Supplement to "Ruth, Anthony, and Clarence"

JEFFREY C. ELY
Department of Economics, Northwestern University

This supplement proves Theorem 2 in the paper.

1. PROOF OF THEOREM 2

Notation We let Σ_i denote the set of stage-game strategies for $i \in \{R, C\}$. A typical strategy is denoted a_i and α_i represents a mixed strategy for player i.

Ruth's stage-game expected payoffs from a given mixed profile $\alpha = (\alpha_R, \alpha_C)$ depends on the probability that Clarence is biased. Denote by $u_R^1(\alpha)$ and $u_R^0(\alpha)$ Ruth's expected stage game payoff from profile α when Clarence is biased and unbiased respectively and then

$$u_R^p(\alpha) = pu_R^1(\alpha) + (1-p)u_R^0(\alpha).$$

Clarence's stage game payoffs $u_C(\alpha)$ do not depend on p. Write $u^p(\alpha) = (u_R^p(\alpha), u_C(\alpha))$ for the expected stage-game payoff vector.

LEMMA A.1. *Consider the mixed strategy $\underline{\alpha}_C$ for Clarence which plays i with probability $2c$ and l with the remaining probability. Regardless of p, voting informatively is a stage best-response for Ruth. When the profile $(i, \underline{\alpha}_C)$ is played Ruth earns stage expected payoff $\underline{v}_R^p = \bar{v} - cp$ and Clarence earns expected stage payoff $\underline{v}_C = 2c(v^* - c) + (1-2c)\bar{v}$.*

PROOF. The following is Ruth's expected payoff from voting informatively against $\underline{\alpha}_C$.

$$2c(1-p)(v^* - c) + (1 - 2c(1-p))(\bar{v} - c)$$

because with probability $2c(1-p)$ Clarence acquires information and votes according to his unbiased signal. In this case Ruth, also voting informatively, earns $v^* - c$. With the remaining probability Clarence's vote is uninformative and Ruth earns $\bar{v} - c$. We can rewrite the payoff as follows

$$\bar{v} + 2c(1-p)(v^* - \bar{v}) - c$$

Recall that $v^* - \bar{v} = \bar{v} = 1/2$. The payoff becomes

$$\bar{v} + 2c(1-p)\bar{v} - c$$
$$= \bar{v} - 2cp\bar{v} + 2c\bar{v} - c$$
$$= \bar{v}(1 - 2cp)$$

which is also equal to Ruth's expected payoff from voting r uninformatively against $\underline{\alpha}_C$. To see why note that with probability $2cp$ Clarence acquires information but votes according to his biased information, i.e. he votes r uninformatively. In this case Ruth earns zero from voting r. With the remaining probability Ruth earns $\bar{v}$ because either Clarence (and Anthony) vote informatively, or Clarence votes l uninformatively counteracting Ruth.

Thus, the strategy $\underline{\alpha}_C$ makes Ruth indifferent between informative voting and voting r. She earns a payoff of $\bar{v}(1-2cp)$ from either strategy. This payoff is smallest when $p=1$ when the payoff equals $\bar{v}(1-2c)$ and when $c<1/4$ always earns at least a payoff of $1/4$ from voting informatively against $\underline{\alpha}_C$. On the other hand Ruth's payoff from voting l against $\underline{\alpha}_C$ is $2c\bar{v}$. That's because with probability $1-2c$ Clarence votes l uninformatively and Ruth voting l would earn zero. With the remaining $2c$ probability Ruth earns $\bar{v}$ because regardless of whether Clarence votes informatively or r (because he is biased) Ruth earns $\bar{v}$ from voting l. Finally since $\bar{v}=1/2$ Ruth's payoff from voting l equals c which is smaller than $1/4$ and therefore smaller than her payoff from voting informatively.

