of these alternatives.)

Fact 2 *If the above assumptions 4-8 on u_1 are satisfied by $u_1 = v_a(t, a) + v_r(r)$ then they are satisfied by $u^E = v_a(t, a) + \delta_1 v_r(r)$ for $0 < \delta_1 \leq 1$.*

Histories, strategies and beliefs

The histories after the i^{th} period are $H_i := (A \times R)^i$. The whole space of histories up to the last period is the disjoint union $H := \bigsqcup_0^{k-1} H_i$.

The respondent observes past play, so his strategy in period i is a function of H_{i-1}. His global strategy is a function of the space of histories H. Take this to be a behaviour strategy, giving a mixed action in $\Delta(R)$ at every history: his stragegy is a function $s_2 : H \to \Delta(R)$ such that $s_2(.)(\sigma)$ is a measurable function $H \to R$ for any measurable $\sigma \subseteq R$. Throughout this paper $\Delta(X)$ for any measure space X denotes the space of probability measures on X.

The signaller, in addition to observing past play, knows his current and previous period-types. So define: $HT_i := (A \times R)^i \times T^{i+1}$; $HT = \bigsqcup_0^{k-1} HT_i$. ("Histories with types".) His strategy is a function $s_1 : HT \to \Delta(A)$ satisfying: $s_1(.)(\alpha)$ is a measurable function $HT \to R$ for any measurable $\alpha \subseteq A$.

Since we are dealing with continuous action spaces for both players the measurability assumptions above are needed to be able to define the progress of the game given the strategies. See the section below on the outcome of the game and the corresponding definitions in the appendix to see why this is so.

The respondent at any history $h_i \in H_i$, $i < k$, has a belief $\beta(h_i) \in \Delta\left(T^{i+1}\right)$ about the signaller's types up to that point.

There is an exogenously given distribution of types in which there is correlation between types from one period to the next. This process will be assumed to be Markov. This correlation will give the motive for the signaller to signal a higher type in a given period: by signalling a higher type he will be thought to be a higher type in the next period. A special case is when types are constant across periods. But we are particularly interested in processes which have full support, so that given any type distribution in a given period, the type distribution in the next period has full support. The equilibrium refinement that I will propose solves the game under this assumption.

The regeneration process is described by the function $\Psi : \sqcup_i T^i \to \Delta(T)$. If types from periods 1 to i are described by $t^i \in T^i$, $\Psi(t^i)$ describes the distribution of types in period $i+1$.

Assumption 9 *Monotonic Markovian type-change:*

$\Psi(.)$ is a Markov process, generated by the function $\psi : T \to \Delta T$ and the initial distribution $\Psi(())$.

$\psi(t)$ is strictly increasing in t in the sense of first-order stochastic dominance.

Assumption 10 *Type regeneration: $\Psi\left(t^i\right)$ has full support for any $t^i \in T^i$.*

The outcome of the game

Define an outcome of the game to be a vector of actions of each player and period-types of the signaller, i.e. an element of $O = (T \times A \times R)^k$, describing the entire progress of the game. Once we have strategies s_1, s_2 and the type regeneration function Ψ we can define from any point in the game ht_i the probability distribution of subsequent play and the probability distribution of outcomes. Call the first distribution the continuation of the game $C^+(s_1, s_2)(ht_i) \in \Delta(T^i \times A^i \times R^i)$ and the second distribution the completion of the game $C(s_1, s_2)(ht_i) \in \Delta\left(T^k \times A^k \times R^k\right)$.

Even though these are familiar notions in game theory, more care than usual needs to be taken since we are dealing with continuous action spaces. Formal definitions of C and C^+ are given in the appendix. The measurability conditions on s_1 and s_2 are needed here for C and C^+ to be defined.

Perfect-Bayesian Nash equilibrium

We shall be examining perfect-Bayesian Nash equilibria of the game described above. Although the concept is standard, I give a formal definition which respects the particular construction of this model.

Definition 3 (s_1, s_2, β) *is a perfect-Bayesian Nash equilibrium if:*

1. For each history with types ht_i, s_1 maximizes $\int U_1 d[C(s_1, s_2)(ht_i)]$

2. At any history h_i, s_2 maximizes:

$\int \int U_2 d[C(s_1, s_2)(h_i, t^{i+1})] d[\beta(h_i)(t^{i+1})]$

3. For any $h_i \in H_i$, $\beta(h_i)$ satisfies:

3.a) $\beta(h_i)$ respects Bayes' rule where applicable:

If $\beta(h_i)$ gives positive probability to some vector $t^{i+1} \in T^{i+1}$ of period-types in the first $i+1$ periods, and $s_1(h_i, t^{i+1})$ gives positive probability to action a, then for any r, $\beta(h_i, (a, r))$ is the usual Bayesian update of $\beta(h_i)$.

3.b)$\beta(h_i)$ respects Ψ between period i and period $i+1$:

$\beta(h_i)(\{t^{i+1}\}) = \beta(h_i)(\{t^i\} \times T).\Psi(t^i)(t_{i+1})$

The integral in 1. is expected utility for the signaller. The integral in 2. is the expected utility of the respondent given that (h_i, t^{i+1}) is reached, integrated over beliefs about t^{i+1}.

Types with the same period-types up to period i but different future period-types behave in the same way up to period i. Condition 3.b) means that even when a zero-probability event is observed by the signaller, he should not doubt the regeneration process Ψ but given his assessment of the period-types in periods 1 to i his assessment of future period-types will be consistent with Ψ.

1.3 Equilibrium selection

The refinement D_ω

The proposed refinement is based on this motivation: out of equilibrium actions that are sub-optimal for all types are considered to be mistakes made by a type in his perception of the response to those actions.[6] The type considered to be making the mistake a is thought to become over-confident about the respondent's immediate response to a and that leads him to do a rather that his optimal equilibrium action. The set of beliefs about responses to a that would cause type t to play a are called the (strictly/weakly) justifying beliefs of a for t. Larger mistakes are considered infinitely more improbable than smaller mistakes, uniformly across types, leading to a type who would have had to make a large error in his perception being considered infinitely less likely than a type who would have had to make a smaller error in order to take action a. The probability of errors is itself infinitely small, and no types make errors in equilibrium. This informal argument supports the following refinement D_ω: if the set of weakly justifying beliefs for type t^* is contained in the set of strictly justifying beliefs for type t^{**}, and t^{**} was considered possible (assigned probability > 0) before a was observed, then t^* is given posterior probability 0. The condition implies that beliefs have to move further from the correct equilibrium beliefs for action a to be justified for t^* than for t^{**}.

Justifying Beliefs

Given a perfect-Bayesian Nash equilibrium (s_1, s_2, β), define at a history h_n the justifying beliefs of an action a for a player with $t^{n+1} \in T^{n+1}$ as follows:

Let $u^* = \int U_1 d[C(s_1, s_2)(h_n, t^{n+1})]$ be the expected utility of the optimal strategy s_1 by the signaller.

Given an action a at history h_n, the strategy of the respondent in the next period is described by some function $\widetilde{r} : R \to \Delta R$ giving the mixed response after the current actions (a, r). If s_2' is a strategy of the respondent, the corresponding function is the function which takes r to $s_2'(h_n, (a, r))$. The reason we describe the respondent's next-period strategy as a function is that he could potentially condition his next-period action

[6] It is also possible to define a refinement based on utility loss, as follows: if a type t_n is assigned positive probability at the beginning of a period and t_n' is some other type and if t_n' would lose more utility from taking the action a than t_n then t_n' is assigned probability 0 after a is observed. Propositions 1 and 2 below still hold and the structure of all arguments can remain the same: additivite separability makes these two refinements work in a similar way. While this refinement has the advantage of simplicity, D_ω is potentially generalizable to the non-additive case and also does not rely on cardinal utility.

on his current action and his current action could be mixed. $\widetilde{r}(.)(\pi)$ will be measurable for measurable π. Now given such an $\widetilde{r}: R \to \Delta R$, define the strategy $s_2'(\widetilde{r})$ as follows: $s_2'(\widetilde{r}) = s_2$ except at $(h_n, (a, r))$ for all r, and $s_2'(\widetilde{r})(h_n, (a, r)) = \widetilde{r}(r)$. This strategy by the respondent is the same as the original one but changed in period $n+1$ to respond to $(h_n, (a, r))$ with $\widetilde{r}(r)$ for any r. So it is changed in the next-period response to the action a by the signaller. Now let $\alpha(a)$ be the set of strategies s_1' of the signaller with $s_1'(h_n, t^{n+1}) = a$. Then $u(\widetilde{r}) = \sup_{s_1' \in \alpha(a)} \int U_1 d[C(s_1', s_2'(\widetilde{r}))(h_n, t^{n+1})]$ is the maximum utility of the signaller in response to s_2', conditional on having to play a at the current history.

Definition 4 *For each binary relation* $\rhd \in \{>, \geq, =\}$, $J^{\rhd}(t^{n+1}, h_n, a) :=$

$$\{\widetilde{r}: R \to \Delta R \text{ with } \widetilde{r}(.)(\pi) \text{ measurable for measurable } \pi, \text{ such that } u(\widetilde{r}) \rhd (u^*)\}$$

Call $J^>$ *the strictly justifying beliefs,* $J^{\geq}$ *the weakly justifying beliefs, and* $J^=$ *the barely justifying beliefs.*

Note that an action a is optimal if the correct belief about the respondent's response given a -the equilibrium strategy - is a barely-justifying belief.

Given h_n call the set $\{a : J^>(t^n, h_n, a) \neq \{\}$ for some $t_n\}$ the justifiable actions. These are the actions that are justified for some type by some possibly erroneous belief about the respondent's response. Some actions may not be justified by any belief, and the respondent's beliefs when confronted with these actions will not be specified by the D_ω criterion below.

Definition of D_ω

Definition 5 *A perfect Bayesian equilibrium* (s_1, s_2, β) *satisfies* D_ω *if:*

For any history $h_n = (h_{n-1}, a_n, r_n)$ *and types* t_1^n, t_2^n,

if $\beta(h_{n-1})$ *assigns positive probability to* t_2^n,

and if $J^{\leq}(t_1^n, h_{n-1}, a_n) \subseteq J^<(t_2^n, h_{n-1}, a_n) \neq \emptyset$,

then $\beta(h_n)$ *assigns probability 0 to* t_1^n.

1.4 The Iterated Riley solution

The Riley map

Consider the standard one-shot monotonic signalling game in which the signaller moves first. Imagine that the signaller has utility $u^E(t, a, r) = v_a(t, a) + \delta_1 v_r(r)$, and the respondent response to the signaller's perceived type is given by a stricly increasing function

$r'' : T \to R$. The Riley equilibrium of this signalling game is the perfect Bayesian equilibrium in which types separate minimally. Separation implies that the lowest type τ_0 must take his myopic optimal action $a^*(\tau_0)$. Each subsequent type takes his myopic optimal action, subject to separating from lower types. Given the monotonicity assumptions, it is sufficient to require each type to separate from the previous type. Define $RILEY(r'')$ to be this equilibrium, specifying an action for each type. Given any strictly increasing function $r'' : T \to R$, $RILEY(r'')$ is defined inductively as follows:

Definition 6 $RILEY(r'') : T \to A$

$$RILEY(r'')(\tau_0) := a^*(\tau_0)$$

$$RILEY(r'')(\tau_i) := \arg\max_{a \in B_i} v_a(t_i, a), \text{ where}$$

$$B_i = \left\{ a \in A : u^E(\tau_{i-1}, RILEY(r'')(\tau_{i-1}), r''(\tau_{i-1})) \geq u^E(\tau_{i-1}, a, r''(\tau_i)) \right\}$$

Assumption 5 ("undesirable $a_{\max}$") guarantees that the set of actions B_i for which a lower type would not want to pretend to be the current type is non-empty. Existence and uniqueness of the arg max above is guaranteed by single crossing and strict quasi-concavity of u^E. (As we saw earlier, u^E must satisfy Assumptions 4-8 since $u_1 = v_a(t, a) + v_r(r)$ does.)

Each type's action is strictly higher than the previous type's by monotonicity (single crossing and preference for higher responses). So by single crossing, if each lower type does not strictly prefer take the subsequent type's action, then any lower type does not strictly prefer to take the action of a higher type. This is why in the function $RILEY$ above it was sufficient to require each type to separate himself from the previous type.

It turns out that the repeated signalling game can be solved by repeated use of the $RILEY$ function. If the signaller after his i^{th} move is thought to be type t, the respondent's action in period $i+1$ will be a function $r''(t)$ of this t. If we only look at actions of the signaller in period i and of the respondent in period $i+1$ we have utility for the signaller given by u^E; this explains the use of the modified utility function u^E in the definition above. Then given the response function r'', the signaller will take Riley separating equilibrium actions $RILEY(r'')$ in period i.

Note that the Riley equilibrium is defined without reference to any distribution of types of the signaller. This fact is very important for analysis of the repeated game and generates history-independence for the signaller.

The Iterated Riley solution

Under assumptions 1-9 we can now define the "Iterated Riley equilibrium", a description of play of both players on the equilibrium path. Assumption 10 (full support) will be used later on to justify the Iterated Riley solution uniquely; it is not necessary to define it. In the Iterated Riley solution the signaller's strategy is a function only of his current type and the stage of the game. His action is given by $\sigma_1 : \{0, ...k-1\} \to T \to A$. $\sigma_1(j)(t_{k-j})$ will define the action of type t_{k-j} of the signaller in the $(k-j)^{th}$ period.

Let $f(\sigma) : T \to T \times A$, $f(\sigma)(\tau) = (\tau, \sigma(\tau))$, so that if σ is a map from period-types to actions, $f(\sigma)$ gives the type and action pair for any type.

σ_1 is defined inductively as follows:

Definition 7 $\sigma_1(0) = a^*$

Given $\sigma_1(j)$*,* $\sigma_1(j+1) := RILEY\left(r''_{k-j}\right)$*, where* $r''_{k-j}(\tau) = r^*\left(\psi(\tau) \circ f(\sigma_1(j))^{-1}\right)$*.*

$f(\sigma_1(j))$ represents the map from types to type-action pairs in period $k-j$. Given type τ was believed to have been the signaller's type in period $k-j-1$, the beliefs about the type in period $k-j$ will be $\psi(\tau)$ and the belief about the type-action pair will be $\psi(\tau) \circ f(\sigma_1(j))^{-1}$. $r''_{k-j}(\tau)$ will be the myopic optimal action of the respondent in period $k-j$, given that the signaller is thought to have been type τ in period $k-j-1$.

Assuming that $\sigma_1(j)$ is a strictly increasing function, r''_{k-i} is a strictly increasing function by assumption 3 (increasing responses to types and actions) and assumption 9 (monotonic Markovian type change) and strictly increasing $\sigma_1(j)$, allowing $RILEY\left(r''_{k-j}\right)$ to be defined, which gives a strictly increasing function $\sigma_1(j+1)$. This justifies the definition.

The Iterated Riley solution can now be defined in terms of σ_1:

Definition 8 s_1, s_2 *are an Iterated Riley equilibrium if for histories on the equilibrium path:*

$s_1\left((a_1, ...a_{i-1}), (r_1, ...r_{i-1}), (t_1, ...t_i)\right) = [\sigma_1(k-i)(t_i)]$*, where* $[a]$ *is the degenerate probability measure placing all weight on* a*.*

$s_2\left((a_1, r_1) ... (a_i, r_i)\right) = r^*(\psi(\sigma_1(k-i)^{-1}(a_i)) \circ f(\sigma_1(k-(i+1)))^{-1})$ *for* $i \geq 1$

$s_2(()) = r^*\left(\Psi(()) \circ f(\sigma_1(k-1))^{-1}\right)$

To understand the nature of the Iterated Riley solution, consider first these conditions on s_1, and s_2 given σ_1. The signaller's strategy (on the equilibrium path) is described very simply by σ_1: $\sigma_1(j)$ gives the map from types to actions in the period $k-j$. It is independent of previous play and only dependent on the period and the current period-type. The respondent's strategy has a more involved definition. In period $i+1$, he looks at the signaller's last action, a_i. Since we are on the equilibrium path this will be in the image of $\sigma_1(k-i)$. Since σ_1 is strictly increasing it is injective and so only one type $\sigma_1(k-i)^{-1}(a_i)$ will ever take that action. Beliefs about the period-type in the next period should be[7] given by $\psi(\sigma_1(k-i)^{-1}(a_i))$. Since $\sigma_1(k-(i+1))$ gives the map

[7] I have not specified beliefs in the Iterated Riley solution, and they are mentioned here as an aid to understand the definition.

from types to actions in the current period $i+1$, the expected type-action pair will be $\psi(\sigma_1(k-i)^{-1}(a_i)) \circ f(\sigma_1(k-(i+1))^{-1})$. The respondent's action is the myopic response r^* to this. In the first period, the expected type distribution is $\Psi(())$ and the respondent's action is then $r^*(\Psi(()) \circ f(\sigma_1(k-1))^{-1})$.

Now consider the definition of σ_1. In the last period, subject to no signalling motives, the signaller takes the myopic optimal action given by the function $\sigma_1(0) = a^*$. If after period $i = k-j$ the signaller is believed to have period-i-type τ, he can expect the response $r^*(\psi(\tau) \circ f(\sigma_1(j))^{-1})$. There is a minimal separating equilibrium looking only at actions in the current period and responses in the next, given by the $RILEY$ map applied to this response function and using utility u^E with the response discounted by the discount factor δ_1.

