Available online at www.sciencedirect.com
ScienceDirect
Games and Economic Behavior 62 (2008) 67–76
GAMES and Economic Behavior
www.elsevier.com/locate/geb

An evolutionary analysis of the volunteer's dilemma

David P. Myatt [a], Chris Wallace [b,*]

[a] *Department of Economics, University of Oxford, UK*
[b] *Department of Economics and Trinity College, University of Oxford, UK*

Received 26 July 2006
Available online 27 April 2007

Abstract

A public good is produced if and only if a volunteer provides it. There are many pure-strategy Nash equilibria in each of which a single player volunteers. Noisy strategy revisions (for instance, quantal responses) allow play to evolve. Equilibrium selection is achieved via the characterisation of long-run play as revisions approximate best replies. The volunteer need not be the lowest-cost player: relatively high-cost, but nonetheless "reliable" players may instead produce the public good. More efficient players provide when higher values are associated with lower costs. Voluntary open-source software provision offers a contemporary application.

JEL classification: C72; C73; H41

Keywords: Volunteer's dilemma; Public goods; Evolution; Equilibrium selection; Quantal response

1. The volunteer's dilemma

A public good is produced if and only if at least one player volunteers to provide it. There are many pure-strategy Nash equilibria involving voluntary provision by a single player. An equilibrium-selection problem arises: who will volunteer?

Studies of the symmetric version of this familiar game (Diekmann, 1985) have often focused on the symmetric mixed-strategy equilibrium and its Bayesian–Nash counterpart for incomplete-information games (Weesie, 1994; Johnson, 2002). Such mixed equilibria have counter-intuitive

* Corresponding author.
E-mail addresses: david.myatt@economics.ox.ac.uk (D.P. Myatt), christopher.wallace@economics.ox.ac.uk (C. Wallace).

0899-8256/$ – see front matter
doi:10.1016/j.geb.2007.03.005

and counter-evidential properties. An (even slightly) asymmetric volunteer's dilemma exemplifies: players with low provision costs volunteer with low probability in order to maintain others' indifference.[1] This is somewhat absurd.

This paper selects a pure-strategy equilibrium via the study of evolving play: strategies are periodically revised by players who usually choose myopic best replies to the current state of play, but occasionally "mutate" against the flow of play (Kandori et al., 1993; Young, 1993). If a revising player chooses to volunteer even when another has already, the process experiences a low probability "birth." Similarly, if a revising player chooses not to volunteer when there is no other provider, then the process experiences a "death."

Any failures to play best replies can be interpreted as equiprobable mistakes. Here, however, a state-dependent specification, which encompasses quantal-response strategy revisions, ensures that the probability of a birth or a death can respond to payoffs.[2] For instance, if a player's cost of volunteering is low then any idiosyncratic benefit from the act of volunteering may overwhelm it; a birth is more likely. Similarly, a volunteer is less likely to die when the public good is highly prized. Birth and death probabilities also depend upon the relative noise in a player's revisions. Under the usual random-utility interpretation, a player with particularly variable payoffs will fail to play a best reply with relatively high probability.

Proposition 1 characterises long-run play when strategy revisions approximate myopic best replies. The player who volunteers in the equilibrium thus selected need not experience the lowest cost. Rather, a combination of enthusiasm (relatively high birth probability) and reliability (relatively low death probability) determines who will provide the public good. Proposition 2 reveals that when enthusiasm and reliability are more positively associated looking across the set of players, then the cost paid in the selected equilibrium is lower.

2. The evolution of voluntary action

In a simultaneous-move n-player binary-action game, player i selects $z_i \in \{0, 1\}$, where $z_i = 1$ represents "volunteering." Looking forward to the strategy-revision process described below, the pure-strategy profile z is described as a "state of play" in the state space $Z \equiv \{0, 1\}^n$.

In a volunteer's dilemma, a public good is provided if and only if at least one player undertakes the costly burden of producing it. Therefore, a player has an incentive to volunteer if and only if no other player does so. This game emerges from the payoff specification

$$u_i(z) = v_i \times \mathcal{I}\left[\sum_{i=1}^{n} z_i \geqslant 1\right] - z_i c_i, \tag{1}$$

[1] This is a common feature of related games including the textbook game of chicken ($n = 2$ here) or the classic war of attrition (Bliss and Nalebuff, 1984; Gradstein, 1992; Gradstein, 1994) in which provision is delayed until a player volunteers. In a global-game (Carlsson and van Damme, 1993) version of the asymmetric chicken game there is a unique equilibrium that approximates one of the (asymmetric) pure-strategy Nash equilibria. Similarly, under a wide variety of equilibrium-selection devices, asymmetric wars of attrition instantly end with the concession of one player (Kornhauser et al., 1989; Riley, 1999; Myatt, 2003).