We have shown that voting informatively is a best response for Ruth against $\underline{\alpha}_C$ and it earns her an expected payoff of $\bar{v}(1-2cp)=\bar{v}-cp=\underline{v}_R^p$. When Ruth votes informatively and Clarence plays $\underline{\alpha}_C$ he earns the payoff from informative voting v^*-c with probability $2c$ and the payoff $\bar{v}$ with the remaining probability. □

Outcomes The *outcome* of the stage game will be an element $y=(y_R,y_C)$ of $Y=\{l,r\}^2$ indicating the vote profile of Ruth and Clarence.[1]

We need notation for the probability distribution over outcomes as perceived differently by Ruth and Clarence. For any profile α and outcome y, let $m^0(y \mid \alpha)$ be the probability that α generates outcome y when Clarence is unbiased. And let $m^1(y \mid \alpha)$ be the probability when Clarence is biased. When Ruth believes that with probability p Clarence is biased, she assigns probability

$$m^p(y \mid \alpha)=pm^1(y \mid \alpha)+(1-p)m^0(y \mid \alpha)$$

to outcome y. Meanwhile Clarence, who is mis-specified, assigns probability $m^0(y \mid \alpha)$

For any belief p and profile α define for each y such that $m^p(y \mid \alpha)>0$ Ruth's Bayesian posterior following the outcome y:

$$\beta(p \mid \alpha, y)=\frac{p \cdot m^1(y \mid \alpha)}{m^p(y \mid \alpha)}$$

Note that because Ruth's and Clarence's votes are statistically independent after any history, the updated belief depends only on Clarence's mixed strategy and Clarence's realized vote.

[1] These are the outcomes on which Ruth and Clarence's continuation payoffs will depend. There will be no benefit to conditioning continuation payoffs on the vote of Anthony so we ignore his vote in the outcome.

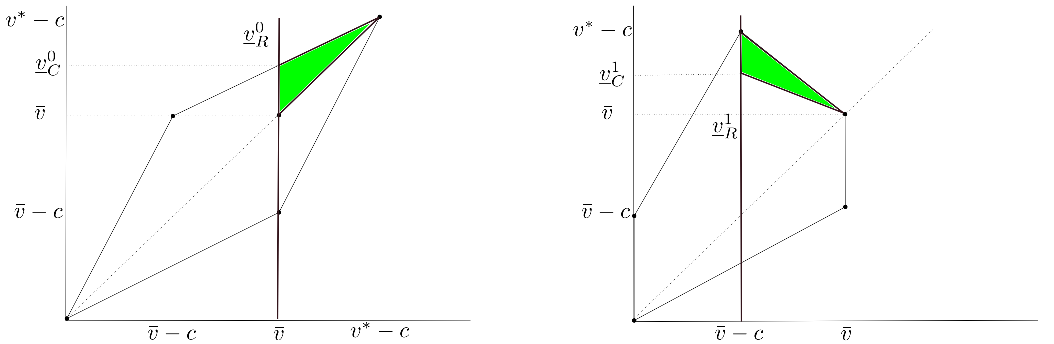

Figure (a) $p = 0$ Figure (b) $p = 1$

FIGURE 1. Feasible Payoff Vectors. Ruth's on the horizontal axis, Clarence's on the vertical axis

A continuation value function is a map $w : Y \to \mathbf{R}^2$ where $w_i(y)$ indicates the expected continuation payoff promised to player i after the vote profile y. Write

$$\mathbf{E}_C(w \mid \alpha) = \sum_Y m^0(y \mid \alpha) w_C(y)$$

and

$$\mathbf{E}_R^p(w \mid \alpha) = \sum_Y m^p(y \mid \alpha) w_R(y)$$

and use

$$\mathbf{E}^p(w \mid \alpha) = (p, \mathbf{E}_R^p(w \mid \alpha), \mathbf{E}_C(w \mid \alpha))$$

to denote the expected outcome vector. Note that we are using Ruth's correct beliefs for the expected posterior, namely p.

Feasible Payoffs Denote by $K = [0,1] \times \mathbf{R}^2$ the set of triples (p, v_R, v_C) such that p is a belief, and (v_R, v_C) is an arbitrary payoff vector for Ruth and Clarence. For any subset $W \subset K$ we will use the notation W^p for the "slice" consisting of all elements whose belief coordinate is p.