It is useful to specify a map F that gives σ_{i+1} in terms of σ_i. Let the space of strictly increasing functions from T to A be $Inc(T, A)$.

Definition 9 $F : Inc(T, A) \to Inc(T, A)$,

$$F(\sigma) := RILEY\,(r''),\ \textit{where}\ r''(\tau) = r^*(\psi(\tau) \circ f(\sigma)^{-1}).$$

Then we have $\sigma_1(i) = F^i a^*$.

Note that the correspondence between the signaller's type and his action in the period j periods from the end is the same across games with a varying number of periods, all other specifications constant.

1.5 The Iterated Riley equilibrium and the D_ω condition

Here I will show the existence and uniqueness of the Iterated Riley equilibrium as a perfect Bayesian-Nash equilibrium satisfying D_ω.

Supportability of the Iterated Riley solution

Proposition 1 *Under assumptions 1 to 9, iterated Riley solution is supportable as a perfect-Bayesian Nash equilibrium satisfying D_ω.*

Proof. See appendix. ■

Note that the Assumption 10 (full support) is not necessary to support the Iterated Riley solution as a perfect-Bayesian Nash equilibrium satisfying D_ω.

A particular Bayesian-Nash equilibrium is defined explicitly in the proof which satisfies the required properties. It has these properties:

At any point in the game (not only on the equilibrium path but at all histories with types) the signaller takes an action given by σ_1. At a history h_i in which the signaller has taken actions $(a_1, ...a_i)$: respondent's belief about the signaller's period-i type is $[\tau_j]$ if $a_i = \sigma_1(k-i)(\tau_j)$. The respondent's belief in any period is a monotonic function of the previous action, and is always supported on a single type. Beliefs about the period-type t_i are unchanged after period i, on and off the equilibrium path. Beliefs about the period types after period i at history h_i are deduced from the Markov process ψ. See the section below on uniqueness for an explanation of why these beliefs satisfy D_ω. Given these beliefs, the respondent then acts myopically based on his beliefs about the type and action he can expect in the current period. This is because the action that he takes has will have no effect on the future course of the game.

Now suppose that separation is from previous types is always binding. Then beliefs have a particularly simple form: if a_i lies in $[\sigma_1(k-i)(\tau_j), \sigma_1(k-i)(\tau_{j+1}))$ beliefs about the type are still $[\tau_j]$: the respondent assumes it is the lower type making a mistake and taking too high an action rather than a higher type taking too low an action. Below $\sigma_1(k-i)(\tau_0)$ the type is believed to be $[\tau_0]$ and above $\sigma_1(k-i)(\tau_h)$ the type is believed to be $[\tau_h]$.[8]

There will be other equilibria than the one checked that satisfy D_ω. But D_ω does specify the equilibrium up to responses to unjustifiable actions. The signaller's strategy must be given by σ_1 for a D_ω equilibrium. The respondent must respond and form beliefs as above after a *justifiable* actions; after an (out-of-equilibrium) unjustifiable action he may form any beliefs and act accordingly.

Uniqueness

Proposition 2 *Under Assumptions 1 to 10, in a perfect-Bayesian Nash equilibrium of the model described in section 1 satisfying D_ω: The signaller's strategy depends only on the period and his type in that period via the function σ_1 defined above, by the equation $s_1\left((a_1, ...a_{i-1}), (r_1, ...r_{i-1}), (t_1, ...t_i)\right) = [\sigma_1(k-i)(t_i)]$.*

The respondent's strategy satisfies $s_2\left((a_1, ...a_i), (r_1, ...r_i)\right) = r^*(\psi(\sigma_1(k-i)^{-1}(a_i)) \circ f(\sigma_1(k-(i+1)))^{-1})$ whenever $a_i \in \mathrm{Im}(\sigma_1(k-i)^{-1})$.

And $s_2(()) = r^*\left(\Psi(()) \circ f(\sigma_1(k-1))^{-1}\right)$.

Proof. See Appendix. ■

Note that this equation for the signaller now holds at every history, not only on the equilibrium path. The proposition implies an equilibrium satisfying D_ω must be an Iterated Riley equilibrium.

[8] Cho, Sobel [6] claim that beliefs of this form always generate a D1 equilibrium of the single-stage monotonic signalling game. This is not quite true: it is only true when separation from previous types is a binding constraint in the Riley equilibrium.

Two facts about D1 equilibria should be called to mind to understand how the Iterated Riley solution is selected by the D_ω criterion. Firstly as discussed earlier the Riley equilibrium selected is independent of the initial type distribution. A second and related fact is that the beliefs of the respondent are categorical and regardless of the initial distribution assign probability 1 to some type [9]. The logic of D1 is strong enough to outweigh any disparities in the probabilities of initial types: to express this in terms of the intuitive understanding given above of the divinity criterion, a larger mistake is infinitely less likely uniformly across types than a smaller one, so if one type would require a larger mistake to justify an observed action than another, then the latter is considered infinitely more probable, and so the first type is given probability 0 regardless of how much more likely he was than the second type before the action was observed.

This same logic applies for D_ω in the repeated game. We will have at every stage a single-stage signalling game and regardless of the history at any particular stage - regardless of the current type-distribution ascribed to the signaller by the respondent - there will be the same map from types to action given by the Riley equilibrium. And while previous action by the signaller will alter the type-distribution expected in a given period, beliefs by the respondent after the current action will be a function of that action only and will be categorical in nature, ascribing probability 1 to a particular type.

The game is solved from the last period and the above logic applied at every stage. Each action by the signaller in period i is paired with the respondent's action in the next period. In the last period there is no signalling incentive, and the signaller takes his myopic optimal action a^*. In period i for the signaller and period $i+1$ for the respondent, given that the game has been solved for the remainder of the game (periods $i+1$ on for the signaller and periods $i+2$ on for the respondent) and generated the Iterated Riley solution there, we can analyze the action in period i and response in period $i+1$ in isolation. This game will be monotonic because the respondent rewards the signaller for signalling a higher type (see the definition of the Iterated Riley solution). The analysis of this restricted game is like the analysis of the one-stage signalling game under D1. Separation comes from the fact that if two players were to pool in equilibrium, by taking slightly higher actions each could discontinuously increase the beliefs about him to beliefs whose support has a minimum of at least the higher type. And minimal separation comes from the fact that if a type were to take an action in equilibrium that is not his myopic optimum given that he has to separate, then by moving to this myopic optimum conditional on separation, he will (at least) maintain beliefs about him, and increase his current period payoff.

[9] For all important actions of the signaller: the "justifiable" ones in my terminology.

1.6 Limit properties of the Iterated Riley equilibrium

First it is useful to note the continuity of the solution with respect to the various primitives defining it.

Fact 3 *The Iterated Riley solution $\sigma_1(i) : T \to A$, for each i, is continuous as a function of v_a, v_r, δ, ψ, r^*.*

To see this, observe that the function $RILEY(r'') : T \to A$ is continuous as a function of u^E and r''. $F : Inc(T, A) \to Inc(T, A)$ is then continuous when considered as a function of v_a, v_r, δ, ψ. a^*. So $\sigma_1(i) = F^i a^*$ is continuous as a function of v_a, v_r, δ, ψ, r^*.

Now let ψ_0 be the degenerate type regeneration function, with $\psi_0(t) = [t]$. Note that the full support assumption was not used in the definition of the iterated Riley solution. We have seen that if we have ψ tend to the degenerate function ψ_0 in which types remain the same with probability 1, the iterated Riley solution will tend to the Iterated Riley solution with $\psi = \psi_0$. Let us now consider the properties of the iterated Riley solution with regeneration function $\psi = \psi_0$, as the number of periods from the end i tends to infinity, for fixed v_a, v_r, δ, r^*.

First it is useful to define "discounted Stackelberg" utility. The undiscounted utility of the signaller with period-type τ in any period if he can and does commit to the action a and is known to be type τ is: $v_a(\tau, a) + v_r(r^*([(\tau, a)]))$. We can call this Stackelberg utility. If the discount factor of the signaller is not 1 it will be more useful to consider the "discounted Stackelberg" utility: $v_S(\delta_1)(\tau, a) := v_a(\tau, a) + \delta_1 v_r(r^*([(\tau, a)]))$. This is the utility for type τ of the action a taken in the current period and plus the discounted utility of the best response in the next period to the type-action pair (τ, a).

Given δ_1 and τ, call the maximum value of this the discounted-Stackelberg payoff (which exists by continuity of all functions involved), and the unique a that maximizes the expression (unique by the concavity assumption on $v_r^*(a)$) the discounted-Stackelberg action $a_S(\delta_1)(\tau)$.

Assumption 11 *$v_S(\delta_1)(\tau, a)$ is strictly quasi-concave in a for each $\tau \in T$.*

This assumption that discounted Stackelberg utility is strictly quasi-concave is important to the limit analysis. It is satisfied for example when v_a and v_r are strictly concave and $r^*([(\tau, a)])$ is linear in a for each τ. In a work incentives example where a is work and $r^*([(\tau, a)])$ is market wage this would be a natural specification.

Proposition 3 *$\sigma_1(i)$ tends to a limit Σ_1 as $i \to \infty$.*

Σ_1 is characterized as follows:

1. $\Sigma_1(\tau_0) = a^*(\tau_0)$

2. Let h *be the highest solution for* x *of:*

$v_a(\tau_j, x) + \delta_1 v_r(r^*([(\tau_{j+1}, x)])) = v_S(\delta_1)(\tau_j, \Sigma_1(\tau_j))$.

(There are one or two solutions.)

$\Sigma_1(\tau_{j+1}) = \max(h, a^*(\tau_{j+1}))$

Proof. See Appendix ■

The proof involves inductive application of a dynamical systems argument. I will explain here some features of the process generating σ_1. Suppose that for a particular type τ_j, $\sigma_1(i)(\tau_j)$ tends to a limit $\Sigma_1(\tau_j)$ as $i \to \infty$. If $\sigma_1(i)(\tau_{j+1})$ also tends to a limit $\Sigma_1(\tau_j)$ it must satisfy $\Sigma_1(\tau_{j+1}) = \max(x, a^*(\tau_{j+1}))$, for some x for which $v_a(\tau_j, \Sigma_1(\tau_j)) + \delta v_r(r^*([(\tau_j, \Sigma_1(\tau_j))])) = v_a(\tau_j, x) + \delta v_r(r^*([(\tau_{j+1}, x)]))$.

This is because $\sigma_1(\tau_{j+1})$ is either eventually given by a binding constraint of separation from the previous type, or by the myopic optimum $a^*(\tau_{j+1})$. The first is the "normal" case; the second is a failure of signalling to have any effect due to types that are too far apart.

Now consider the equation in x. The first part $v_a(\tau_j, \Sigma_1(\tau_j)) + \delta v_r(r^*([(\tau_j, \Sigma_1(\tau_j))]))$ is the converged period-utility of type τ_j. The second part $v_a(\tau_j, \Sigma_1(\tau_{j+1})) + \delta v_r(r^*([(\tau_{j+1}, \Sigma_1(\tau_{j+1}))]))$ is the converged utility of pretending to be $\Sigma_1(\tau_{j+1})$. Assuming that the need to separate is a binding constraint, these two must be equal. But even if we know that $\sigma_1(i)(\tau_{j+1})$ converges we have not found what it converges to yet because this equation may have more than one solution: it may have one or two solutions. The lower solution lies below $\Sigma_1(\tau_j)$ if the respondent has a direct preference for rewarding higher types, and so in this case we can rule it out because the limit map from types to actions must be weakly increasing. But in the reputation case where the respondent does not care directly about the signaller's type it is not so easy to rule out the lower solution. It may be the case that both $\Sigma_1(\tau_j)$ and some higher action are solutions to the equation above. For a description of how this is resolved and the higher solution is chosen, see the section on reputation below.

Note that the convergence is not monotonic: this has been confirmed by numerical computation of an example.

1.7 Reputation

Now make the assumption that the respondent does not care directly about the signaller's type, only about his action:

Assumption 12 $u_2(t, a, r)$ *is a function of* a *and* r *only*

It follows that $r^*(\mu_{ta})$ only depends on the probability distribution over actions. Define $v_r^*(a) := v_r(r^*(\mu_t * [a]))$, which is independent of μ_t.

Let the highest action that gives the same Stackelberg utility for type τ as action a be $\overline{a}_S(\delta_1)(\tau, a)$.

Corrolary 1 *$\sigma_1(i)$ tends to a limit Σ_1 as $i \to \infty$.*

Σ_1 is characterized as follows:

1. $\Sigma_1(\tau_0) = a^*(\tau_0)$

2. For each j, $\Sigma_1(\tau_{j+1}) = \max(\overline{a}_S(\delta_1)(\tau_j, \Sigma_1(\tau_j)), a^*(\tau_{j+1}))$

This follows simply from proposition 3, noting that $r^*([(t, x)])$ is independent of t and so that h in proposition 3 is equal to $\overline{a}_S(\delta_1)(\tau_j, \Sigma_1(\tau_j))$. In words, if $\Sigma_1(\tau_j)$ is weakly above the discounted-Stackelberg action of type τ_j, then $\Sigma_1(\tau_{j+1}) = \Sigma_1(\tau_j)$, assuming this is above the myopic-optimal action $a^*(\tau_{j+1})$. Otherwise $\Sigma_1(\tau_{j+1})$ jumps up above the discounted-Stackelberg action of type τ_j to $\overline{a}_S(\delta_1)(\tau_j, \Sigma_1(\tau_j))$.

Assume that separation is binding (as I show in the proof of proposition 4, this will be true if types are close together). Let us continue the discussion of proposition 3 and examine why if $\Sigma_1(\tau_j)$ is below the discounted-Stackelberg action, $\Sigma_1(\tau_{j+1})$ is equal to $\overline{a}_S(\delta_1)(\tau_j, \Sigma_1(\tau_j))$ and not the lower $\Sigma_1(\tau_j)$. Both are solutions of the equation $v_a(\tau_j, \Sigma_1(\tau_j)) + \delta v_r(r^*([(., \Sigma_1(\tau_j))])) = v_a(\tau_j, \Sigma_1(\tau_{j+1})) + \delta v_r(r^*([(., \Sigma_1(\tau_{j+1}))]))$. There are two facts that combine to give this result. Firstly at every stage $\sigma_1(i)$ is strictly monotonic: there is separation of types. (A separation that does not always occur in the limit as we have seen.) Secondly $\Sigma_1(\tau_j)$ is below the Stackelberg action of type τ_j (this is the case we are considering): it follows that for type τ_j small increases in expectations about his action are more valuable than small increases in his action are painful, and so in order to be thought to be taking the action $\Sigma_1(\tau_j) + x$ for small x, type τ_j would be willing to increase his action to $\Sigma_1(\tau_j) + y$ where $y \geq x$. This is the logic that generates type τ_{j+1}'s action: what would type τ_j be willing to do in order to be thought of as taking type τ_{j+1}'s action. So if $\sigma_1(i)(\tau_j)$ is close to $\Sigma_1(\tau_j)$ and $\sigma_1(i)(\tau_{j+1}) - \sigma_1(i)(\tau_j)$ is small, then $\sigma_1(i+1)(\tau_{j+1}) - \sigma_1(i+1)(\tau_j)$ must be larger. The first difference is the increase in expected action that type τ_j will gain in pretending to be type τ_{j+1}; the second is the increase in action that is necessary. This means that $\sigma_1(i)(\tau_j)$ can never become close to $\sigma_1(i)(\tau_{j+1})$ and is pushed away from $\Sigma_1(\tau_j)$.

Now we can see what happens when types become dense: the main reputation result is for this case. Suppose that u_1 and u_2 are defined continuously over an interval $\overline{T} = [\tau_{\min}, \tau_{\max}]$ and satisfy the relevant assumptions above. Define the function S as follows:

Definition 10 $S : \overline{T} \to A$

$S(\delta_1)(\tau) = \max(\overline{a}_S(\delta_1)(\tau_{\min}, a^*(\tau_{\min})), a_S(\delta_1)(\tau))$

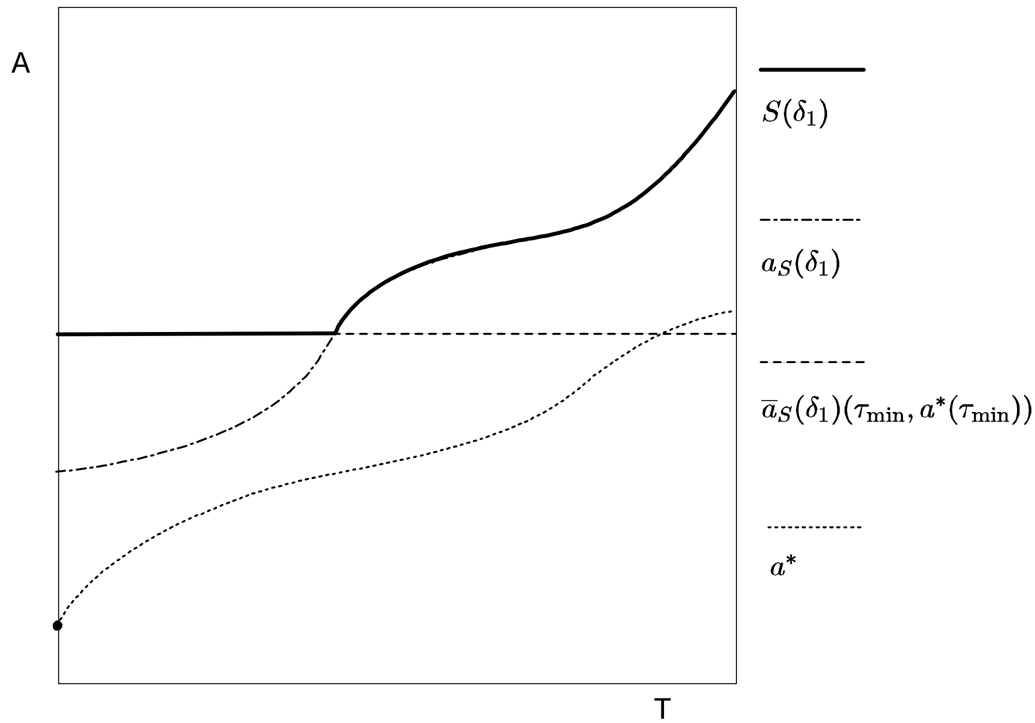

Figure 1.1: the limit map $S(\delta_1)$

I.e. $S(\delta_1)(\tau)$ gives maximum discounted Stackelberg utility to type τ over the set of actions that give type $\tau_{\min}$ at most the discounted Stackelberg utility of $a^*(\tau_{\min})$. See the diagram Fig 1. Proposition 4 asserts that as types become dense, the limit map from actions to types $\Sigma_1$will converge to $S(\delta_1)$, apart from the lowest type who must take the action $a^*(\tau_0)$.