[2] Quantal responses (McKelvey and Palfrey, 1995) were exploited by Blume (1995, 1997, 2003) and Blume and Durlauf (2001), who studied logit-driven evolution (one of the specifications considered here).

where $\mathcal{I}[\cdot]$ is the indicator function and where $v_i > c_i > 0$ for all i. Thus player i's private valuation for the public good is v_i, and the private cost of volunteering is c_i.[3] The pure-strategy Nash equilibria are the subset $Z_1 \subset Z$ of n states in which a single player provides.

Attention turns to evolving play. At each time t the state of play $z^t \in Z$ is updated via a one-step-at-a-time strategy-revision process: a player i is randomly selected and responds to the current play of others. This generates a Markov chain on Z. The transitions involve single steps up and down in the state space. A step up is the "birth" of a new volunteer, and is a (myopic) best reply by the revising player if $z^{t+1} \in Z_1$; that is, whenever there are no other volunteers. Otherwise, a birth is against the flow of play. Similarly, a step down is the "death" of an existing volunteer; this is against the flow of play when $z^t \in Z_1$.

If strategy revisions were myopic best replies then the process would lock in to pure-strategy equilibria. Here, however, revisions are occasionally against the flow of play: player i volunteers with some birth probability $b_i^\varepsilon > 0$ even when other providers exist; similarly, player i ceases to be the lone volunteer (or fails to volunteer when no other player is doing so) with some death probability $d_i^\varepsilon > 0$. Such "mutations" allow play to escape from Nash equilibria and move around the state space. The strategy-revision process is an ergodic Markov chain, and a unique stationary distribution reveals how often each state is played in the long run.

The birth and death probabilities are indexed by a noise parameter $\varepsilon > 0$, and satisfy $b_i^\varepsilon \to 0$ and $d_i^\varepsilon \to 0$ as $\varepsilon \to 0$. Thus, for small ε, strategy revisions approximate best replies, and most time is spent in the Nash-equilibrium states. A standard approach (Foster and Young, 1990; Kandori et al., 1993; Young, 1993) is to examine the limit of the ergodic distribution as $\varepsilon \to 0$, when it places all weight on a "stochastically stable" subset of states; when this subset is a single pure-strategy equilibrium then that equilibrium is "selected."

One possibility is $b_i^\varepsilon = d_i^\varepsilon = \varepsilon$, so that ε is a state-independent error probability. This approach is not fruitful here, since evolution treats the players symmetrically; each member of Z_1 attracts probability $1/n$ as $\varepsilon \to 0$.[4] Instead, these birth and death probabilities differ from each other and across players, and are "state dependent" in the sense of Bergin and Lipman (1996): they decline at different rates as noise is reduced. This means that the ratio of any two distinct birth or death probabilities either explodes or vanishes as $\varepsilon \to 0$.

Some formal notation is useful here. For two functions $f(\varepsilon) > 0$ and $g(\varepsilon) > 0$, write

$$f(\varepsilon) \succ g(\varepsilon) \quad \Leftrightarrow \quad \lim_{\varepsilon \to 0} \frac{f(\varepsilon)}{g(\varepsilon)} = \infty \quad \text{and} \quad f(\varepsilon) \simeq g(\varepsilon) \quad \Leftrightarrow \quad \lim_{\varepsilon \to 0} \frac{f(\varepsilon)}{g(\varepsilon)} = K > 0. \quad (2)$$

When $f(\varepsilon)$ and $g(\varepsilon)$ both vanish as $\varepsilon \to 0$, then they decline at the same rate if $f(\varepsilon) \simeq g(\varepsilon)$, and $g(\varepsilon)$ vanishes more quickly if $f(\varepsilon) \succ g(\varepsilon)$. $f(\varepsilon) \succeq g(\varepsilon)$ if $f(\varepsilon) \succ g(\varepsilon)$ or $f(\varepsilon) \simeq g(\varepsilon)$.

It is assumed that either $b_i^\varepsilon \succ b_j^\varepsilon$ or $b_i^\varepsilon \prec b_j^\varepsilon$ for all $i \neq j$. Given that this is so, it is without loss of generality to label players so that $b_1^\varepsilon \succ \cdots \succ b_n^\varepsilon$. This means that $b_1^\varepsilon > \cdots > b_n^\varepsilon$ for all ε small enough: player 1 experiences the highest birth probability and is the most enthusiastic; the remaining players are ordered by declining enthusiasm. Turning to deaths, it is assumed that $d_i^\varepsilon \succ d_j^\varepsilon$ or $d_i^\varepsilon \prec d_j^\varepsilon$ for all $i \neq j$. A player with a relatively low death probability is relatively

[3] The important feature is the game's best-reply structure. It is straightforward to restate the results for games with the same best-reply structure: for example, if costs are shared between multiple volunteers.