Each slice $W^{p'}$ is isomorphic to W^p for all p', p by the linear bijection $\varphi_{p,p'}$ that maps (p, v_R, v_C) in W^p to

$$\varphi_{p,p'}(v) = \left(L^{p'} + (v - L^p)\right) \in W^{p'}$$

We will also sometimes make use of the restriction of φ to its payoff components:

$$\varphi_{p,p'}(v_R, v_C) = \left(v_R + (p' - p)(v_R^1 - v_R^0), v_C\right) \tag{A.1}$$

Let V^p denote the set of feasible stage payoffs at p, and

$$F = \prod_{p \in [0,1]} V^p \subset K$$

the correspondence of beliefs to feasible stage payoffs. Refer to Figure 1. The feasible payoffs stage payoffs for $p = 0$ and $p = 1$ are depicted. For $p = 0$, we have the complete information, symmetric feasible payoff set. The extreme points are $(v^* - c, v^* - c)$ achieved by the informative voting profile, $(0,0)$ achieved by Ruth and Clarence casting the same, uninformed, vote, and $(\bar{v}, \bar{v} - c)$ achieved by Clarence voting informatively and Ruth voting uninformatively (along with the permuted payoffs/profile).

The extreme points of the stage feasible payoffs when $p = 1$ are as follows. When Ruth and Clarence both vote informatively Clarence earns $v^* - c$ and Ruth earns $\bar{v} - c$. Note that the latter payoff is also Ruth's minmax payoff $\underline{v}_R^1$. When Ruth and Clarence polarize, the each earn $\bar{v}$. When Clarence votes informatively but Ruth votes l they earn $\bar{v} - c$ and $\bar{v}$ respectively. On the other hand when Clarence votes informatively and Ruth votes r Clarence again earns $\bar{v} - c$ but Ruth earns 0. Finally the payoff vector $(0,0)$ is again achieved by Ruth and Clarence casting the same uninformed vote.

Clarence's first-best payoff is $v^* - c$. The associated payoff for Ruth depends on p and equals $(1-p)(v^* - c) + p \cdot \underline{v}_R^1$.

THEOREM 2. *For any $\varepsilon > 0$ for large enough discount factors there exists an equilibrium yielding payoffs within ε of Clarence's first-best.*

Note that for p close to 0, i.e. in the neighborhood of complete information, both Ruth and Clarence earn payoffs approaching first-best $v^* - c$.

Definitions Consider the two triangles depicted in Figure 1. The first, T_0 is the subset of the slice W_0 with vertices $(\bar{v}, \bar{v}), (v^* - c, v^* - c), \left(\underline{v}_R^0, \underline{v}_C\right)$. The second, T_1, is the subset of the slice W_1 which has as its set of vertices $(\bar{v}, \bar{v}), (\bar{v} - c, v^* - c), \left(\underline{v}_R^1, \underline{v}_C\right)$.

Pick $v_C \in (\bar{v}, v^* - c)$ and take v_R^0, and v_R^1 such that $v^0 = (v_C, v_R^0)$ belongs to the interior of the T_0 and $v^1 = (v_C, v_R^1)$ belongs to the interior of T_1. Consider the line L in K connecting the points $\left(0, v^0\right)$ and $\left(1, v^1\right)$, with L^p denoting the point on the line with belief coordinate p.

Note that the point L^p belongs to the interior of the triangle, call it T_p, in W^p whose vertices are $(\bar{v}, \bar{v}), (v^* - c, v^p - c), \left(\underline{v}_R^p, \underline{v}_C\right)$.

Pick a radius $\varepsilon > 0$ such that the balls of radius ε around $\left(0, v^0\right)$ and $\left(1, v^1\right)$ are contained in the interiors of their respective triangles. The "cylinder"

$$W = \bigcup_{p \in [0,1]} B(L^p, \varepsilon)$$

is included in the interior of F, and each $B(L^p, \varepsilon)$ is included in the interior of the triangle T_p.

DEFINITION A.1. *A profile α is enforceable at p with respect to a subset $U \subset K$ and discount factor δ if there exists a payoff vector v and a function $w : Y \to U$ such that*

1. $w_0(y) = \beta(p \mid \alpha, y)$ *for each outcome y such that* $m_R^p(y \mid \alpha) > 0$.
2. $v_i = (1-\delta)u^p(\alpha) + \delta \mathbf{E}^p(w_i \mid \alpha)$ *for each i*

3. $v_i \geq (1-\delta)u^p(a_i, \alpha_{-i}) + \delta \mathbf{E}_i^p(w_i \mid a_i, \alpha_{-i})$ *for all* a_i.

We say that w enforces α at p and we say that (p, v) is generated by the set U given δ.