Proposition 4 *Given $\eta > 0$, we can find $\epsilon > 0$ such that for any finite set $T \subset \overline{T}$ with* $\max_{\overline{\tau}} \min_{\tau}(\{|\tau - \overline{\tau}|, \tau \in T, \overline{\tau} \in \overline{T}\}) < \epsilon$*, the limit solution Σ_1 satisfies*

$|\Sigma_1(\tau) - S(\delta_1)(\tau)| < \eta$ *for $min(T) \neq \tau \in T$.*

Proof. See appendix ■

My result is that as the number of periods from the end tends to infinity the given limit holds. The model studied is finitely repeated and over games with different numbers of periods but the same specifications otherwise play is determined by the number of periods from the end. An implication is that for any levels of patience, as the number of periods $k \to \infty$, play in period p converges to the limit found. If the number of periods is large, reputation will take a long time to die out.[10]

[10]"Reputation effects" actually are strong up to the penultimate period. However the given limit properties are only realized further back in the game.

Discussion

A modified Stackelberg property

The reputation result above combines a separation property for low types with a (discounted) Stackelberg property and logic. The separation property is that types above the lowest type must separate from the lowest type making him unwilling to move from his myopic optimal action $a^*(\tau_{\min})$ and pretend to be a high type, where this willingness is evaluated with (discounted) Stackelberg utility. And types whose (discounted) Stackelberg actions lie above this point take these actions. One can think of the actions that are in $Im(\Sigma)$ - actions that are taken in the limit - as the actions that the signaller can commit to: by taking an action that is in (or more exactly close to something in) this set far from the end of the game, he we be expected to take (close to) the same action in the next period. Thought of in this way the reputation result is that the limit Σ exists and $Im(\Sigma)$ becomes dense in $[\overline{a}_S(\delta_1)(\tau_{\min}, a^*(\tau_{\min}), a_S(\delta_1)(\tau_h)]$ but has a gap in $(a^*(\tau_{\min}), \overline{a}_S(\delta_1)(\tau_{\min}, a^*(\tau_{\min}))$ and these actions no-one can commit to. This results in the lower types pooling at $\overline{a}_S(\delta_1)(\tau_{\min}, a^*(\tau_{\min}))$.

This result is distinguished from standard reputation models in that these models will just generate a Stackelberg property for the normal type or types: the actions that a normal type can effectively commit to are those actions which get played by behavioural types and these tend to be assumed to include the Stackelberg action of the normal type.

Depending on the context both parts of the curve $S(\delta_1)$ may be interesting, or only one of the two, pooling or Stackelberg. If we think, following a line of thought that is found in the standard reputation literature, that some of our types are "normal" (probable) and others improbable, and that we can imagine a type that is so low that he would rather take his myopically optimal action than commit to a Stackelberg action of a normal type, and we give this low type some positive bur low probability, then our "normal" types will take their Stackelberg actions.

In the lower part of the curve $S(\delta_1)$, types pool at a point determined by separation from the lowest type. This action is higher than the actions that they would like to commit to. But by taking an action even slightly lower than this action, they pay a heavy cost: they are considered to be the lowest type. If they take at least the action $\alpha = \overline{a}_S(\delta_1)(\tau_{\min}, a^*(\tau_{\min}))$, they will be expected to take at least this action in the next period, while if they take an action less than α they will be expected to do $a^*(\tau_{\min})$ in the next period. This is a discontinuity that makes taking at least the action α very important. It is appropriate to call α a *reputational standard*, a mark that it is important to reach in order to prevent one's reputation from being destroyed altogether - at least for the next period which is as long as reputations last in this model.

This *reputational standard*, a novel consequence of the repeated signalling model, can potentially be used to explain various situations in which there is a standard of behaviour that can be thought of as a standard necessary to live up to in order maintain a reputation. For example obedience to some social (legal or moral) or business norms can be understood as a requirement for establishing that one is not a bad (criminal or untrustworthy or undependable) type. The point is that these norms are often not continuous but discrete: either one complies with them or one does not.

Let me offer some potential examples in more detail. High actions by the signaller could represent good behaviour by and the respondent could be society; low actions of society could be imprisonment for protection of society, high actions the ability to participate fully in society. Low types are criminal; high types are upstanding citizens. The lowest type will commit crimes; higher types follow the "norm" of the society, which is the particular standard of behaviour, determined endogenously α by the need to separate from the most criminal type. Many people will follow this norm and do no more than this: they follow it because they do not want to be considered the criminal type. And there may be high types that do more both out of natural inclination and the rewards of being thought of as especially trustworthy and good to deal with.[11] Business norms may be thought of in the same way, with the low type being the laziest or least trustworthy. Relatively lazy types may for example prefer to work shorter hours than 9 to 5, and would be willing to take cuts in pay to do this, but they do not because the reputation cost of not coming up to the fixed standard is steep.

Technical comparison to other models of reputation

The repeated signalling model is fully calculable at all histories of the game. While this is true of some standard reputation models ([19]) it is rare and usually even the actions of the reputation-builder in equilibrium are often not fully specified: limit results tend to be in the form of lower bounds on payoffs of the normal type rather than convergence of his behaviour. By contrast the modified Stackelberg property above specifies the signaller's action in the limit.

The calculability of the model extends to arbitrary discount rates of both players. The respondent's discount factor has no effect on the course of the game. The signaller's discount factor does and the reputation results are given in terms of δ_1. A simple modification of Stackelberg payoffs into "discounted Stackelberg payoffs" is all that is required. This generality is very unusual in the reputation literature, which invariably requires that the patience of the reputation-builder goes to 1. Stackelberg results obtain when the respondent is short-lived ([12]) and there exist limit results in the case when the uninformed player's patience tends to 1 but when the informed player's becomes at the same time infinitely more patient than the uninformed ([11]): $(\delta_2, (1-\delta_1)/(1-\delta_2)) \to (1, 0)$. The special case of strictly conflicting interests ([7]) is an exception, as is the reputational

[11] In contracts, Socrates, justice is of use. - *Plato, The Republic*

It may be more realistic to assume that in addition to treating a person favourably for being expected to behave well, society will be well disposed to a person who is thought to be "good". The repeated signalling model can still deal with this situation. We can still use the same mathematical definition of the *reputational standard* and it will still be the action necessary to avoid being thought of as the worst type. However there will no longer be pooling at this action in the limit.

bargaining model of Abreu, Gul [1], in which both discount factors tending to 1 with differences in patience tending to a limit. In general, however, reputation results require the informed player to become infinitely patient and infinitely more patient than the uninformed player. When the opponent is long lived with a fixed discount factor player may be able to establish reputations for complex strategies under certain conditions and do better than the static Stackelberg payoff ([8]). This does happen in the repeated signalling model because reputation is established along one dimension only.

Standard reputational models are often completely general in the stage games studied, while the repeated signalling model analyzes only a class of games in which stage game payoffs are monotonic and additively separable. But within the class I define the model can be completely solved and the question of what happens for any levels of patience of both players addressed, questions which are not addressed in the standard literature. I find that a reputation can be established against a patient player, even by a player that is less patient. And I find that a "discounted Stackelberg" result applies when the informed player is not patient (subject to separation from the lost type). The discounted Stackelberg action is a novelty and just as easy to calculate as the Stackelberg action and can easily be applied to situations in which players are thought of as impatient. One implication of the discounted Stackelberg result for high types is that a small reduction in the discount factor from 1 has a second order effect on (limit) payoffs of the informed player.

The way in which reputation is established is quite different in the repeated signalling model from the standard approach to reputation. Reputation is a one-period property of the repeated signalling model, with the expectation of the respondent being based on the previous action of the signaller, and the signaller can gain or lose it immediately at any time. This happens because it is easier, given the full support assumption, for a type to change than to make a (larger) mistake. For a discussion of signalling without type regeneration, in which this logic does not apply, see section 1.9. The property that reputation can be gained or lost at any point is shared with reputational models with imperfect observability but with much more sudden gains and losses. Without imperfect observability, in standard reputation models, either reputation is lost immediately if at all (revelation of the normal type) or the play from any point in the game tree may be very unknown.

The nature of the types I consider to be an advantage of the repeated signalling model. The commitment types of reputation models are often considered to be an unsatisfactory element, out of place in a theory based on strategy and rationality. On the other hand it has been argued[12] as a genericity assumption they make models involving them at small levels more reasonable than purely "rational" models without. My view is that including behavioural types at small levels is an unobjectionable and valid method, but that the order of the limits involved in reputational models restricts how small the probabilities of behavioural types can be to be effective. The results are found in general under a limit as the informed player becomes infinitely patient for a given probability of behavioural

[12] This argument is made I believe by Fudenberg; I will give a reference and exact quotation when I have located the article.

types. If this probability is very small, the required patience may be very large indeed. If we look at the set of discount factors which result in payoffs a certain distance from the Stackelberg (assuming a model that gives a Stackelberg result) as a function of the probability p of behavioural types, we only know that this set contains $(\delta(p), 1)$ for some $\delta(p)$, and this could tend to the empty set as $p \to 0$. For any specified situation with a given high level of patience we will need the probability of behavioural types to be high enough to justify applying a reputation result to expect actions that are near Stackelberg.

The workings of the equilibrium differ from those in reputational models in that the normal type in reputational models pools with commitment types while in the repeated signalling model "normal types" separate from each other in each period.

Mailath and Samuelson [21] they have a model of reputation in which the lowest type's action is fixed and the higher type establishes a reputation by separating himself from the lower type. They find that the higher type will take higher actions than he would otherwise, which may be higher or lower than the Stackelberg action. This holds true in this model with two types: here the lowest type's action is effectively fixed, although he is not a behavioural type, and the higher type separates and may take an action that is more or less than the Stackelberg action depending on the distance between the types. But the most interesting results in the repeated signalling model come from having a large number of types. While with two types separation from the lowest type determines the answer: with more types this model gives both the logic of separation from the lowest type (for low types) and Stackelberg actions (for high types).

1.8 The general case

Suppose now that $u_2(t, a, r)$ depends on all three arguments so that $r^*([t, a])$ is a function of both t and a. Consider for example the work-incentives model above. The observable productivity a of the worker (signaller) could be measured by the market (respondent) as quantity of writing, or some other other easy and imperfect measure. Suppose that t is the ability of the worker, with more able workers being better able to produce more writing. It is a reasonable assumption that a worker who produces a given amount of writing has a value to the market that is an increasing function of his ability. The wage $r^*([t, a])$ paid will then be strictly increasing in both t and a.

Assumption 13 *$r^*([(\tau, a)])$ is continuously differentiable in (τ, a) with both partial derivatives strictly positive. $v_r(r)$ and $r^*[(\tau, a)]$ are continuously differentiable. $v_a(\tau, a)$ is differentiable with respect to a, with derivative continuous in (τ, a).*

Definition 11 $G : \overline{T} \to A$

$$G(\delta_1)(\tau_{\min}) := \overline{a}_S(\delta_1)(\tau_{\min}, a^*(\tau_{\min}))$$

$$G(\delta_1)'(\tau).[\tfrac{\delta}{\delta a} v_a(\tau, a) + \delta_1 \tfrac{\delta}{\delta a}(v_r \circ r^*)([(\tau, a)])] + \delta_1 \tfrac{\delta}{\delta \tau}(v_r \circ r^*)([(\tau, a)])$$

The limit will now be given by G instead of S. $G(\delta_1)$ is differentiable and lies above the discounted Stackelberg curve $a_S(\delta_1)$.

Proposition 5 *Given $\eta > 0$, we can find $\epsilon > 0$ such that for any finite set $T \subset \overline{T}$ with $\max_{\overline{\tau}} \min_{\tau}(\{|\tau - \overline{\tau}|, \tau \in T, \overline{\tau} \in \overline{T}\}) < \epsilon$, the limit solution Σ_1 satisfies*

$|\Sigma_1(\tau) - G(\delta_1)(\tau)| < \eta$ for $min(T) \neq \tau \in T$.

Proof. See Appendix ■

Consider the one-stage signalling game with the same specifications with the signaller moving first and utility given by u^E. In the reputation case we get the discounted Stackelberg action as the solution since the signaller is a Stackelberg leader and type-inference has no significance. In the general case the solution is given by H, say, where H satisfies the differential equation $H'(\tau).[\frac{\delta}{\delta a}v_a(\tau, a) + \frac{\delta}{\delta a}(v_r \circ r^*)([[(\tau, a)]])] + \frac{\delta}{\delta \tau}(v_r \circ r^*)([[(\tau, a)])$ with initial condition $H(\tau_{\min}) = a_S(\delta_1)(\tau)$. The differential equation is the same as the differential equation in the limit above but the initial condition is the discounted Stackelberg action of the lowest type rather than the higher $\overline{a}_S(\delta_1)(\tau_{\min}, a^*(\tau_{\min}))$.

The standard one-stage signalling game can be thought of as combining commitment and pure type-signalling - commitment trivially because the respondent observes the signaller's action before moving and "pure" type-signalling because given an known action the signaller would still want to be thought as being a higher type. The limit of the repeated signalling game with simultaneous moves also combines commitment and pure type-signalling. The two solutions satisfy the same differential equation but the initial condition (describing types close to the lowest) is higher in the repeated signalling game because the lowest type takes his myopic optimal rather than his discounted Stackelberg action. Thus apart from the lowest type there is more costly signalling in the repeated game with simultaneous actions than in the one-stage game with the signaller moving first.

1.9 Further work and extensions

The fixed type case

If there is no type-regeneration what will happen in the repeated signalling game? The Iterated Riley solution can be defined in this case, but is no longer selected uniquely by the equilibrium refinement D_w. Moreover in an Iterated Riley solution for many histories (off the equilibrium path) the respondent will believe that the signaller's type has changed. If we were to require that the respondent's beliefs always assign probability 0 to type-change, then the Iterated Riley solution is ruled out. (Here the difference is the difference between having a type space T^n with types staying the same with probability 1 and having a type space T. The space of Bayesian Nash equilibria is different because the restrictions on beliefs after probability zero events are different.)

The criterion D_w will be weak because it relies on full support at a given stage for its strength: if a type has probability zero before period i and but has a larger set of justifying beliefs for a given action than any other type, D_w does not specify what beliefs will be after that action is observed. Another criterion would do better. One possibility that I have partly analyzed is measuring for a given type and history the error at each period of the types actions, combining these errors into a real-valued total error via some norm, and specifying that the respondent's beliefs about the signaller's type after a given history have support in the set of types with the least total error. I find that if there is a "reasonable" solution in the sense that higher actions generate higher beliefs about the type and a continuity property holds, then all signalling must happen in the first period: that is to say, on the equilibrium path, myopic actions are taken except in the first period. Equilibrium play is then independent of the norm and measure of error above. However the existence and necessity of such an equilibrium have however not yet been shown.

Kaya [16] has a model of repeated signalling in which the signaller has a fixed type. Rather than using any refinement she calculates the "least cost" separating equilibrium, the separating equilibrium where it exists that is most preferred by all types. Kaya finds that when types are ranked in the convexity of their payoff functions, there will be such an equilibrium, and in this equilibrium signalling will be either spread out or all in the first period depending on the direction of the convexity ranking. The signaller moves first in the stage game and so the model does not study reputational issues.

Multi-dimensional A

If the action set A of the signaller is a product of intervals in $\mathbb{R}^n$ rather than $\mathbb{R}$, it will be necessary to find an assumption on v_a that gives a (uniquely) definable and strictly increasing function $RILEY\left(r''\right)$ for any strictly increasing $r'' : T \to R$. If this can be found then propositions 1 and 2 go through with no changes. It will then be interesting to see what the limit properties are. Preliminary work indicatest that if the methods can be extended to $\mathbb{R}^n$ there will be a curve $T \to A$ and the current Stackelberg result will apply when restricted to this curve, but that the curve will be defined by a differential equation unrelated to any commitment property. In the one dimensional case it is the necessity of separating from the previous type at every stage that determines the solution and gives the reputation property. In more dimensions there will be a whole range of actions at any stage that just separate from the previous type; the exact specification of the subsequent type's utility on this set, which will give his Riley action, then becomes significant.