Given a belief p, a profile α and a continuation value function w, let $B_i^p(\alpha \mid w)$ be the set of pure strategies a_i that satisfy item 3 with equality, i.e the set of *best-replies* for player i. We will make use of the following observation.

LEMMA A.2. *For any* p*,* α*, and* w *there are at most two pure strategies* a_i *for player* i *in* $B_i^p(\alpha \mid w)$.

PROOF. Denote by i, l, and r, the pure stage-game strategies of informed voting and uninformed voting l and r respectively. By Lemma 2 in the text no other stage-game strategy is enforceable. When $c < 1/4$ at most one of l and r can belong to $B_i^p(\alpha \mid w)$. Because if both are elements, then also a 50-50 mixture between them is a best-reply, say for Ruth. The expected stage-payoff for Ruth from a 50-50 mixture is

$$q\bar{v} + (1-q)\frac{1}{4}$$

where q is the probability under α, given p, that Clarence is unbiased and votes informatively. Whereas voting informatively would yield

$$q(v^* - c) + (1-q)(\bar{v} - c)$$

and since $c < 1/4$, we have $v^* - c > \bar{v}$ and $\bar{v} - c > 1/4$. Thus informed voting yields a strictly larger stage-game expected payoff but exactly the same distribution over voting outcomes and hence continuation payoffs. □

We will show that for all discount factors close enough to 1, the set W is *self-generating* That is, there exists $\bar{\delta} < 1$ such that for all $\delta > \bar{\delta}$ every (p, v) is generated by W given δ. By standard arguments this implies that for every element of $(p, v) \in W$, there exists an equilibrium of the game beginning with prior p that yields Ruth and Clarence the expected discounted average payoffs v_R and v_C.

Consider a point (p, v) on the boundary of W, and the plane $\bar{H}$ that is tangent to W at (p, v). The halfspace H defined by $\bar{H}$ and that includes W is called the tangent halfspace at (p, v).

Similarly, if v is on the boundary of W^p for some p, then the (two-dimensional) halfspace tangent to the slice W^p and includes W^p is denoted H^p. That halfspace consists of all payoff vectors v' such that $(v - L^p) \cdot v' \leq (v - L^p) \cdot v$ and $v - L^p$ is the normal vector to H^p.

Note that a point (p, v') belongs to the halfspace H tangent to W at (p, v) if and only if the payoff vector v' belongs to the the two-dimensional halfspace tangent to W^p at v.

Moreover, these two-dimensional halfspaces are isomorphic under the projection mapping $\tilde{\varphi}_{p,p'}$. Specifically the inequality

$$(\tilde{\varphi}_{p,p'}(v) - L^{p'}) \cdot \tilde{\varphi}_{p,p'}(v') \leq (\tilde{\varphi}_{p,p'}(v) - L^{p'}) \cdot \tilde{\varphi}_{p,p'}(v)$$

is equivalent to

$$(v - L^p) \cdot [L^{p'} + (v' - L^p)] \leq (v - L^p) \cdot [L^{p'} + (v - L^p)]$$

which after canceling terms is equivalent to

$$(v - L^p) \cdot v' \leq (v - L^p) \cdot v.$$

It follows from these observations that a vector (p, v) belongs to a halfspace H tangent to W if and only if the projection $\varphi_{p,p'}(p, v)$ belongs to H as well. This fact will be used in the proof of the next proposition.

PROPOSITION A.1. *The set W is decomposable on tangent halfspaces. That is, for every (p, v) on the boundary of W, there is a profile α such that $(p, u^p(\alpha))$ is separated from W by the halfspace H that is tangent to W at the point (p, v) and α is enforceable with respect to H at p using a continuation function w that satisfies*

$$\mathbf{E}^p(w \mid \alpha) \in \bar{H} \tag{A.2}$$

PROOF. If v is on the boundary of W^p, say that v is separated from a payoff vector v' if the line tangent to W^p at v separates the point v' from W^p. Precisely, if (a, b) is a vector normal to the line tangent to W^p at v then $(a, b) \cdot v' \geq (a, b) \cdot v \geq w$ for all $w \in W^p$. Say that v is *regular* if the normal vector (a, b) is such that $a, b \neq 0$.

For every p, every point on the boundary of W^p is separated from at least one of the following four payoff vectors: $(0, 0)$, $(\bar{v}, \bar{v})$, $\left(\underline{v}_R^p, \underline{v}_C\right)$, and $(v^p - c, v^* - c)$ where $(v^p - c)$ is the expected payoff to Ruth from informative voting when p is the probability Clarence is biased. This is by construction because the slice W^p is a subset of the triangle T_p.

Case 1 Fix p and consider any point (p, v) such that v is separated from $(\bar{v}, \bar{v})$. Note that polarized voting (Ruth voting l, and Clarence voting r, both without acquiring information) is a stage Nash equilibrium that yields payoff vector $(\bar{v}, \bar{v})$ and produces the outcome (l, r) with probability 1. We can pick a continuation function w that is constant, e.g. $w(\cdot, \cdot) = (p, v)$. Since polarized voting is a stage Nash equilibrium the continuation function w enforces polarized voting. Equation A.2 is trivially satisfied.

Case 2 Next consider a regular point v that is separated from $(0, 0)$. The pure profile (r, r) yields payoff vector $(0, 0)$ and the outcome (r, r) with probability 1 so to enforce (r, r) we begin by setting $w(r, r) = (p, v)$. To satisfy item 3, we pick $x > 0$ large enough that

$$(1 - \delta)\bar{v} + \delta\,(v - x) < \delta v. \tag{A.3}$$

and

$$(1 - \delta)(\bar{v} - c) + \delta\left(\frac{1}{2}v + \frac{1}{2}(v - x)\right) < \delta v. \tag{A.4}$$

and set $w(l, r) = (p, v - x, v + \Delta_C)$ and $w(r, l) = (p, v + \Delta_R, v - x)$. The left-hand sides of Equation A.3 and Equation A.4 are then the expected payoffs to either Ruth or Clarence from deviating to l or i respectively, and item 3 is satisfied. Because v is regular we may pick the "rewards" $\Delta_R > 0$ and $\Delta_C > 0$ to ensure that $w(l, r)$ and $w(r, l)$ lie on the plane tangent to W at (p, v).

Finally, since the outcome (l, l) can arise only from a double deviation, we may choose $w(l, l)$ to be an arbitrary point on that plane. Condition Equation A.2 is trivially satisfied because the continuation $w(r, r) = (p, v)$ occurs with probability 1.

Case 3 Next consider any point v separated from $(v^p - c, v^* - c)$. We will take α to be informative voting. This is always a stage best-reply for Clarence. Indeed if $p \leq 1 - 2c$ then it is also a stage best-reply for Ruth, hence a stage Nash equilibrium and we can proceed as we did with polarized voting.

If instead $p > 1 - 2c$ then informative voting is not a stage best-reply for Ruth (voting l is the unique best-reply) and we therefore have to design continuation values to incentivize Ruth. In this case suppose v maximizes Clarence's payoff within the slice W^p. In other words $a = 0$ and $b > 0$. Choose $\underline{w}_R$ and $\bar{w}_R$ to satisfy

$$(1 - \delta)(v^p - c) + \delta\left(\frac{\underline{w}_R + \bar{w}_R}{2}\right) = (1 - \delta)\bar{v} + \delta\underline{w}_R. \tag{A.5}$$

The left-hand side will be the expected payoff from voting informatively and the right-hand side the expected payoff from voting l. These continuation values will therefore make Ruth indifferent between these two stage-game strategies and hence by Lemma A.2 they enforce informative voting.

Pick a constant continuation value for Clarence equal to v_C. Note that any vector of the form (v'_R, v_C) belongs to the line tangent to W^p because $a = 0$ and $b > 0$ implies $(a, b) \cdot (v'_R, v_C) = (a, b) \cdot v$. In particular this is true of the continuation payoff vectors $(\underline{w}_R, v_C)$ and $(\bar{w}_R, v_C)$.

Now set

$$\begin{aligned} w(l, r) &= \varphi_{p,\beta(p|\alpha,r)}(\underline{w}_R, v_C) \\ w(l, l) &= \varphi_{p,\beta(p|\alpha,l)}(\underline{w}_R, v_C) \\ w(r, r) &= \varphi_{p,\beta(p|\alpha,r)}(\bar{w}_R, v_C) \\ w(r, l) &= \varphi_{p,\beta(p|\alpha,l)}(\bar{w}_R, v_C). \end{aligned}$$

Because Clarence's continuation value is constant and informed voting is a stage best-reply we have satisfied item 3 for Clarence. As for Ruth, since the projection mapping $\varphi_{p,p'}$ shifts her payoffs by a constant (see Equation A.1), item 3 follows from Equation A.5. Equation A.2 is satisfied because the expected continuation is $(p, \mathbf{E}^p_R(w \mid \alpha), v_C)$ which is on the line tangent to W^p.