Two-way uncertainty, non-additive utility

Additive utility is a reasonable specification for some situations but not others. Current analysis indicates that what is needed is an assumption that gives a unique equilibrium at each point in time in the actions of both players. When a good assumption is found the game is solveable by backward induction as before, with monotonicity and single crossing being preserved at each stage. The history-independence property will no longer hold.

Two way uncertainty in the repeated signalling model will require the action of each player to be both dependent on his type and to be a reward for a higher expected type or action of the other player. With additive separability for both players the model will involve two-way type-signalling only. A potential model is reciprocation, with types representing something like good-will and both players being made more generous by their own and the other player's revealed good will. Current work studies such a model, although outside of a limited informational setting it turns out there are difficult problems with preserving monotonicity as the game is solved backwards.

Two-way uncertainty without the additivity assumption would open up a large class of models, with the potential to study reputational incentives on both sides. Very interesting two-way reputational models in have been studied in the context of bargaining ([1], [2]). Applications of two-way signalling models include oligopoly and work incentives in teams.

A compact set of types

Ramey[25] solves the one-stage signalling game using D1 with any compact set of types in $\mathbb{R}$, generalizing Cho, Sobel [6]. The repeated signalling can be defined and extended in the same way and propositions 1 and 2 (existence and uniqueness) will hold. It will involve a more complex notation and an adapted version of the Riley equilibrium. The arguments for the limit results will need a new approach, but it is possible that something in the spirit of the current inductive proof will work. The benefit of compact continuum of types is that the limit should be exactly the function S and we should not need to take a limit as types become dense. And a continuum of types will often be a more natural specification of a given repeated signalling situation, with a continuum being a natural way to model types with more or less of a certain predisposition, and a continuum of potential types being part of the limit reputation result (u_1 and u_2 being defined on the real interval $\overline{T}$ of types). It is unlikely that this generalization will add new insights into repeated signalling and reputation.

1.10 Appendix

Definition 12 $P^{HT \to H} : HT \to H$ *projects from histories with types to histories in the natural way.*

$P^{HT \to \sqcup T^j} : HT \to \sqcup T^j$ *projects from histories with types to vectors of types in the natural way.*

$P_j^{T^i} : T^i \to T$ *for* $i \geq j$ *projects to the* j^{th} *period-type.*

Definition 13 *The completion of the game $C(s_1, s_2) : HT \to \Delta\left(A^k \times R^k \times T^k\right)$, given strategies of each player, is defined as follows:*

1. *Let $C(s_1, s_2)(ht_i)[\{\tau \times \left(A^{k-j} \times R^{k-j} \times T^{k-(j+1)}\right)]$ over measurable sets τ of $\left(A^j \times R^j \times T^{j+1}\right)$ be denoted $\Phi_j(\tau)$.*

Φ_j for $j \in \{i, ...k-1\}$ is defined inductively by:

a. *$\Phi_i([(a_1, ...a_i), (r_1, ...r_i), (t_1, ...t_{i+1})]) = 1$. I.e. the existing history happens with probability 1 in the completion.*

b. *If α, β, γ are measurable in A, R, T respectively,*

$$\Phi_{j+1}[\tau \times (\alpha \times \beta \times \gamma)] = \int_{ht_j \in \tau} s_1(ht_j)(\alpha).s_2\left(P^{HT \to H} ht_j\right)(\beta).\Psi\left(P^{HT \to \sqcup T^j} ht_j\right)(\gamma) d\Phi_j(ht_j).$$

This defines a probability measure Φ_{j+1} given Φ_j.

2. *Then over measurable sets τ of $\left(A^{k-1} \times R^{k-1} \times T^k\right)$, and for α, β are measurable in A, R,*

$$C(s_1, s_2)(ht_i)[\tau \times \alpha \times \beta] = \int_{ht_{k-1} \in \tau} s_1(ht_{k-1})(\alpha).s_2\left(P^{(HT \to H)} ht_{k-1}\right)(\beta) d\Phi_{k-1}(ht_{k-1}).$$

This defines the probability measure $C(s_1, s_2)(ht_i)$.

The integrals above exist by the measurability assumptions on s_1 and s_2. They define measures on the appropriate product spaces.[13] 1. defines Φ_j inductively from $j = i$ to $j = k-1$. 2 defines the probability measure $C(s_1, s_2)(ht_i)$ given Φ_{k-1}.

Definition 14 *The continuation play $C^+(s_1, s_2)$ of the game, given strategies of each player, is defined as follows:*

$C^+(s_1, s_2) : HT \to \sqcup\Delta\left(A^i \times R^i \times T^i\right)$, with $C^+(s_1, s_2)(ht_j) \in \Delta\left(A^{k-j}, R^{k-j}, T^{k-j}\right)$

$C^+(s_1, s_2)(ht_i) := C(s_1, s_2)(ht_i) \circ P^{-1}$, where P projects $(A \times R \times T)^k$ onto the last $k - i$ coordinates $(A \times R \times T)^{k-i}$.

Proof. Proof of Proposition 1: Supportability of the iterated Riley solution as a perfect Bayesian equilibrium satisfying $\mathbf{D}_\omega$

Define s_1, s_2, β as follows:

For any $ht_i = ((a_1, ...a_i), (r_1, ...r_i), (t_1, ...t_{i+1}))$, $h_i = ((a_1, ...a_i), (r_1, ...r_i))$:

$s_1(ht_i) = [\sigma_1(k-(i+1))(t_{i+1})]$.

[13] For uniqueness of the defined measure, note that we have defined the measure on all product sets, which are a π-system generating the product σ-algebra. For existence, we have defined a measure on product sets. Extend additively to a measure on the class of finite disjoint unions of product sets (uniquely). Applying the monotone convergence theorem, this measure is countably additive on this class, which is a ring of sets generating the product σ-algebra, and so extends to a measure on the product σ-algebra by Caratheodory's extension theorem.

Define $r''_i(\tau) := r^*(\psi(\tau) \circ f(\sigma_1(k-(i+1)))^{-1})$.

Define $\underline{r}_i(\tau, a) := \inf\{r : u^E(\tau, a, r) \geq u^E(\tau, \sigma_1(k-(i+1))(\tau), r''_i(\tau))\}$. (The infimum is taken over the set R so that $\inf(\emptyset) = r_{\max}$.). $\underline{r}_i(\tau, a)$ will be the minimum response in the next period that would justify action a for type τ.

For $j \leq i$, let $\beta_j(h_i) = [\sup(\arg\min_\tau \underline{r}_j(\tau, a_j))]$. This defines beliefs over the type in the first i periods. Beliefs about the full k-period type are generated from these beliefs by ψ.

Note that $\underline{r}_i(\tau, \sigma_1(k-(i+1))(\tau)) = r''_i(\tau)$ trivially, while by the definition of the Riley equilibrium, $\underline{r}_i(\tau', \sigma_1(k-(i+1))(\tau)) \geq r''_i(\tau)$ with strict inequality for $\tau' > \tau$: any higher type would strictly lose on moving to the action of a lower type. So if $a_j = \sigma_1(k-j)(\tau)$ then $\beta_j(h_i) = [\tau]$.

Let $s_2(h_i) = r''_i(\beta_i(h_i))$, player 2's myopic best response response to the action a_i in period n given above beliefs.

Now I will show that (s_1, s_2, β) is a perfect Bayesian equilibrium. It follows that it is an Iterated Riley equilibrium and so the proposition is proved.

Claim 1 *(s_1, s_2, β) is a perfect Bayesian equilibrium*

Optimality of s_1

Let $\widehat{r}_{j+1}(a_i)$ be player 2's action in period $j+1$ in response to player 1's action a_j in period j: $\widehat{r}_j(a) := r''_j(\sup(\arg\min_\tau \underline{r}_j(\tau, a)))$.

Player 1's utility at history ht_i given player 2's strategy above, when he takes strategy s'_1 is

$K_1 + \delta^{i+1}[v_r(\widehat{r}_{i+1}(a_i)))] + \delta_1^{i+1}[v_a(t_{i+1}, a_{i+1}) + \delta_i v_r(\widehat{r}_{i+2}(a_{i+1}))] + ... + \delta_1^{k-1}[v_a(t_{k-1}, a_{k-1}) + \delta_i v_r(\widehat{r}_k(a_{k-1}))] + \delta_1^k[v_a(t_k, a_k)]$, integrated over $C(s'_1, s_2)(ht_i)$, where K is a constant (utility up to period i) independent of s'_1.

This expression equals $K_2 + \delta_1^{i+1}[v_a(t_{i+1}, a_{i+1}) + \delta_1 v_r(\widehat{r}_{i+2}(a_{i+1}))] + ... + \delta_1^{k-1}[v_a(t_{k-1}, a_{k-1}) + \delta_1 v_r(\widehat{r}_k(a_{k-1}))] + \delta_1^k[v_a(t_k, a_k)]$, where $K_2 = K_1 + \delta^{i+1}[v_r(\widehat{r}_{i+1}(a_i)))]$ is independent of s'_1.

I will show that $a_j = \sigma_1(k-j)(\tau)$ maximizes $[v_a(\tau, a_j) + \delta_1 v_r(\widehat{r}_{j+1}(a_j))]$ for any τ. It follows from this that player 1's strategy maximizes each component $[v_a(t_j, a_j) + \delta_1 v_r(\widehat{r}_{j+1}(a_j))]$ of the expression above since it puts probability 1 on $a_j = \sigma_1(k-j)(t_j)$. And so it maximizes expected utility of player 1 after any history with types ht_i.

By construction of $\sigma_1(k-j)$, $\widehat{r}_{j+1}$ as the Riley equilibrium with utility $v_a + \delta_1 v_r$, $a_j = \sigma_1(k-j)(t_j)$ maximizes $[v_a(t_j, a_j) + \delta_1 v_r(\widehat{r}_{j+1}(a_j))]$ over $\mathrm{Im}(\sigma_1(k-j))$. (As shown above, $\widehat{r}_{j+1}(\sigma_1(k-j)(\tau)) = r''_j(\tau)$.) It needs to be shown that $a_j = \sigma_1(k-j)(t_j)$ maximizes the expression over all A.

Suppose that for type τ_m, action $a = \alpha$ gives a higher value of $v_a(\tau_m, a) + \delta_1 v_r(\widehat{r}_{j+1}(a))$ than $a = \sigma_1(k-j)(\tau_m)$. Then this is also true for any type with a lower value of $\underline{r}_i(\tau, a)$.

So take without loss of generality $\tau_m = \sup(\arg\min_\tau \underline{r}_j(\tau, \alpha))$. Then $\widehat{r}_i(\alpha) = r_i''(\tau_m)$ given beliefs of player 2 specified above.

We know that $\sigma_1(k-j) \geq a^*$. If $\sigma_1(k-j)(\tau_m) \leq \alpha$ then $a^*(\tau_m) \leq \sigma_1(k-j)(\tau_m) \leq \alpha$ and by quasi-concavity of v_a we must have decreasing v_a above $a^*(\tau_m)$. So $v_a(\sigma_1(k-j)(\tau_m)) \geq v_a(\alpha)$. The actions $\sigma_1(k-j)(\tau_m)$ and α both generate the same response by player 2 and so the former gives a (weakly) higher value of $v_a(\tau_m, .) + \delta_1 v_r(\widehat{r}_{j+1}(.))$. This contradicts our assumption, so we must have $\alpha < \sigma_1(k-j)(\tau_m)$.

Now suppose that $\alpha < \sigma_1(k-j)(\tau_{k-1}) < \sigma_1(k-j)(\tau_k)$. Since τ_k prefers $\sigma_1(k-j)(\tau_k)$ and the corresponding response to $\sigma_1(k-j)(\tau_{k-1})$, and the single crossing condition holds between α and $\sigma_1(k-j)(\tau_{k-1})$, we must have $\underline{r}_j(\tau_{k-1}, \alpha) < \underline{r}_j(\tau_k, \alpha)$ contradicting $\tau_m = \sup(\arg\min_\tau \underline{r}_j(\tau, \alpha))$. So τ_k is the minimal type such that $\alpha < \sigma_1(k-j)(\tau_k)$.

We can rule out that $\sigma_1(k-j)(\tau_k) = a^*(\tau_k)$ because if this were so moving below $a^*(\tau_k)$ to α and being thought of as the same type will hurt type τ_k.

So $\sigma_1(k-j)(\tau_k) > a^*(\tau_k)$, which implies by construction of the Riley equilibrium that type τ_{k-1} must exist (that $\tau_k \neq \tau_0$) and that $\sigma_1(k-j)(\tau_k)$ is optimal for type τ_{k-1} (indifference between own action and the next type's).

So by single crossing, comparing $\sigma_1(k-j)(\tau_k)$ with α, we must have $\underline{r}_j(\tau_{k-1}, \alpha) < \underline{r}_j(\tau_k, \alpha)$, which again contradicts the definition of τ_k.

So s_1 must be optimal.

Optimality of s_2

It is sufficient for each action in the support of player 2's strategy at any history to be optimal given the rest of player 2's strategy.

It is clear that player 2's action at any history does not affect any subsequent play either of player 1 or of player 2.

Therefore the myopic best response is optimal.

Consistency of beliefs

1. Bayesian updating: If $t^k = (t_1, ...t_k)$ and $\beta(h_i)(t^k) > 0$, then we must have beliefs assign probability 1 to type $t_1, ...t_i$ and which are generated by ψ afterwards. If some type the action a_{i+1} in period $i+1$ with positive probability then $\sigma_1(k-i-1)(t_{i+1}) = a_{i+1}$ for a unique t_{i+1}. On observing a_{i+1}, player 2 assigns probability 1 to $t_1, ...t_{i+1}$ and beliefs about future types are generated by ψ. This is the Bayesian update on the information that the current period-type is t_{i+1}.

2. $\beta(h_i)$ respects Ψ after period i: by definition.

Therefore (s_1, s_2, β) is a perfect Bayesian equilibrium.

Claim 2 *(s_1, s_2, β) satisfies D_ω.*

Suppose not. Then for some history $h_n = ((a_1, ...a_n), (r_1, ...r_n))$, with history $h_{n-1} = ((a_1, ...a_{n-1}), (r_1, ...r_{n-1}))$ in the previous period, and for some $t_*^n = (t_1^*, ...t_n^*)$ and $t_{**}^n = (t_1^{**}, ...t_n^{**})$ in T^n, $J^{\leq}(t_*^n, h_{n-1}, a_n) \subseteq J^{<}(t_{**}^n, h_{n-1}, a_n) \neq \emptyset$ and $\beta(h_n)$ assigns non-zero probability to t_*^n.

Consider the justifying beliefs for the action a_n for type $t^n = (t_1, ...t_n) \in \{t_*^n, t_{**}^n\}$ at history h_{n-1}.

Let $u^*(t^n) = \int U_1 d[C(s_1, s_2)(h_{n-1}, t^n)]$ be expected equilibrium utility for player 1 of type t^n at the beginning of period n.

Given $\widetilde{r} : R \to \Delta R$ (satisfying the measurability requirement), let $s_2'(\widetilde{r}) = s_2$ except at $(h_{n-1}, (a_n, r))$ for all r with $s_2'(\widetilde{r})(h_{n-1}, (a_n, r)) \equiv \widetilde{r}(r)$. Let $\alpha(a_n)$ be the set of strategies s_1' of player 1 with $s_1'(h_{n-1}, t^n) = a_n$. Then let $u(t^n)(\widetilde{r}) = \sup_{s_1' \in \alpha(a_n)} \int U_1 d[C(s_1', s_2'(\widetilde{r}))(h_{n-1}, t^n)]$, the maximum utility of player 1 in response to s_2' conditional on having to play a at the current history.

Let $s_1' = s_1$ except at (h_{n-1}, t^n) where $s_1'(h_{n-1}, t^n) = a_n$.

Now since $s_2'(\widetilde{r}) = s_2$ from period $n+1$ on, and player 2's action in period n affects player 1's utility additively, s_1' is optimal within $\alpha(a_n)$ against $s_2'(\widetilde{r})$.

Then $C(s_1, s_2)(h_{n-1}, t^n)$ and $C(s_1', s_2'(\widetilde{r}))(h_{n-1}, t^n)$ differ only in the period n actions by player 1, which are $\sigma_1(k-n)(t_n)$ and a_n respectively, and the period $n+1$ responses by player 2, which are $r''(t_n)$ and $\widetilde{r}(s_2(h_{n-1}))$ respectively.

So $u(t^n)(\widetilde{r}) - u^*(t^n) = [\delta_1^n v_a(t_n, a_n) + \delta_1^{n+1} v_r(\widetilde{r}(s_2(h)))] - [\delta_1^n v_a(t_n, s_1(h_{n-1}, t^n)) + \delta_1^{n+1} v_r(s_2(h_{n-1}, s_1(h_{n-1}, t^n), s_2(h_{n-1})))]$, regarding $s_1(h_{n-1}, t^n)$ and $s_2(h_{n-1})$ as elements of A and R since they are degenerate probability measures.

Take $\widetilde{r}^* \in J^{=}(t_{**}^n, h_{n-1}, a_n)$, which is possible since $J^{<}(t_{**}^n, h_{n-1}, a_n) \neq \emptyset$.

$[v_a(t_n^{**}, a_n) + \delta_1 v_r(\widetilde{r}^*(s_2(h_{n-1})))] = [v_a(t_n^{**}, s_1(h_{n-1}, t_n^{**})) + \delta_1 v_r(s_2(h_{n-1}, s_1(h_{n-1}, t_n^{**}), s_2(h_{n-1})))]$, extending v_r here to expected utility over probability measures.

So $[v_a(t_n^{**}, a_n) + \delta_1 v_r(\widetilde{r}^*(s_2(h_{n-1})))] = [v_a(t_n^{**}, \sigma_1(k-(i+1))(t_n^{**})) + \delta_1 v_r(r_i''(t_n^{**}))]$

So $u^E(t_n^{**}, a_n, \widetilde{r}^*(s_2(h_{n-1}))) = u^E(t_n^{**}, \sigma_1(k-(i+1))(t_n^{**}), r_i''(t_n^{**}))$.

Let $\bar{r} \in R$ s.t. $v_r(\bar{r}) = v_r(\widetilde{r}^*(s_2(h_{n-1})))$. $\bar{r}$ exists uniquely by continuity and monotonicity of v_r.

Then $u^E(t_n^{**}, a_n, \bar{r}) = u^E(t_n^{**}, \sigma_1(k-(i+1))(t_n^{**}), r_i''(t_n^{**}))$ so $\underline{r}_i(t_n^{**}, a_n) = \bar{r}$.

Since $\beta(h_n)$ assigns non-zero probability to t_*^n, $J^{\leq}(t_*^n, h_{n-1}, a_n) \subseteq J^{<}(t_{**}^n, h_{n-1}, a_n)$ by assumption and so $\underline{r}_i(t_n^*, a_n) \leq \underline{r}_i(t_n^{**}, a_n)$.

So $\underline{r}_i(t_n^*, a_n) \leq \bar{r}$ and

$u^E(t_n^*, a_n, \bar{r}) \leq u^E(t_n^*, \sigma_1(k-(i+1))(t_n^*), r_i''(t_n^*))$

$u^E(t_n^*, a_n, \widetilde{r}^*(s_2(h_{n-1}))) \leq u^E(t_n^*, \sigma_1(k-(i+1))(t_n^*), r_i''(t_n^*))$

So $\widetilde{r}^* \in J^{\leq}(t^n_*, h_{n-1}, a_n)$.

So $\widetilde{r}^*$ is in $J^{\leq}(t^n_*, h_{n-1}, a_n)$ but not in $J^{<}(t^n_{**}, h_{n-1}, a_n)$, contradicting our assumption.

QED ■

Proof. Proof of proposition 2 (Uniqueness)

The inductive step, proposition $P(i)$, is defined as follow:

For all histories $ht_j \in HT^j$ for $j \geq i$, i.e. from the $(i+1)^{\text{th}}$ period on, player 1's strategy is described by $\sigma_1(j)$, and this strategy is optimal against s_2 even if player 2's strategy is altered from the equilibrium one in the $(i+1)^{\text{th}}$ period only.

It is trivial that in the last period player 1's strategy is described by a^*. So $P(1)$ is true.

Assume $P(i)$.

Then for $j \geq i$, player 2's action at history $h_j \in H^j$ in period $j+1$ when player 1's action in period j $\left(P_j^{H_j \to A} h_j\right)$ is in the image of $\sigma_j(k-j)$ is as in section II.2.ii.

This is because player 2 at the beginning of period j gave a positive probability to all of player 1's possible types in that period:

If the sub-history of h_j at the beginning of period j is $h_{j-1}\left(= P^{H_j \to H_{j-1}} h_j\right)$, 2's beliefs after observing h_{j-1} about the type of player 1 in period $j-1$ are $\mu = \beta\left(h_{j-1}\right)\left(P_{j-1}^{T^k}\right)^{-1}$.

By the definition of perfect-Bayesian Nash equilibrium, $\beta\left(h_{j-1}\right)$ respects ψ after period $j-1$, so:

The beliefs about player 1's j-type after observing h_{j-1} are $\beta\left(h_{j-1}\right)\left(P_j^{T^k}\right)^{-1} = \int \psi(t)(.)d\mu$, which has full support over T since $\psi(t)$ does for each t.

Since only one type takes each action in $\operatorname{Im}\left(\sigma_j\right)$ (by $P(i)$), player 2 observing such an action a in period j assigns probability 1 to the appropriate type $\left(\sigma_j\right)^{-1}(a)$ (as being player 1's j-type), and forms beliefs ν as in section 1.4 about player 1's type and action in period $j+1$.

Since by the assumption $P(i)$ player 2's actions have no effect on player 1's future actions, player 2 acts myopically in each period after i, and so in period $j+1$ takes the myopic best response $r^*(\nu)$ to player 1's expected type-action pair ν.

Now consider a history ht_{i-1}, with type τ_1 at period i. Suppose that at history with types ht_{i-1} we replaced equilibrium strategy for player 1 by $a \in A$, and at history with types $ht_i \oplus (a, r, t)$ for any r, t replaced equilibrium strategy for 1 by $\alpha(r, t)$ for 2 by $\hat{r}(r)$. Assume α is optimal. Then $\hat{r}$ is a justifying belief for type τ_1 at history t_{i-1} if the utility is now at least as great as it was before.

We can take α to be the strategy $\sigma_i(i)$ because we know this is optimal by $P(i)$. Now the strategy of player 1 after period i is fixed, and player 2's actions after period $i+1$ as a function of player 1's type are fixed, independent of a and $\hat{r}$.

Player 1 would then get a continuation utility as a function of a and $\hat{r}$ given by $v_a(\tau_1, a) + s_2(h_{i-1})[v_r] + \delta(\hat{r}(s_2(h_{i-1}))[v_r]) + const1 = v_a(\tau_1, a) + \delta(\hat{r}(s_2(h_{i-1}))[v_r]) + const2$.

So the utility just depends on the direct utility of a via $v_a(\tau_1, a)$ and the expectation of the reward in the next period $\hat{r}(s_2(h_{i-1}))[v_r]$ via the value of the reward v_r.

We now have the level of simplicity of the two-period signalling game where player 1 moves first and player 2 responds.

No pooling:

Suppose at history h_{i-1} action $\bar{a}$ is in the support of two types with period $-\,i$ types τ and τ', with $\tau < \tau'$, where τ' is maximal. Observing $\bar{a}$, player 2 forms beliefs β about player 1's i-period type that are strictly less than $[\tau']$, resulting in a response $\bar{r}$ in the next period. The action $\bar{a}$ is weakly justified for types τ , τ' by the belief $\hat{r}'(r) = [\bar{r}]$ in period $i+1$. Consider player 2's best response r'' to the belief $[\tau']$ about 1's period $-\,i$ type. Let $\hat{r}''(r) = [r'']$ in period $i+1$. Since $r'' > r'$, $\hat{r}''$ strictly justifies $\bar{a}$ for types τ , τ', so for small $\epsilon \in \mathbb{R}^n$, $\hat{r}''$ strictly justifies $a'' = \bar{a} + \epsilon$ for types τ , τ'. We will 2's actual strategy in period i+1 as a response to a'' is going to be at least $\hat{r}''$.

Suppose $\hat{r}'''$ weakly justifies a'' for type τ. Let $\hat{r}'''(s_2(h_{i-1}))[v_r]$ be V''' and $\hat{r}'(s_2(h_{i-1}))[v_r]$ be V'.

Then $v_a(\tau, a'') + \delta V''' \geq v_a(\tau, \bar{a}) + \delta V'$. Then by single-crossing $v_a(\tau', a'') + \delta V''' > v_a(\tau', \bar{a}) + \delta V'$ since $a'' > \bar{a}$ and $\tau' > \tau$. So $\hat{r}'''$ strictly justifies a'' for type τ'.

So type τ is assigned probability 0 by player 2 after observing a'', by criterion D_ω, since any belief that would weakly justify his taking action a'' would strictly justify type τ'.

Also note that if $\hat{r}$ weakly justifies action a'' for any type $\tau^- < \tau'$ then it must strictly justify action a'' for type τ', because type τ's utility in equilibrium is at least his utility on taking action $\bar{a}$. So any type $\tau^- < \tau'$ is also assigned 0 probability. So player 2's belief is supported on $\{t \in T : t \geq \tau\}$, so strictly justifies type τ', so the action $\bar{a}$ could not have been optimal for type τ'.

Minimal separation

If a current-type $t(i)$'s strategy involved taking an action that did not maximize $u^E(t(i), a, r''(t(i)))$ subject to separating from lower types then he could change his action to the action that does maximize this subject to separating from lower types and still being perceived as at least type $t(i)$, so can raise utility. (Expand)

The result then follows for period i, and by induction for all periods. ■

Proof. Proof of proposition 3 (Convergence of $\sigma_1(i)$)

The proof is by induction on the type number, applying dynamical systems arguments for each type assuming the previous type's action converges.

$\sigma_1(i)(\tau_0) = a^*(\tau_0)$ is constant, so tends to the limit $a^*(\tau_0)$.

Suppose $\sigma_1(i)(\tau_j)$ tends to a limit Σ as $i \to \infty$.

σ is defined by:

1. $\sigma_1(0) = a^*$

2. $\sigma_1(i)(\tau_0) = a^*(\tau_0)$

3. Given $\sigma_1(i+1)(\tau_j)$, $\sigma_1(i)(\tau_j)$, $\sigma_1(i)(\tau_{j+1})$, let h be the solution above $a^*(\tau_j)$ of:

$v_a(\tau_j, \sigma_1(i+1)(\tau_j)) + \delta v_r(r^*([(\tau_j, \sigma_1(i)(\tau_j))])) = v_a(\tau_j, h) + \delta v_r(r^*([(\tau_{j+1}, \sigma_1(i)(\tau_{j+1}))]))$

Then $\sigma_1(i+1)(\tau_{j+1}) = \max(\{h, a^*(\tau_{j+1})\})$.

Write the sequence $\sigma_1(0)(\tau_j), \sigma_1(1)(\tau_j), ...$ as $x_0, x_1, ...$ $x_j \to \Sigma$.

Write the sequence $\sigma_1(0)(\tau_{j+1}), \sigma_1(1)(\tau_{j+1}), ...$ as $y_0, y_1, ...$

Let $A(x) := v_a(\tau_j, x)$ and $B_i(x) := v_a(\tau_j, x_{i+1}) + \delta v_r(r^*([(\tau_j, x_i)])) - \delta v_r(r^*([(\tau_{j+1}, x)]))$. A is defined on $[a^*(\tau_j), a_{\max}]$ and is strictly decreasing and continuous on this set.

Then $y_{i+1} = \max(A^{-1}B_i(y_i), a^*(\tau_{j+1}))$.

Define $B_\infty(x) := v_a(\tau_j, \Sigma) + \delta v_r(r^*([(\tau_j, \Sigma)])) - \delta v_r(r^*([(\tau_{j+1}, x)]))$.

An eventual lower bound on the sequence y_{i+1}.

Consider the function $F_i = A^{-1}B_i$. (So that $y_{i+1} = \max(F_i(y_i), a^*(\tau_{j+1}))$.)

Given α and $i \in \{0, 1, ...\}$, consider the set $S_i(\alpha) = \{x : F_i(x) \geq x + \alpha\} = \{x : B_i(x) \leq A(x + \alpha)\}$.

$S_i(\alpha) = \{x : v_a(\tau_j, x_{i+1}) + \delta v_r(r^*([(\tau_j, x_i)])) \leq \delta v_r(r^*([(\tau_{j+1}, x)])) + v_a(\tau_j, x + \alpha)\}$.

Since $v_r(r^*([(\tau_{j+1}, x)]))$ is concave by assumption and $v_a(\tau_j, x + \alpha)$ is concave in x, $S_i(\alpha)$ is convex, i.e. an interval.[14]

Define $S_\infty(\alpha)$ similarly in terms of B_∞ and we get $S_\infty(\alpha)$ convex too.

($S_\infty(\alpha) = \{x : v_a(\tau_j, \Sigma) + \delta v_r(r^*([(\tau_j, \Sigma)])) \leq \delta v_r(r^*([(\tau_{j+1}, x)])) + v_a(\tau_j, x + \alpha)\}$.)

So we have that for $i \in \{0, 1, ...\} \cup \{\infty\}$, $F_i(x) - x$ is quasi-concave. We can see also that since $v_a(\tau_j, x + \alpha)$ is strictly quasi-concave, $F_i(x) - x$ must be strictly quasi-concave.

Let χ_i be the value of x that maximizes $F_i(x) - x$. It follows from quasi-concavity of $F_i(x) - x$ that for $x \leq y \leq \chi_i$, $F_i(y) - F_i(x) \geq y - x$.

Now consider $S_\infty(0)$. $S_\infty(0)$ contains Σ.

[14] These assumptions have been weakened in the main text and there is one point at which this proof needs minor adjustments, to be added shortly.

Let $l_\infty = inf(S_\infty(0)) \in [-\infty, \Sigma]$ and $h_\infty = sup(S_\infty(0)) \in [\Sigma, \infty)$.

Suppose that $l_\infty < h_\infty$. I will show that in this case y_{i+1} is bounded away from l_∞ eventually.

Define l_i, h_i similarly when $S_i(0) \neq \{\}$.

Take n large enough so that for $m \geq n$:

1a. $S_m(0) \neq \{\}$

1b. $\|F_m - F_\infty\| < \epsilon$. (Uniform metric here as above.)

1c. $|\chi_m - \chi_\infty| < \epsilon$

Where ϵ is chosen such that:

2a. $2\epsilon < \chi_\infty - l_\infty$.

2b. $F_\infty(\chi_\infty - \epsilon) - \chi_\infty > 0$.

By condition 2b, if $x \geq \chi_\infty - \epsilon$, $F_\infty(x) > \chi_\infty$ and so by condition 1b, $F_i(x) > \chi_\infty - \epsilon$. So $\max(F_m(x), a^*(\tau_{j+1})) > \chi_\infty - \epsilon$. This implies that if sequence y_i ever leaves the set $(-\infty, \chi_\infty - \epsilon)$, it never returns.

Suppose the sequence y_i remains inside $(-\infty, \chi_\infty - \epsilon)$ for ever. Otherwise y_i has eventual lower bound $\chi_\infty - \epsilon$.

$y_i > x_i$ is a general property of the iterated Riley solution.

We had earlier $F_m(y_m) - F_m(x_m) \geq y_m - x_m$ for $x_m \leq y_m \leq \chi_m$.

$y_{m+1} - x_{m+1} = \max(F_m(y_m), a^*(\tau_{j+1})) - x_{m+1} \geq F_m(y_m) - F_m(x_m) \geq y_m - x_m$ for $x_m \leq y_m \leq \chi_m$.

The first inequality holds because $\max(F_m(y_m), a^*(\tau_{j+1})) \geq F_m(y_m)$ and $F_m(x_m) \geq x_{m+1}$.

$\chi_\infty - \epsilon < \chi_m$ for $m \geq n$.

So for $m \geq n$, $y_m - x_m \geq \delta := y_n - x_n > 0$

Since $x_m \to \Sigma$, y_m is eventually bounded below by $\Sigma + \delta/2 > l_\infty$.

Conclusion 1 *If $l_\infty < h_\infty, y_m$ is bounded below eventually by a lower bound $\underline{\mathbf{b}}$ strictly above l_∞.*

Now we can show that $y_m \to M = \max(h_\infty, a^*(\tau_{j+1}))$

Case 1 *Suppose $l_\infty < h_\infty$.*

Let $G_i(x) := \max(F_i(x), a^*(\tau_{j+1}))$ for $i \in \{0, 1, ...\} \cup \{\infty\}$.

F_∞ on the set $[\underline{b}, a_{\max}]$ has $F_\infty > x$ below M and $F_\infty < x$ above M and no other fixed points.

M is a global attractor and so the sequence defined by $y_{m+1} = G_m(y_m)$ converges to M since $G_m \to G_\infty$ and y_m remains in $(\underline{b}, a_{\max})$.

Case 2 *Now suppose* $l_\infty = h_\infty$

Then $l_\infty = h_\infty = \Sigma$ since $\Sigma \in S_\infty(0)$.

Let $M = \max(h_\infty, a^*(\tau_{j+1})) = \max(\Sigma, a^*(\tau_{j+1}))$ as before.

On $[\Sigma, a_{\max}]$ G has only one fixed point M, and above this $G(x) < x$.

Take $\epsilon > 0$. Let $\inf_{x \in [M+\epsilon, a_{\max}]}(x - G(x)) = \delta$.

Choose n such that for $m \geq n$, $\|G_m - G_\infty\| < \delta$.

When $y_m \leq M + \epsilon$, $y_{m+1} = G_m(y_m) \leq G_m(M+\epsilon) < G(M+\epsilon) + \delta \leq M + \epsilon$

So if y_m reaches the set $(-\infty, M+\epsilon]$, it stays there.

And above this set, y_m decreases by at least δ each time. So y_m remains in the set $(-\infty, M+\epsilon]$ eventually.

We also know that $y_m > x_m \to \Sigma$, and $y_m \geq a^*(\tau_{j+1})$, so $y_m > \Sigma - \epsilon$ eventually .

And $y_m > a^*(\tau_{j+1})$ always. So $y_m > M - \epsilon$ eventually.

Since ϵ is arbitrary, y_m converges to M.

QED ■

Proof. Proof of Proposition 4

ϵ and η as in the statement.

On $\overline{T}$, $a_S(\delta_1)(\tau) > a^*(\tau)$. Since both a_S and a^* are continuous and $\overline{T}$ compact, we can take ϵ_1 such that $a_S(\delta_1)(\tau) > a^*(\tau')$ for $|\tau' - \tau| < \epsilon_1$.

Take $\epsilon \leq \epsilon_1$.

Claim 3 $\Sigma(\tau) > a^*(\tau)$ *for* $\tau \in T = \{\tau_0, ...\tau_n\}, \tau > \tau_0$.