Let's finally consider points v that are separated *only* from $(\underline{v}^p_R, \underline{v}_C)$ and not from any of the other three points, i.e. none of cases 1-3 apply. I claim that for such a point v, the tangent line $\bar{H}^p$ has normal vector (a, b) with $a \neq 0$ and $b \geq 0$. For if $a \geq 0$ and $b < 0$ then $(a, b) \cdot (\bar{v}, \bar{v}) \geq (a, b) \cdot (\underline{v}^p_R, \underline{v}_C)$ and case 1 applies. And if $a < 0$ and $b < 0$, then v is normal and $(a, b) \cdot (0, 0) > (a, b) \cdot (\underline{v}^p_R, \underline{v}_C)$ and case 2 applies. Finally if $a = 0$ and $b > 0$ then $(a, b) \cdot (v^p - c, v^* - c) > (a, b) \cdot (\underline{v}^p_R, \underline{v}_C)$ and case 3 applies.

We will take α to be $(i, \underline{\alpha}_C)$. Pick $\bar{w}_C$ and $\underline{w}_C$ to satisfy

$$(1-\delta)(v^* - c) + \delta\left(\frac{\bar{w}_C + \underline{w}_C}{2}\right) = (1-\delta)\,\bar{v} + \delta \bar{w}_C. \tag{A.6}$$

and note that $\bar{w}_C > \underline{w}_C$. The left-hand side is Clarence's expected payoff from voting informatively and the right-hand side is his expected payoff from voting l. These continuation values thus make Clarence indifferent between these two stage-game strategies and by Lemma A.2 the mixture $\underline{\alpha}_C$ is a best-reply for Clarence.

Next pick an expected continuation value for Ruth, y, such that

$$(y, z) \in \bar{H}^p \tag{A.7}$$

where

$$z = (1-2c)\,(\bar{w}_C) + 2c\left(\frac{\bar{w}_C + \underline{w}_C}{2}\right)$$

is Clarence's expected continuation value. To achieve expected continuation value y for Ruth we then choose values $\bar{w}_R$ and $\underline{w}_R$ so that

$$y = (1-2c)\,(\bar{w}_R) + 2c\underline{w}_R$$

and so that the continuations $(\bar{w}_R, \bar{w}_C)$ and $(\underline{w}_R, \underline{w}_C)$ belong to a line, call it $\tilde{H}^p$, parallel to $\bar{H}^p$. This is possible because $a \neq 0$.

Notice that Ruth's expectation of Clarence's continuation value is smaller than Clarence's expectation of his own continuation value

$$z' := (1-2c)\,(\bar{w}_C) + 2c\underline{w}_C < z$$

because $\bar{w}_C > \underline{w}_C$.

Now the vector (y, z') is Ruth's expectation of a pair of vectors that belong to $\tilde{H}^p$ and therefore $(y, z') \in \tilde{H}^p$. And since $b \geq 0$ and $z' < z$ we have

$$(a, b) \cdot (y, z) \geq (a, b) \cdot (y, z') = (a, b) \cdot (\bar{w}_R, \bar{w}_C) = (a, b) \cdot (\underline{w}_R, \underline{w}_C)$$

where the equalities follow from the fact that (a, b) is the normal of $\bar{H}^p$ and therefore also of $\tilde{H}^p$. The inequalities establish that the continuation value vectors belong to the halfspace H^p.

As such we can use the continuation function

$$w(r, r) = w(l, r) = \varphi_{p, \beta(p|\alpha, r)}(\underline{w}_R, \underline{w}_C)$$
$$w(r, l) = w(l, l) = \varphi_{p, \beta(p|\alpha, l)}(\bar{w}_R, \bar{w}_C)$$

which then belong to the halfspace H tangent to W at the point (p, v) and satisfy Equation A.2.