By corollary (N), for $j > 0$, $\Sigma_1(\tau_j) \geq \overline{a}_S(\delta_1)(\tau_{j-1}, \Sigma_1(\tau_{j-1})) \geq a_S(\delta_1)(\tau_{j-1})$.

So $\Sigma_1(\tau_j) \geq a_S(\delta_1)(\tau_{j-1}) > a^*(\tau_j)$ since $|\tau_j - \tau_{j-1}| < \epsilon \leq \epsilon_1$. This proves the claim.

So we know $\Sigma_1(\tau_0) = a^*(\tau_0)$, and for each j, $\Sigma_1(\tau_{j+1}) = \overline{a}_S(\delta_1)(\tau_j, \Sigma_1(\tau_j))$.

$\overline{a}_S(\delta_1)(\tau_j, \Sigma_1(\tau_j)) = \Sigma_1(\tau_j)$ when $\Sigma_1(\tau_j) \geq a_S(\delta_1)(\tau_j)$. So on $\{\tau_j : a_S(\delta_1)(\tau_j) \leq \overline{a}_S(\delta_1)(\tau_0, a^*(\tau_0))\}$, $\Sigma_1(\tau_j) = \overline{a}_S(\delta_1)(\tau_0, a^*(\tau_0))$. Let the last such τ_j be τ_α.

Take $\eta_2 < \eta/2$ such that for $a_S(\delta_1)(\tau) - \eta_2 \leq x \leq a_S(\delta_1)(\tau)$, $\overline{a}_S(\delta_1)(\tau, x) < a_S(\delta_1)(\tau) + \eta/2$. This is possible because as $\overline{a}_S(\delta_1)(\tau, a_S(\delta_1)(\tau) - \eta_2)$ as a function of τ is continuous, is increasing in η_2 and converges pointwise to $a_S(\delta_1)(\tau)$ as $\eta_2 \to 0$, so converges uniformly to $a_S(\delta_1)(\tau)$ as $\eta_2 \to 0$.

Now take ϵ_2 such that for $|\tau' - \tau| < \epsilon_2$, $|a_S(\delta_1)(\tau') - a_S(\delta_1)(\tau)| < \eta_2$. This is possible because $a_S(\delta_1)$ is continuous and so uniformly continuous on $\overline{T}$.

Suppose that $a_S(\delta_1)(\tau_j) - \eta_2 \leq \Sigma_1(\tau_j) < a_S(\delta_1)(\tau_j) + \eta/2$. Assume that $\epsilon < \epsilon_1, \epsilon_2,$.

If $\Sigma_1(\tau_j) \geq a_S(\delta_1)(\tau_j)$, $\Sigma_1(\tau_{j+1}) = \overline{a}_S(\delta_1)(\tau_j, \Sigma_1(\tau_j)) = \Sigma_1(\tau_j)$, so $a_S(\delta_1)(\tau_{j+1}) - \eta_2 \leq a_S(\delta_1)(\tau_j) \leq \Sigma_1(\tau_j) = \Sigma_1(\tau_{j+1}) < a_S(\delta_1)(\tau_j) + \eta/2 < a_S(\delta_1)(\tau_{j+1}) + \eta/2$.

If $\Sigma_1(\tau_j) < a_S(\delta_1)(\tau_j)$ and $a_S(\delta_1)(\tau_{j+1}) - \delta_2 < a_S(\delta_1)(\tau_j) < \Sigma_1(\tau_{j+1}) = \overline{a}_S(\delta_1)(\tau_j, \Sigma_1(\tau_j)) < a_S(\delta_1)(\tau_j) + \delta/2 < a_S(\delta_1)(\tau_{j+1}) + \eta/2$.

So $a_S(\delta_1)(\tau_{j+1}) - \eta_2 \leq \Sigma_1(\tau_{j+1}) < a_S(\delta_1)(\tau_{j+1}) + \eta/2$.

Let $S'(\tau) = \max(\overline{a}_S(\delta_1)(\tau_0, a^*(\tau_0)), a_S(\delta_1)(\tau))$.

It follows that $|\Sigma_1(\tau_j) - S'(\tau_j)| < \eta/2$ for $j > 0$.

Take $\epsilon_3 \leq \epsilon_1, \epsilon_2$ small enough so that $|\overline{a}_S(\delta_1)(\tau, a^*(\tau)) - \overline{a}_S(\delta_1)(\tau_{\min}, a^*(\tau_{\min}))| < \eta/2$ for $|\tau - \tau_{\min}| < \epsilon_3$. Then $|S' - S| < \delta/2$.

So for $\epsilon < \epsilon_3$, $|\Sigma_1(\tau_j) - S(\tau_j)| < \eta$. QED ■

Proof. Proof of Proposition 5 (sketch)

1. The differential equation is solvable and has a solution that lies strictly above the discounted Stackelberg curve. Let the difference be at least ϵ.

Note that above ϵ above the discounted Stackelberg curve $\frac{\delta}{\delta a} v_S(\delta_1)$ is bounded above away from 0: $\frac{\delta}{\delta a} v_S(\delta_1)(\tau, a) \leq k < 0$ for $a \geq a_S(\delta_1)(\tau) + \epsilon$.

2. Consider the process starting at $\Sigma_1(\tau_0) = \overline{a}_S(\delta_1)(\tau_{\min}, a^*(\tau_{\min}))$ and generated by:

$v_a(\tau_j, \Sigma_1(\tau_{j+1})) + \delta_1 v_r(r^*([(\tau_{j+1}, \Sigma_1(\tau_{j+1}))])) = v_S(\delta_1)\,(\tau_j, \Sigma_1(\tau_j))$.

This is equivalent to:

$[v_S(\tau_j, \Sigma_1(\tau_{j+1})) - v_S(\delta_1)\,(\tau_j, \Sigma_1(\tau_j))] + \delta_1 [v_r(r^*([(\tau_{j+1}, \Sigma_1(\tau_{j+1}))])) - v_r(r^*([(\tau_j, \Sigma_1(\tau_{j+1}))]))] = 0$

3. Using the intermediate value theorem, write this as: $(\Sigma_1(\tau_{j+1}) - \Sigma_1(\tau_j)) \frac{\delta}{\delta a} v_S(\tau_j, \alpha) + \delta_1(\tau_{j+1} - \tau_j) \frac{\delta}{\delta \tau}[v_r(r^*([(\beta, \Sigma_1(\tau_{j+1}))]))]$, where $\alpha \in [\Sigma_1(\tau_j), \Sigma_1(\tau_{j+1})]$ and $\beta \in [\tau_j, \tau_{j+1}]$.

It follows that Σ_1 above τ_0 has uniform Lipschitz constant $K = \frac{\max(\frac{\delta}{\delta \tau}(v_r(r^*([.]))))}{k}$, assuming it remains ϵ above the discounted Stackelberg curve.

4. Rewrite the equation above as:

$$\frac{\Sigma_1(\tau_{j+1}) - \Sigma_1(\tau_j)}{\tau_{j+1} - \tau_j} = \delta_1 \frac{\frac{\delta}{\delta \tau}[v_r(r^*([(\beta, \Sigma_1(\tau_{j+1}))]))]}{\frac{\delta}{\delta a} v_S(\tau_j, \alpha)} \in Conv[\delta_1 \frac{\frac{\delta}{\delta \tau}[v_r(r^*([(\tau, a)]))]}{\frac{\delta}{\delta a} v_S(\tau, \alpha)}]$$

5. Now consider the solution to the differential equation G:

$\frac{G_1(\tau_{j+1})-G_1(\tau_j)}{\tau_{j+1}-\tau_j} \in Conv[\delta_1 \frac{\frac{\delta}{\delta\tau}[v_r(r^*([(\tau,a)]))]}{\frac{\delta}{\delta a}v_S(\tau,\alpha)}]$, where this is taken over $(\tau, a) \in [\tau_j, \tau_{j+1}] \times [\Sigma_1(\tau_j), \Sigma_1(\tau_{j+1})]$.

6. The range of this convex hull tends to zero uniformly as the distance between types tends to zero. Take any $\zeta > 0$, then choose $\eta_2 > 0$ such that when all types are within η_2 of each other, the convex hull above has range at most $\zeta < \frac{\epsilon}{2(\tau_{\max}-\tau_{\min})}$.

Therefore $(\Sigma_1(\tau_j) - \Sigma_1(\tau_0)) - (G(\tau_{j+1}) - G(\tau_0)) \leq 2\zeta(\tau_j - \tau_0)$, as long as Σ_1 remains above ϵ above the discounted Stackelberg curve.

7.Combining 3 and 6, the solution remains above ϵ above the discounted Stackelberg curve.

8. $\Sigma_1(\tau_1)$ is arbitrarily close to $\overline{a}_S(\delta_1)(\tau_{\min}, a^*(\tau_{\min}))$ for η_2 small enough, which is arbitrarily close to $G_1(\tau_1)$.

The result follows. ■

Chapter 2

A Repeated Signalling Model of Reciprocity

2.1 Introduction

Modelling reciprocation

There exist a variety of attempts to explain reciprocal behaviour and cooperation, including standard repeated game theory, psychological games, and behavioural models. The model of reciprocation given here involves several views on the nature of reciprocal interaction. First that in evaluating the generosity of actions, people will look for the motives behind the action. It is not a model where actions inspire actions directly. Second that more generous actions result from more generous dispositions. This is more direct reasoning than a repeated games approach to repeated interactions, where any ostensibly generous actions are the result of potentially complex strategies and reasoning about future consequences. Third it takes a rational rather than behavioural approach, supporting a simple idea about the workings of reciprocation in an expected utility model.

In the model more generous actions will encourage more generous responses, and this is what supports more generosity than would exist otherwise, a process that requires a dynamic model rather than a static one such Rabin's [24]. Types represent intrinsic generosity and players show generosity to each other in order to gain more favourable treatment in the future. Players factor out this effect and make correct deductions about types.

The effective generosity of each player is composed of his intrinsic generosity and his deductions about the other's generosity in the past. Types across time are assumed to be independent, which represents people reacting to being treated surprisingly well or badly: type is effectively shown as how high up a player is in the probability distribution over actions, where higher actions are more generous. (Another motivation for this could be that a player is interested in whether the other was more or less generous than could be expected given what has happened in the game.) At any rate there is volatility in the model which makes it appropriate for some observed situations and not others.

The model is a signalling model of reciprocation, with both parties signalling intrinsic generosity to each other by taking kinder (higher) actions. A solution concept is proposed for the game (dynamic Riley) which is well-defined and calculable, and a refinement proposed that selects it. The signalling incentives are always present and "cooperation" will be observed with higher actions would exist without signalling concerns.

Modelling repeated signalling games

The paper adapts Roddie [26] to a situation with two signallers, with additive separability in the utilities of both players. The important property that allows the tractable solution of repeated signalling games is that single crossing and monotonicity can be preserved through time. That allows us to define select a simple dynamic version of the Riley equilibrium, just as single crossing and monotonicity allow the static Riley equilibrium to be defined and selected. It turns out that correlation between types gives potential problems with monotonicity: by signalling a higher type today, player A may reduce the signalling incentive for player B tomorrow (which is bad for A) because B may think that with a higher type there is less need to signal to A. While conditions may be found that allow monotonicity to be preserved, this paper takes the approach of assuming independent types.

Several alternative approaches to repeated signalling are explored in this paper. Types are either finite or continuous, and the refinement is designed to apply to either. Continuous types has the potential to generate cleaner solutions. The game is not required to be finite, and Markov equilibria of the infinite game are explored in addition to solutions of the finite game.

2.2 Framework

Two players, P1 and P2, play in a repeated game with k periods, where $k \in \mathbb{N} \cup \{\infty\}$. If $k < \infty$, label the periods $-k+1$ to 0. Players take actions a_i^{P1}, a_i^{P2} in period i simultaneously at each stage. Actions a_i^p, $p = P1, P2$ lie in the set $A^p = [a^p_{\min}, a^p_{\max}]$. Higher actions are more favourable to the other player.

Types t_i^p are drawn independently from distributions μ_{T^p} on T^p with full support, real intervals or finite sets of reals. In the continuous case the distributions will (without loss of generality) uniform over $T^{P1} = T^{P2} = [0, 1]$. Types exist from the period before the initial play period, if an initial period exists.

Players have discount expected utility with discount factors δ_p. Payoffs in the stage game are $U_p(t_i^p, t_{i-1}^{-p}, a_i^{P1}, a_i^{P2}) = v_p(t_i^p, t_{i-1}^{-p}, a_i^p) + w_p(a_i^{-p})$, where "$-$" permutes P1 and P2.

Assumption 1 *v_p and w_p are continuously differentiable*

Assumption 2 *$\int v_p(t_i^p, ., a_i^p) d\mu$ for any measure μ on T^{-p} and any t_i^p is strictly quasi-concave in a_i^p.*

Assumption 3 *w_p is strictly increasing*

Assumption 4 *Single crossing: if $a'' > a'$ and $((t^p)'', (t^{-p})'') > ((t^p)', (t^{-p})')$ then , $v_p((t^p)'', (t^{-p})'', a'') - v_p((t^p)'', (t^{-p})'', a') > v_p((t^p)', (t^{-p})', a'') - v_p((t^p)', (t^{-p})', a')$*

Assumption 5 *undesirable $a_{\max}^p : v_p(t_i^p, t_{i-1}^{-p}, a_{\max}^p) + w_p(a_{\max}^{-p}) \leq v_p(t_i^p, t_{i-1}^{-p}, a^p) + w_p(a_{\min}^{-p})$ for some a^p, for all t_i^p, t_{i-1}^{-p}*

Perfect Bayesian Nash equilibrium

Actions are observable, while only current and past types of self are observed. So strategies and beliefs can condition on previous actions of both players and current and past types of self. Strategies map into mixtures over actions, and satisfy a measurability requirement that the probability of taking actions in any measurable set is a measurable function of past actions and past and current types, as in [26]; this is needed to define a probabilistic outcome of the game from any point in the game.

Beliefs must always conform to Bayes' rule from any point in the game (off or on the equilibrium path) and always ascribe the same independent probability distributions μ_{T^p} to future types.

2.3 Riley equilibria

Static Riley equilibrium

Take the static signalling game where the signaller receives utility $u(t, a, \widehat{t})$, where $t \in T$ is actual type, $a \in A$ is message sent, and $\widehat{t}$ is type inferred by the respondent. Suppose that u satisfies single crossing, monotonicity, strict quasi-concavity in a, and that the highest action is not optimal even if beliefs are maximal at the highest action and minimal elsewhere. Then the associated Riley equilibrium exists.

Suppose T is finite, $T = \{t_0, ...t_n\}$. Then the Riley equilibrium is given by:

$\Re(u) : T \to A$, $\Re(u)(t_i) = \arg\max_a u(t_i, a, [t_i])$ over a s.t. $u(t_j, a, [t_i]) \leq u(t_j, \Re(u)(t_j), [t_j])$ for $0 \leq j = i - 1$

If T is a continuous interval $[t_0, t_h]$, then the Riley equilibrium is defined by the differential equation:

$\Re(u)(t_0) = \arg\max_a u(t_0, a, [t_0])$

$\Re(u)'(t).u_2(t, \Re(u)(t), [t]) + u_3(t, \Re(u)(t), [t]) = 0$, where u_2 and u_3 are partial derivatives w.r.t the 2^{nd} and 3^{rd} arguments of u.

In our context we will define $\Re^p$ to take $T = T^p$, $A = A^p$.

Note that while the monotonicity of u for $\widehat{t} \in \Delta(T)$ are important, in calculating $\Re(u)$ only degenerate distributions of $\widehat{t}$ are used.

Dynamic Riley equilibrium

In the dynamic version of the Riley equilibrium, both players take strategies that are functions of current own type, the preceding (believed) type of the other and the number of periods left in the game. So strategies are given by $\sigma_i^p(t_i^p, \widehat{t_{i-1}^{-p}})$, where i is the the period number, t_i^p is current own type, and $\widehat{t_{i-1}^{-p}}$ is the belief about the other's type in the previous period.

Suppose that from periods $n+1$ on, (in every history) strategies conform to σ_i^p and that $\sigma_i^p(t_i^p, \widehat{t_{i-1}^{-p}})$ is strictly increasing in $\widehat{t_{i-1}^{-p}}$. Suppose that from periods $n+1$ on (in every history), type is revealed after moves are observed.

Then in period n, actions of player p and subsequent beliefs about his type have an effect only on actions in the next period: in periods $n+2$ on, all actions are based on types from periods $t+1$ on (since by the assumption types are revealed after moves are observed, so type beliefs from periods $n+2$ on about types from periods $n+1$ on depend on actual types from periods $n+1$ on. And in the next period $n+1$, player p's action is a function of his type in that period and p's beliefs about player $-p$'s type in period n, which is independent of the period-n action of player p and subsequent beliefs about player p's type. So in period n, actions of player p and subsequent beliefs about his type affect only the next period action of player $-p$.

Suppose in period n player p, with type t_n^p, believes player $-p$'s type last period was $\widehat{t_{n-1}^{-p}}$ and expects $-p$ to take action $\widehat{a_n^p}$. If player p takes action a_n^p and is believed to be type $\widehat{t_n^p}$, then utility is $\int v_p(t_n^p, t_{n-1}^{-p}, a_n^p) d\widehat{t_{n-1}^{-p}}(t_{n-1}^{-p}) + \delta_p \int w_p(\sigma_{n+1}^{-p}(t_{n+1}^{-p}, \widehat{t_n^p})) d\mu_{T^p}(t_{n+1}^{-p}) + c$, where c is independent of a_n^p and $\widehat{t_n^p}$. This satisfies the requirements of single crossing (first term), strict quasiconcavity in action (first term), monotonicity (second term increases in $\widehat{t_n^p}$) and undesirable $a_{\max}^p$, so we can apply the Riley function $\Re^p$ to define:

$$\sigma_n^p(t_n^p, \widehat{t_{n-1}^{-p}}) = \Re^p(u)(t_n^p), \text{ where:}$$
$$u(t, a, \widehat{t}) := \int v_p(t, t_{n-1}^{-p}, a) d\widehat{t_{n-1}^{-p}}(t_{n-1}^{-p}) + \delta_p \int w_p(\sigma_{n+1}^{-p}(t_{n+1}^{-p}, \widehat{t})) d\mu_{T^p}(t_{n+1}^{-p})$$

This defines σ_n^p in terms of σ_{n+1}^{-p}. Note that $\sigma_n^p(t_n^p, \widehat{t_{n-1}^{-p}})$ will also be strictly increasing in $\widehat{t_{n-1}^{-p}}$. This is because increasing $\widehat{t_{n-1}^{-p}}$ increases the preference of player p for higher actions, a single crossing condition on u w.r.t. $\widehat{t_{n-1}^{-p}}$, holding type t fixed (deriving from the single crossing assumption on v_p). This generates a strictly higher actions of each type[1] in the Riley equilibrium $\Re^p(u)$. This is a simple exercise to show. In the discrete case, by induction on the types, with preferences of each type being strictly more inclined to higher actions the separation condition being always (strictly) stronger. In the continuous case, if the action of any type for a higher $\widehat{t_{n-1}^{-p}}$ ever equals the action for a lower $\widehat{t_{n-1}^{-p}}$, the gradient of $\Re^p(u)$ at that point must be higher.

This allows us to extend the dynamic Riley equilibrium backwards by one period indefinitely. A dynamic Riley equilibrium is defined to be an equilibrium in which the relationships between periods above hold at every stage, with full revelation at every stage and strategies given by $\sigma_i^p(t_i^p, \widehat{t_{i-1}^{-p}})$, strictly increasing in $\widehat{t_{i-1}^{-p}}$. The behaviour of p in period i will depend on the behaviour of $-p$ in period $i+1$, which depends on the behaviour of p in period $i+2$, and so on.

Assuming that actions are optimal in the final period (if it exists), a Dynamic Riley equilibrium must be a perfect Bayesian Nash equilibrium. Beliefs are updated correctly: at every stage each current type takes a different action (since $\Re$ is separating) and so the true type is revealed (held with probability 1) after one of these actions, as assumed. Given beliefs, each action is optimal because given future strategies, current incentives are the static signalling incentives given above and $\Re^p$ gives an equilibrium of the static signalling game so optimal actions when applied here.

Note that apart from the initial period if it exists, current beliefs of p about $-p$'s previous type have support on a single point, so the structure of the game is given by the simple object $\sigma_n^p({}_n^p, [t_{n-1}^{-p}])$, specified only for degenerate beliefs $[t_{n-1}^{-p}]$. Then the relationships simplify slightly to:

$$\sigma_n^p(t_n^p, [t_{n-1}^{-p}]) = \Re^p(u)(t_n^p), \text{ where:}$$
$$u(t, a, \widehat{t}) := v_p(t, t_{n-1}^{-p}, a) + \delta_p \int w_p(\sigma_{n+1}^{-p}(t_{n+1}^{-p}, \widehat{t})) d\mu_{T^p}(t_{n+1}^{-p})$$

Finite game: Iterated

In a finite game dynamic Riley equilibrium behaviour in the last period last period (0) is given by static optimization of $v_p(t_0^p, t_{-1}^{-p}, a_0^p)$, required by Perfect Bayesian Nash equilibrium, giving $\sigma_0^p(t_0^p, \widehat{t_{-1}^{-p}})$ as the strategy in the final period, strictly increasing in $\widehat{t_{-1}^{-p}}$. Previous periods are defined iteratively as above.

[1]Except possibly the lowest. Assume there are at least two types.

Analytic limit properties have not yet been found for the iterated Riley equilibrium. Non-persistence of types requires a different method from Roddie [26]. Work is in progress to explore either dynamical systems arguments or contraction mapping methods that might yield limit properties.

Nevertheless the model is calculable for a finite game, and a computer program that does so for the finite type case can be simply a list of the above definitions.

Infinite game: Markov

In an infinite game (infinite in the future direction, infinite or finite in the past), a Markov Riley equilibrium is a dynamic Riley equilibrium in which σ_n^p is independent of n. Write $\sigma_n^p = \sigma^p$. This additional assumption may not in fact be necessary and it may be the case that any dynamic Riley equilibrium with an infinite horizon is a automatically a Markov Riley equilibrium, but this has not been shown as yet. The question is connected with limit arguments in the finite game. Take the closed set of weakly increasing functions $(\sigma_0^{P1}, \sigma_0^{P2})$. These generate an associated functions $(\sigma_{-1}^{P1}, \sigma_{-1}^{P2})$ by the process described above, a continuous function generating a smaller closed set. Applying this process repeatedly, if there exists a unique limit that is independent of play in the last period, this set must in the limit shrink to a point (note that it is important that these sets are closed to make this conclusion). This point itself will then be a Markov Riley equilibrium, and any other point if it is a dynamic Riley equilibrium must be outside this set after a finite number of iterations and so cannot be supported in an infinite horizon.

Strategies given by σ, with $\sigma^p(t_i^p, \widehat{t_{i-1}^{-p}})$ strictly increasing in $\widehat{t_{i-1}^{-p}}$, are a Markov Riley equilibrium if the following holds for each player p, giving two symmetric relations between P1 and P2:

$$\sigma^p(t_0^p, [t_{-1}^{-p}]) = \Re^p(u)(t_0^p), \text{ where:}$$
$$u(t, a, \widehat{t}) := v_p(t, t_{-1}^{-p}, a) + \delta_p \int w_p(\sigma^{-p}(t_1^{-p}, \widehat{t}))d\mu_{T^p}(t_1^{-p})$$

Claim 1 *There exists a Markov Riley equilibrium*

Proof. First consider the finite type case. Consider the function F taking $(\sigma_{n+1}^p(., [.]), \sigma_{n+1}^{-p}(., [.]))$ to $(\sigma_n^p(., [.]), \sigma_n^{-p}(., [.]))$ as described above. Defined above on on functions σ_{n+1}^p where $\sigma_n^p(t_{n+1}^p, [t_n^{-p}])$ is strictly increasing in $[t_n^{-p}]$, F can be extended to functions that are weakly increasing in $[t_n^{-p}]$ which are a closed set by extending $\Re(u)$ to $u(t, a, \widehat{t})$ only weakly increasing in $\widehat{t}$[2]. This space of functions is compact because types are finite and action spaces are intervals. And F is continuous. So there exists a fixed point by Brouwer's fixed point theorem, which is a Markov Riley equilibrium.

[2] The function can be defined exactly in the original form. There may not be separation whenever u is not strictly increasing in $\widehat{t}$, so it would not be called a Riley equilibrium.

Now consider a continuum and approximate both type spaces by finite types, taking probability measures on the finite types that converge (weak topology) to the original uniform distribution on the interval. The important fact is that the Riley eqilibrium probability distribution - the distribution generated by $\Re(u)$, converges to the distribution generated in the continuous case when the the type space converges as above and u converges (see immediately below). So the finitely defined Riley equilibrium converges to the Riley equilibrium defined by the differential equation.

Convergence of u is in understood like this: $u(t, a, \widehat{t})$ is defined for t in the continuum, a in A, and $\widehat{t}$ in the finite approximation. As a function of $\widehat{t}$ it is always weakly increasing. So consider it as a function on the continuum, takin value at any point given by the sup of values up to that point. Then we have u in the same space and define convergence as uniform convergence. So by taking limits of the Markov Riley equilibria in the finite case (by sequential compactness of the space of weakly increasing strategies σ_n^p), with corresponding limits of u defined above, so that the u given by the finite case tends to the u of the limit strategy σ^{-p}. Then by continuity of the Riley map we have that the limit is a Markov Riley equilibrium. ■

Consider the continuous case with types drawn from a uniform distribution over $[0, 1]$.

Then the differential equation defining $\Re$ gives:

$$(v_p)_a(t_0^p, t_{-1}^{-p}, \sigma^p(t_0^p, [t_{-1}^{-p}]))\frac{\delta\sigma^p}{\delta t_0^p}(t_0^p, [t_{-1}^{-p}]) + \delta_p\frac{\delta}{\delta t_0^p}[\int_0^1 w_p(\sigma^{-p}(t_1^{-p}, t_0^p))dt_1^{-p}] = 0$$

where $(v_p)_a$ is the partial derivative of v_p w.r.t. action.

We also have initial conditions $\sigma^p(0, [t_{-1}^{-p}])$ given by taking the myopic optimum of the lowest type 0.

In the symmetric case where v_p and w_p are independent of p we can look for a symmetric solution $\sigma^{P1} = \sigma^{P2} = \sigma$ and look to solve a single equation in $\sigma(x, y)$ with boundary values $\sigma(0, y)$ given.

$$v_a(x, y, \sigma(x, y))\frac{\delta\sigma}{\delta x}(x, y) + \delta_p\frac{\delta}{\delta x}[\int_0^1 w(\sigma(\overline{y}, x))d\overline{y}] = 0$$

2.4 Refinement and equilibrium selection

Refinement

The refinement D_ω in [26] deals with the case of a single signaller and finite types. The refinement selects beliefs by ascribing infinitely smaller probability to greater mistakes, where these mistakes are measured by the beliefs needed to justify suboptimal actions.

Beliefs are over the next-period response of the respondent and this refinement is difficult to apply to a situation with two signallers since altering (in the mind of the signaller who makes a mistake) the next-period play for either player should now change beliefs in the following period, since the player with altered actions may not now be separating. Instead the refinement here makes a mistake about the type inferred from his action, and supposes when he takes the suboptimal action that there will be a equilibrium from that point with the wrong belief as an initial condition, where this equilibrium satisfies the refinement.

Like D_ω, additive separability between actions of both players in the utility functions of both signallers make the refinement equivalent to a cardinal utility loss criterion (a type who loses more utility from a mistake is infinitely less likely to have made it). While the utility loss criterion has the important advantage of simplicity, belief-based refinements are potentially more generalizable into situations without additive separability. The proofs given will work for either refinement.

Suppose the game is finite. Take a perfect Bayesian Nash equilibrium. Call the vector of present and past types of p the type vector t^p. Let equilibrium utility for type t^p be $u(t^p)$. At any history of actions, suppose that p makes an action a_i, which may or may not be in his equilibrium strategy, Alter the equilibrium in the future so that player $-p$'s beliefs about player p's i-type after a_i is observed are $\widetilde{t}_i^p \in \Delta(T^p)$, and so that there is a perfect Bayesian Nash from that point satisfying the refinement. Let the supremum over all such equilibria give utility $u(t^p, \widetilde{t}_i^p)$. The justifying beliefs for a type t^p to take action a_i are the set of $\widetilde{t}_i^p$ for which $u(t^p, \widetilde{t}_i^p)$ is greater than $u(t^p)$.

The refinement is that if there are two disjoint intervals A and B containing types, and the (non-empty) intersection of the weakly justifying beliefs of types in A contains the union of the strictly justifying beliefs of types in B, then types in B are ruled out after observing action a_i.

After action a_i, player p expected to come to his mind and use his equilibrium strategy and beliefs, forming correct beliefs about the type he is thought to be. So player $-p$ best-responds to the original strategy of p given beliefs given by the refinement.

In the infinite horizon case, instead of assuming that there is a Bayesian Nash equilibrium satisfying the refinement after coming to a wrong belief, it is assumed that players use the same strategies $\sigma^p(t_i^p, t_{i-1}^{-p})$.

Equilibrium selection

Proposition 1 *Each refinement, belief or utility-based, selects the iterated/Markov Riley equilibrium uniquely in the finite case out of perfect Bayesian equilibria, in the infinite horizon case perfect Bayesian Markov equilibria with strategies* $\sigma_n^p(t_n^p, [t_{n-1}^{-p}])$.

Proof. Appendix ■

2.5 Discussion

A model of reciprocation is presented in which players are more generous in order to signal intrinsic good-will to the other to encourage a response. A method is given for solving the game with finite or continuous types and with a finite or infinite horizon, with Markov strategies explored in the latter case. There is minimal separation of types of each player at each stage, and this is selected by the refinement, adapted from [26]. While the iterated Riley equilibrium of the finite game is tractable further work needs to be done to show limit properties and to show uniqueness of the Markov Riley equilibrium. The connections between the questions have been clear and they can now be addressed from more than one angle.

If the model could allow correlation between types over time it could potentially give it greater scope and then with the relaxing of additive separability allow generalizing to various important situations of joint reputation. It may be that assumptions can be found under which monotonicity is preserved through the game even with type correlation.

2.6 Appendix

Proof. Equilibrium selection:

i. **Uniqueness**

Consider the belief-based refinement. At any history take an action a_i^p of player p and consider adjusting beliefs as in the refinement. In the finite case, by induction assume that strategies conform to the iterated Riley equilibrium in games of length one less than the time remaining, and so in future periods as well as in the mistake-belief of the refinement. In the infinite case we have Markov strategies. Either way when beliefs are adjusted about player p's type after taking action a_i^p only the opponent's action in the next period changes, which has an additive effect on utility, and beliefs can be changed continuously to give a range of utility changes sufficient to select the dynamic Riley equilibrium. This is why the utility criterion is equivalent.

Suppose there is pooling at any stage. Then this must occur with an interval of period-types with supremum s, because by single crossing types' actions must be weakly increasing. (Given future play described by strategies of the form σ_i^p) Consider a slight increase in action. And the belief mistakes/utility bonuses that would justify it. This is a strictly decreasing below s but never an empty set below s. So the refinement implies that every closed interval of types below s is ruled out. So taking an action slightly above the pool discontinuously improves beliefs and so utility. That implies separation.

For the finite type case, to get the Riley equilibrium at each stage we also need to show minimal separation. If a type τ_j is not maximizing utility subject to separating from the preceding type, consider taking the maximizing action instead. This action is strictly justified by belief τ_j, while the beliefs that strictly justify the action for the preceding type do not include τ_j. So the set of beliefs that weakly justifies the action for type τ_j includes the set of beliefs that strictly justify the action for type τ_{j-1}, so by moving to this action beliefs are at least preserved, so the original action was not optimal.

Since we have minimal separation and signalling incentives (assuming the future is as described) are as discussed in the section 2.3, play in period i in terms of play in period $i+1$ is given by the Riley map, as described in that section.

ii. **Existence**

That the Iterated/Markov Riley equilibrium is a perfect Bayesian Nash equilibrium has already been explained in section 2.3. The proof that it satisfies the refinement is a minor variation on [6] and [26] and is omitted. ■

Chapter 3

Mobility and redistribution with non-linear taxes

3.1 Synopsis

Mobility between regions imposes a constraint on redistribute. The model in this paper shows that if taxes are very general, this constraint is very strong. It is common in simple models of taxation for taxes to be linear. Then they can be studied under democratic policy choice, with both taxes and ability aligning along some line and a single-crossing condition generating a median voter result. In this way linear taxes give a tractable framework for political economic models. Epple and Romer [17] study redistribution and mobility with linear taxes and conclude that some redistribution is possible, but limited by stratification and the benefit of attracting richer populations.

The model here weakens the restriction to linear taxes, allowing very general taxes, non-linear taxes which can discriminate between residents and immigrants. It also weakens policy choice assumptions: there is no assumption of democratic choice, only local Pareto optimality. Pareto optimality, commonly used (for better or worse) as a normative principle when applied to policy choice, is used here as a relatively weak positive assumption. With general taxes there is more competition than with a linear tax, since it is possible to target particular classes of people, and local Pareto optimality is enough to obtain this effect.

In addition to studying mobility the model incorporates a choice of work by workers and so can study the moral hazard problems of optimal taxation and mobility together. Attention is given to the question of mobility but the results do show how to separate the two issues. Mobility implies that regions will not have workers who would be better off doing an efficient amount of work and being fully compensated. Then the moral hazard problem can be studied with this as an additional constraint.

The result that regions will not have workers who would be better off doing an efficient amount of work and being fully compensated should not come as a great surprise.

However it needs to be shown coherently in a model, and modelling the effect poses technical difficulties, largely related to multiplicity and general intractability of equilibrium after regions have made tax schedule decisions. The problem of multiplicity of equilibria is dealt with here by considering only changes in populations in which one region poaches residents from one or more other regions. If no "poaching" is possible, the equilibrium is considered stable.

3.2 Model

Several independent regions, labelled 1 to n, control their own tax policy. There is mobility between regions and no region can prevent its residents from leaving, although regions may restrict immigrants.

People are described by a characteristic a lying in a compact interval A, taken to represent ability. There is a measure μ giving the distribution of characteristics in the whole population, with $\mu(A)$ being the total population. Population is taken as continuous so non-integer values of μ are meaningful.

Each person regardless of which region he resides in choose a level of productivity w from a compact interval W. A resident of region i of type a who does w work earns gross income $I(w)$, where I is increasing. Total income in region i is $F_i(\omega_i) = \int I(w)d\omega_i(w)$ where the measure ω_i describes work done by the population in region i. Preferences $\preccurlyeq_a$ of type a are over w and net income x and derive from a continuous utility function $U(a, w, x)$.