Finally, since Ruth's continuation value is independent of her own vote and by Lemma A.1 voting informatively is a stage best-response for Ruth, we have satisfied item 3 for Ruth. As for Clarence, since his component of the projection mapping $\varphi_{p,p'}$ is the identity (see Equation A.1) item 3 follows from Equation A.6.

□

PROPOSITION A.2. *The set W is locally self-generating. That is, for every $(p,v) \in W$ there exists δ and an open neighborhood U of (p,v) such that each $w \in U \cap W$ is generated by W given δ.*

LEMMA A.3. *If α is enforceable at p with respect to a halfspace H and some discount factor δ using a continuation function w that satisfies Equation A.2 then there is a constant κ such that for all discount factors δ' and payoff vectors v' there exists a continuation function $w' : Y \to H'$ where $H' = H + [(p,v') - \mathbf{E}^p(w \mid \alpha)]$ such that w' enforces α with respect to H' and δ' and*

1. $(p,v') = \mathbf{E}^p(w' \mid \alpha) \in \bar{H}'$
2. $\|\varphi_{w_0(y),p} w'(y) - (p,v')\| < \kappa \frac{1-\delta'}{\delta'}$ *for all* y.

In light of Lemma A.3 we can talk about a profile being enforced by a halfspace without reference to any specific discount factor.

PROOF. Let w be the continuation value function that enforces α at belief with respect to H and δ. Write $\bar{w} = \mathbf{E}^p(w \mid \alpha)$ and set

$$w'(y) = \varphi_{p,w_0(y)} v' + \frac{\delta(1-\delta')}{\delta'(1-\delta)} \left[w(y) - \varphi_{p,w_0} \bar{w} \right].$$

Straightforward computation shows that w' enforces α with respect to H' and δ' and that item 1 and item 2 in the statement are satisfied. □

PROOF OF PROPOSITION A.2. Pick any (p,v) on the boundary of W, i.e. v belongs to the boundary of W^p. Since W is decomposable on tangent planes, there is a profile α such that $(p, u^p(\alpha))$ is separated from W by the halfspace H that is tangent to W at the point (p,v) and there exists $w : Y \to H$ which enforces α at p with $\mathbf{E}^p(w \mid \alpha) \in \bar{H}$. For any δ, pick v' so that

$$v = (1-\delta) u^p(\alpha) + \delta v'$$

and note that once δ is close enough to 1, the vector (p,v') belongs to the interior of W^p due to the fact that H separates $(p, u^p(\alpha))$ from W.

By Lemma A.3 we may choose $w'(\cdot)$ to enforce α at p with respect to δ and satisfy.

1. $w'(y) \in H' = H + [(p,v') - \mathbf{E}^p(w \mid \alpha)]$
2. $(p,v') = \mathbf{E}^p(w' \mid \alpha) \in \bar{H}'$
3. $\|\varphi_{w'_0(y),p} w'(y) - (p,v')\| < \kappa \frac{1-\delta}{\delta}$ for all y.

The remainder of the proof follows FLM. The distance from (p,v') to the boundary of W^p in the direction of $(p,u^p(\alpha))$ is of order $(1-\delta)/\delta$, as is the distance from (p,v') to each $\varphi_{w_0'(y),p}w'(y)$. But because H is tangent to W, the distance from (p,v') along H^p is of order no less than $\sqrt{(1-\delta)/\delta}$. Thus, for δ close enough to 1, each $\varphi_{w_0'(y),p}w'(y)$ belongs to the interior of W^p, and therefore each $w'(y)$ belongs to the interior of $W^{w_0'(y)}$.

That is, the vector (p,v) is generated by the interior of W given δ using the profile α and the continuation value function w'. Let us now parameterize a family of continuation value functions as follows.

$$\hat{w}_i(y) = w_i'(y) + \mu_i + \nu_i \mathbf{1}_{y_i=r}$$

for $i = R, C$ and

$$\hat{w}_0(y) = w_0'(y) = \beta(p \mid \alpha, y).$$

For $\mu_R = \mu_C = \nu_R = \nu_C = 0$ we have the continuation value function w'. We can express the conditions item 3 and item 2 from Definition A.1 in the following form:

$$(1-\delta)u_i^p(a_i,\alpha_{-i}) + \delta\left[\mathbf{E}_i^p(\hat{w}_i \mid a_i,\alpha_{-i})\right] = (1-\delta)u_i^p(a_i',\alpha_{-i}) + \delta\left[\mathbf{E}_i^p(\hat{w}_i \mid a_i',\alpha_{-i})\right]$$

for each i and for distinct pure strategies a_i and a_i', if any, in the support of α_i (recall that by Lemma A.2 there can be at most two of these), and

$$v_i = (1-\delta)u_i^p(a_i,\alpha_{-i}) + \delta\left[\mathbf{E}_i^p(\hat{w}_i \mid a_i,\alpha_{-i})\right]$$

Using our parameterization, the first condition reduces to

$$(1-\delta)u_i^p(a_i,\alpha_{-i}) + \delta\left[\mathbf{E}_i^p(w_i \mid a_i,\alpha_{-i}) + \mu_i + \nu_i m_i(r \mid a_i)\right] = (1-\delta)u_i^p(a_i',\alpha_{-i}) + \delta\left[\mathbf{E}_i^p(w_i \mid a_i',\alpha_{-i}) + \mu_i + \nu_i m_i(r \mid a_i')\right]$$

or

$$\nu_i = \frac{1-\delta}{\delta}\left[\frac{u_i^p(a_i',\alpha_{-i}) - u_i^p(a_i,\alpha_{-i})}{m_i(r \mid a_i) - m_i(r \mid a_i')}\right]$$

The second condition can be rearranged to

$$\mu_i + m_i(r \mid a_i)\nu_i = v_i - (1-\delta)u_i^p(a_i,\alpha_{-i}) - \delta\mathbf{E}_i^p(w_i \mid a_i,\alpha_{-i})$$

In particular the parameters μ_i and ν_i are continuous functions of the vector (p,v). For $(\tilde{p},\tilde{v})$ close enough to (p,v), the enforcing continuation values $\hat{w}(\cdot,\cdot)$ remain in W. It follows that there exists an open neighborhood U such that every $(\tilde{p},\tilde{v}) \in U \cap W$ is generated by W, concluding the proof. □

PROPOSITION A.3. *There exists a $\delta' < 1$ such that for all $\delta \geq \delta'$ the set W is self-generating with respect to δ. As a result, for every prior belief $p \in [0,1]$ and every payoff vector $v \in$*

W^p, there exists an equilibrium of the repeated game in which Ruth and Clarence earn expected payoff profile v.

PROOF. Since W is locally self-generating there exists an open covering of W such that for every open set U in the covering there exists a discount factor δ_U such that each (p, v) belonging to U is generated by W given δ_U. Since W is compact there exists a finite sub-covering. Let δ' be the largest among the discount factors associated with sets in the finite sub-covering. Then W is self-generating with respect to δ' and all discount factors larger.

For any (p, v) in W we now construct an equilibrium strategy profile yielding the payoff vector v. Because W is self-generating given δ' there is a profile α and a continuation value function w which maps to W and generates (p, v). In the initial period the players play the stage-game profile profile. Following the outcome y, by the definition, the posterior belief will be $w_0(y)$ and since $w(y) \in W$ there exists a profile α^y which is enforced given δ' by a further continuation value function mapping to W and generates $w(y)$. We specify that α^y is played after the single-period history y. Continuing in this way we construct the stage profile to be played after every (public) history, i.e. the full strategy profile. By the one-stage deviation principle this strategy profile is an equilibrium and by the standard telescoping argument it yields the payoff vector v. □

Proof of the theorem

PROOF OF THEOREM 2. Pick v_C within ε of Clarence's first-best payoff $v^* - c$. There exist payoffs v_R^1 and v_R^0 such that $v_R^1 > \underline{v}_R^1$, $v_R^0 > \underline{v}_R^0$ and both (v_R^1, v_C) and (v_R^0, v_C) are in the interior of the stage feasible payoff sets V^1 and V^0 respectively. Refer to Figure 1.

Take any cylinder centered around the line connecting (v_R^1, v_C) and (v_R^0, v_C) of small enough radius to be included in the correspondence of feasible payoffs. By Proposition A.3 for discount factors close enough to 1 the cylinder is self-generating and every point it contains is an equilibrium payoff for the associated prior. In particular this includes all points on the line connecting (v_R^1, v_C) and (v_R^0, v_C), and each of these is within ε of Clarence's first-best payoff vector at the associated prior. □