There are two stages. Population is initially distributed among regions according to μ_i, giving a measure over A for each region i. First governments simultaneously set tax schedules. They may have the ability to set quotas at the same time. Then people decide whether to stay in their region or migrate, and at the same time how much to work.

Regional governments set tax schemes $T_i : [n] \times W \to \mathbb{R}$[1] with $-T_i$ upper semicontinuous. $T_i(j, w)$ is the tax paid by a person originally from region j who does work w. Taxes are bounded. Taxes are allowed to discriminate between immigrants and residents. Total government income is $\int T_i d\nu_i$, where ν_i is the distribution of work and region of origin in region i. Additionally governments posses independent resources $r_i > 0$. The role of the assumption is to rule out regions with zero population. When population is very low, a region becomes very attractive. Taxes are redistributed equally as handouts *to residents who remain in the region*, so the net income of resident i who remains in region i and works w is $I_i(w) - T_i(i, w) + \frac{1}{\nu_i(\{i\} \times A)}(\int T_i d\nu_i + r_i)$.

The precise way that taxes are determined - whether by bargaining, democratic vote with full or partial suffrage, or some other type of political process - is not modelled, but it is assumed that the choices are locally Pareto optimal for the region. That is, given expectations about other regions' choices, each region's choice is locally Pareto optimal,

[1] Taxes are functions of work rather than income, but if I is one to one these are equivalent.

in the sense that no other policy would be expected to strictly improve some residents' utilities and weakly improve all residents' utilities.

Quotas for immigrants, if allowed, are described by $q_i(j)$, measures on W, with $q_i(j)(S)$, $S \subseteq W$ being the quota in region i for immigrants doing work in S from region i. A quota is broken by measure x if there exists $S \subseteq W$ such that $x(S) > q_i(j)(S)$. A quota is filled for work $w \in W$ when immigration from region is given by measure x if for any $\epsilon > 0$, $x + 1_w\epsilon$ breaks the quota.

Assumption 1 *Utility is quasi-linear.*

Assumption 2 *Single crossing: If* $w_1 < w_2$ *and* $a_1 < a_2$ *and* $(w_1, x_1) \preccurlyeq_{a_1} (w_2, x_2)$ *then* $(w_1, x_1) \preccurlyeq_{a_2} (w_2, x_2)$

Assumption 3 *There are no point masses of ability*

The assumption is used for in the following way:

Lemma 1 *For any tax schedule in any region and for any population, there is a unique work-choice equilibrium level of handouts, and handouts increase continuously with resources* r_i.

Proof. By quasi-linearity, preferences over work are independent of handouts. By assumption 2, the correspondence between ability and work is strictly increasing. So the abilities for which the correspondence is many-valued has measure zero. (Since an increasing function must have measure zero points of discontinuity.) ■

An efficiency equilibrium is a tax schedule choice by each region, which given that future play is optimal, is locally efficient for each region. The results presented here have to do with equilibria when population is stable - there is no movement. What generates limits on the power to redistribute is the ability of regions to offer better deals to residents of other regions who are having income redistributed away from them. We will call these residents "discontent" (for want of a better word).

3.3 The movement phase

When tax schemes have been set people decide what region to be in and how much to work. This phase can be described by a measure M over $[n] \times [n] \times A \times W$ with (i, j, a, w) representing a person with ability a from region i moving to region j and doing work w. In the absence of quotas the optimality condition is that there is not a positive measure of people who would gain utility by changing their choice of region and/or work. With quotas the optimality condition is that there is not a positive measure of people (i, j, a, w) who would gain by changing to some choice (j', w') of region and work, where the quota for w' in region j' is unfilled.

Claim 1 *Given any choice of tax schedules and quotas, there exists an subgame equilibrium in movement and work choice.*

Proof. First approximate the ability set A by a finite set A^F and the initial distribution of populations by μ_i^F, with support on A, and work by a finite set W^F, approximating taxes by the lowest set between the previous and the next point in W^F.

Let D^F be the set of feasible distributions over $[n] \times [n] \times A^F \times W^F$ (with correct total population and satisfying quotas). D is a convex and compact set embedded in $\mathbb{R}^k$ for some k.

Given an outcome $d \in D^F$ there are hand-outs associated with any region of origin, destination and work, and an optimal choice of destination and work exists for any person with any ability for any region of origin.

So given $d^F \in D^F$ there is a non-empty set $f(d) \subseteq D^F$ giving optimal movement and work choice given income functions above. The function f is upper semicontinuous by continuity of hand-outs with population, taking hand-outs to be infinite when the population is zero.

So by Kakutani's fixed point theorem there exists a fixed point f^F. This is an equilibrium of work and movement for the finite approximation μ_i^F.

Now let μ_i^F tend to μ_i (weak topology), taking finer and finer sets A^F, and W^F tending to W. Take a limit of fixed points by sequential compactness of distributions over $[n] \times [n] \times A \times W$. This limit is an equilibrium of movement and work choice.

Infinite handouts when population is zero rules out equilibria with zero population (assuming that all regions had positive population to start with). ■

3.4 Definitions

Definition 1 *A person of ability a is **content** if his utility is at least* $\max_w U(a, w, I(w))$.

The notion of contributing (giving more in tax then is received in benefits) will be useful to the arguments presented later. The difference between contributing and being discontent arises because of the moral hazard problem in each region, with workers potentially not doing an efficient amount of work.

Definition 2 *A person doing work w is **contributing** if his net income is less than $I(w)$.*

Lemma 2 *If a person is contributing, he must be discontent.*

The exploited residents may be poached by other regions:

Definition 3 *Take a tax schedule decisions T_j of all regions, and optimal population movement and work decisions M that involve no movement. A* ***poaching*** *by region i from a set S with $i \notin S$ is an adjustment $T_i \to T_i'$ by region i and corresponding optimal population movement and work decisions $M\prime$ with the only movement being from S to i, and the existing population in region i having Pareto-improved utility.*

The results presented here consider "stable" equilibria, in which poaching is not possible.

Definition 4 *A* ***stable*** *efficiency equilibrium is an equilibrium with no population movement in which no region can poach from other regions.*

Note that this is a stronger requirement than just requiring a subgame-perfect efficiency equilibrium, since when region i changes its tax schedule from T_i to T_i', after T_i' there may be multiple equilibria of population work and movement choice. The idea of poaching is that there is no tax scheme T_i' which *could* be a successful poaching.

3.5 A weaker result

Proposition 1 *In a stable efficiency equilibrium with quotas, there exist at least two regions in which all residents are content.*

Proof. Suppose not. Then take a region i such that all other regions have some discontent residents. Take one such resident of of region $j \neq i$, of ability a.

Let $w_e = \arg\max_w U(a, w, I(w))$. Set a new tax rate $T_i'(j, w) = -K$ for $w \neq w_e$, and $T_i'(j, w_e) = \epsilon$, where $\epsilon > 0$ but $U(a, w_e, I(w_e) - \epsilon)$ is strictly greater than the original utility of a in region j. This is possible because a is exploited. K is chosen to be large enough to discourage all abilities from doing work other than w_e.

Now there is a set S of residents in region j that strictly prefers to move to region i and do work w_e than remain in region j and do the current work, gaining utility more than $\alpha > 0$. α is chosen such that this set still includes a. Since U is continuous in a, the utility of residents of j is continuous in a, and the utility of moving to region i and doing w_e is also continuous in a, so an open set around a is in S, so S has positive measure.

Now take an arbitrary parametrization S_x, with $S_x \subseteq S$ and $\nu_j(\{j\} \times S_x) = x$, and S_x increasing, defined on some interval $x \in [0, \delta]$. So S_x is an increasing set within S with measure x in region j. Take δ to make sure that S_x has measure less than S.

Consider the revenue available for distribution when S_x is removed. The change in revenue is a combination of stopping benefits to S_x (a continuous change) and removing tax revenue from S_x (again a continuous change, Lipschitz since taxes are bounded).

This is continuous, with $max_w I(w)$ so is lower semi-continuous from above. (It will be continuous but that is not necessary to show.) So the level of hand-outs to residents when S_x is removed are continuous in x.

1. Suppose that for some region j and for all x hand-outs are weakly increased. This is true if non-contributing residents are being poached. Then take x small enough that the increase in hand-outs is small enough that no-one's utility in region j is increased by more than α by the increase in hand-outs. Then poach from the single region j, taking S_x, setting taxes as above and setting a quota x for work w_e and 0 for all other work. Revenue available for existing residents of i is increased by ϵx, giving by lemma 1 a Pareto improvement for region i. Since members of S_x gained more than α utility from the original utility, and after the poaching utility from remaining is increased by less than a, it is still optimal for them to move to region j. And since hand-outs in region j have increased, people remaining there still have no incentive to move to a third region. All other incentives remain the same, so the poaching is an equilibrium of movement and work choice.

2. Suppose instead that for each region j and for $x > 0$, handouts are reduced when S_x is removed. This happens if contributing residents are being poached. Then for some $\varepsilon > 0$, values of x can be chosen for each other region such that hand-outs are reduced by ε, since handouts are lower semi-continuous from above. Poach from all other regions, taking for the various values of x corresponding residents S_x, setting quotas to the values of x for work w_e (which is dependent on region) and zero elsewhere. Revenue from immigrants generates a Pareto improvement for existing residents (lemma 1). ■

The method is very roughly: if there are discontent non-contributors in a single region, then poach them. If there are not, then poach from all regions in such a way as to reduce benefits equally in all other regions. This is important because if benefits were reduced more in region B than region C some residents of region B might then have an incentive to move to region C. Not only does this upset the definition of poaching but it could upset the equilibrium altogether that is Pareto improving for the poaching region: if those residents move to region C and were not contributors in region A, benefits might rise in region A, making it possibly no longer in the interest of the originally poached residents to be poached. The following section will assume that non-contributors are never attracted and disallow this sort of equilibrium breakdown, allowing for poachings that are not equally from all regions and allowing a stronger result.

Definition 5 *A tax schedule T_i is ungenerous if $T_i(j, .) \geq 0$ for all $j \neq i$.*

In order to guarantee the existence of a poaching that reduces benefits equally in all other regions it was necessary to allow for quotas to restrict immigration. A criticism of stable equilibrium with quotas is that population can in effect be selected by the poaching region: the composition of the types filling the quota can be just right to preserve the poaching equilibrium. However in the proof above the way in which the quota is filled (given by S_x) is arbitrary, and could potentially be taken to be those residents who would gain the most from being in the quota, or some other plausible group. However the technical problem that quotas solve is relatively minor and they are not used in the following result.

3.6 A stronger result

Proposition 2 *In a stable efficiency equilibrium in which all regions set ungenerous tax schedules, all residents of all regions are content.*

Proof. Suppose there are discontent residents in region j and let $i \neq j$.

Choose one discontent resident of region j a, and let $w_e = \arg\max_w U(a, w, I(w))$, and choose ϵ such that $U(a, w_e, I(w_e) - \epsilon)$ is greater than the original utility.

Let $T_i'(j, w) = \min_{k \neq i} T_k(j, w) \geq 0$ for $j \neq i$ be, except $T_i'(j, w_e) = \epsilon$.

If population moves from region j to i, there is still no incentive for a resident of region j to choose some other region over region i, because j has lower tax schedules for immigrants from i. Also there is still no incentive for a resident of region i to move to any other region, since handouts have increased in region i since the tax schedule is ungenerous.

It follows that we can find an equilibrium in which population moves from j to i, by applying a Kakutani fixed point argument only considering the space of population moves from j to i. Let this space be D', a convex subset of distributions over $[n] \times [n] \times A^F \times W^F$. Consider the map from populations to handouts to populations after optimal movement and work given the handouts. Because of the incentive considerations above, under this map D' maps into itself. Using finite approximations as in the original argument for existence of equilibrium, applying the Kakutani fixed point theorem, and taking limits, we get a movement equilibrium d in D'.

In d it must be that some residents from j have move to region i and taken work w_e. If not then handouts in region j must have strictly increased (since otherwise ability a would prefer to move and work w_e). In that case all residents have a strict incentive to stay in region j rather than take any work in region i other than w_e; otherwise with original lower handouts ability b would have a strict incentive to move to region i and do $w' \neq w_e$, but then ability b would have had a strict incentive to move to region k in the original equilibrium, where $T_i'(j, w') = T_k(j, w')$. So there is no movement, and so handouts cannot have strictly increased.

So (T_i', d) is a poaching in which a positive measure immigrates from region j to region i and does work w_e, adding to government revenue in region i and giving a Pareto improvement. ■

Proposition 3 *If regions set ungenerous tax schedules which in isolation make all residents (weakly/strictly) content, (a/ the unique) resulting population movement and work choice involves no movement.*

Proof. No movement is an equilibrium because if any resident with ability a moves to another region he will make less than $\max_w U(a, w, I(w))$ while he is making weakly more than this by remaining since he is content.

Now suppose all residents are strictly content. To see that the equilibrium is unique, note that if a resident is strict content then he must be strictly non-contributing, so removing a positive measure of strictly content residents strictly increases handouts in that region, and if any immigrants join the region, since taxes are ungenerous handouts increase further. So there is even less incentive to leave than before, and so there is no equilibrium in which a positive measure of residents leaves a region. ■

3.7 Discussion

For all residents to be content, none can be contributing, so instead of redistribution, tax policy concerns distribution of independent resources r_i. Moreover not all distributions of resources will be possible: there is a moral hazard problem and distributions favouring lower levels of work (or equivalently income) will reduce the level of handouts, and may make some residents discontent, even though they are not contributing. A mechanism-designer government with given preferences must solve the moral hazard problem subject to keeping all abilities content.

Problems of multiple equilibria are worked around by assuming that "poaching" can be done successfully. An alternative approach that could be to assume that movement/work equilibria vary continuously in response to tax schedule choice. A competitive equilibrium or auction approach to buying citizens might encapsulate some results and allow generalization of the model, but it would be hard to incorporate work incentives. Replacing handouts with a public good, and having economies (income functions) that adjust to the population distribution are important issues that await further exploration.

Bibliography

[1] D. Abreu, & F. Gul (2000): "Bargaining and Reputation", Econometrica

[2] D. Abreu, , & D. Pearce (2006): "Reputational Wars of Attrition with Complex Bargaining Postures", mimeo

[3] A. Ambrus and P. Pathak (2007): "Cooperation Over Finite Horizons: A Theory and Experiments"

[4] I.-K. Cho (1987): "A refinement of sequential equilibrium", Econometrica

[5] I.-K. Cho & D. Kreps (1987): "Signaling games and stable equilibria", Quarterly Journal of Economics

[6] I.-K. Cho & J. Sobel (1990): "Strategic stability and uniqueness in signaling games", Journal of Economic Theory

[7] M.W. Cripps, E. Dekel & W. Pesendorfer (2005): "Reputation with Equal Discounting in Games with Strictly Conflicting Interests", Journal of Economic Theory

[8] M. Celentani, D. Fudenberg, D.K. Levine, & W. Pesendorfer (1996): "Maintaining a reputation against a long-lived opponent", Econometrica

[9] M.W. Cripps., G.J. Mailath & L. Samuelson (2007): "Disappearing private reputations in long-run relationships," Journal of Economic Theory

[10] M. Dufwenberg & G. Kirchsteiger (2004): "A theory of sequential reciprocity", Games and Economic Behavior

[11] R. Evans & J.P. Thomas (1997): Reputation and Experimentation in Two-Player Repeated Games with Discounting, Econometrica

[12] D. Fudenberg & D. Levine (1989): Reputation and equilibrium selection in games with a patient player. Econometrica

[13] D. Fudenberg & D. Levine (1992): Maintaining a Reputation when Strategies are Imperfectly Observed, Review of Economic Studies

[14] D. Fudenberg & J. Tirole (1986): "A "Signal-Jamming" Theory of Predation", RAND Journal of Economics

[15] B. Holmstrom (1999): "Managerial incentive problems: A dynamic perspective", Review of Economic Studies

[16] A. Kaya (2007): "Repeated Signaling Games", mimeo

[17] D. Epple & T. Romer (1991): "Mobility and Redistribution", Journal of Political Economy

[18] E. Kohlberg & J.-F. Mertens (1986): "On the Strategic Stability of Equilibria," Econometrica

[19] D. Kreps & R. Wilson (1982): "Reputation and imperfect information", Journal of Economic Theory

[20] G.J. Mailath, M. Okuno-Fujiwara, and A. Postlewaite (1993): "Belief based refinements", Journal of Economic Theory

[21] G.J. Mailath & L. Samuelson (2001): "Who Wants a Good Reputation?", Review of Economic Studies

[22] J.A. Mirrlees (1971): "An Exploration in the Theory of Optimal Income Taxation", 1971, Review of Economic Studies

[23] G. Noldeke & L. Samuelson (1995): "A Dynamic Model of Equilibrium Selection in Signaling Markets", working paper

[24] M. Rabin (1993): "Incorporating Fairness Into Game Theory and Economics", American Economic Review

[25] G. Ramey (1996): "D1 Signalling Equilibria with Multiple Signals and a Continuum of Types", Journal of Economic Theory

[26] C. Roddie (2008): "Repeated Signalling and Reputation", mimeo (Chapter 1 of this thesis)

[27] J. Sobel (2007): "Signaling games", technical report, UCSD

[28] M. Spence (1973): "Job Market Signaling", Quarterly Journal of Economics