[4] This is also true when a state-independent process is applied to a "stag-hunt" coordination game: with two players it takes exactly one mutation to flip between the two pure-strategy equilibria. Kandori et al. (1993) circumvented this by allowing members of an n-strong population to match randomly and subsequently play a symmetric 2×2 game. In effect, the population members are players engaged in an n-player binary-action game. It then takes more than one mutation to move between equilibria.

unlikely to quit; such a player is reliable. Player r satisfying $d_r^\varepsilon \prec d_i^\varepsilon$ for all $i \neq r$ is the most reliable player. No direct association between enthusiasm and reliability is imposed. However, it is assumed that either $b_i^\varepsilon \succ d_j^\varepsilon$ or $b_i^\varepsilon \prec d_j^\varepsilon$ for all i and j.[5]

3. Strategy revisions as quantal responses

The model of state-dependent mutations considered here arises naturally when revising players choose quantal responses in the sense popularised by McKelvey and Palfrey (1995). For instance, the choice of a revising player arises from a logit quantal response when

$$\log\left[\frac{1-b_i^\varepsilon}{b_i^\varepsilon}\right] = \frac{c_i}{\varepsilon} \quad \text{and} \quad \log\left[\frac{1-d_i^\varepsilon}{d_i^\varepsilon}\right] = \frac{v_i - c_i}{\varepsilon}. \tag{3}$$

Logit responses approximate best replies when $\varepsilon \to 0$. The rates at which b_i^ε and d_i^ε vanish depend on the payoffs; for instance, $b_i^\varepsilon \succ b_j^\varepsilon$ if and only if $c_i < c_j$. Hence, ordering players by decreasing enthusiasm is equivalent to ordering them by efficiency: $c_1 < \cdots < c_n$, so that the most enthusiastic player experiences the lowest cost of provision. Similarly, player i is more reliable than player j (so that $d_i^\varepsilon \prec d_j^\varepsilon$) if and only if $v_i - c_i > v_j - c_j$. Clearly, if all players value the public good in the same way, then a more enthusiastic player is also more reliable. However, when valuations differ the most enthusiastic player may be unreliable.

The specification in (3) carries a random-utility interpretation: if the incentive of a player to volunteer is perturbed by a logistic error, then the logit is obtained. Of course, other random-utility specifications are available. Suppose, for instance, that the payoffs of a revising player are perturbed by normal noise, so that $\tilde{c}_i \sim N(c_i, \varepsilon^2 \times \xi_i^2)$ and $\tilde{v}_i \sim N(v_i, \varepsilon^2 \times \sigma_i^2)$. Writing ρ_i for the correlation coefficient between $\tilde{c}_i$ and $\tilde{v}_i$, this probit specification yields

$$b_i^\varepsilon = 1 - \Phi\left(\frac{c_i}{\varepsilon \times \xi_i}\right) \quad \text{and} \quad d_i^\varepsilon = 1 - \Phi\left(\frac{v_i - c_i}{\varepsilon \times \sqrt{\xi_i^2 + \sigma_i^2 - 2\rho_i \xi_i \sigma_i}}\right), \tag{4}$$

where $\Phi(\cdot)$ is the cumulative distribution of the standard normal. Note that $b_i^\varepsilon \succ b_j^\varepsilon$ if and only if $(c_i/\xi_i) < (c_j/\xi_i)$. The enthusiasm of player i is determined not only by the provision cost but also by the variance term ξ_i^2; players with particularly variable random utility shocks (high values of ξ_i^2 and σ_i^2) experience high birth and death probabilities.[6] Such players are more enthusiastic and yet less reliable.

Reliability is also influenced by the correlation between players' costs of provision and valuations for the public good. For instance, holding other parameters constant across the player set, $d_i^\varepsilon \prec d_j^\varepsilon$ if and only if $\rho_i > \rho_j$. Reliability is enhanced by positive correlation between cost and value shocks: any short-term increase in provision cost is tempered by a contemporaneously high valuation for the public good.

4. Long-run play

When ε is small, play almost always follows the direction of best reply and tends to "lock in" to pure-strategy Nash equilibria. These n states in Z_1 form (singleton) limit sets from which a

[5] This is for simplicity only. Admitting the possibility of ties presents no particular technical difficulties for the proofs, but the statement of the main result in Proposition 1 becomes more involved.

[6] The probit is a natural way in which to introduce heteroskedasticity to the population. Of course, it would be equally legitimate to introduce heteroskedasticity to the logit specification of (3).

noiseless (pure myopic best reply) strategy-revision process cannot escape. As ε vanishes, however, only a stochastically stable subset $Z^{\dagger} \subseteq Z_1$ retains weight in the ergodic distribution. This section characterises the stochastically stable set and identifies the players who, for vanishing noise and as part of long-run play, volunteer to provide the public good.

This stochastically stable set $Z^{\dagger}$ consists of the states in Z_1 that are hardest to leave. For $z \in Z_1$, there are two methods of escape: the active player i satisfying $z_i = 1$ is selected to revise, dies with probability d_i^{ε}, and is replaced; or some other $j \neq i$ is selected to revise, volunteers with probability b_j^{ε}, and so supplants i. For the latter escape route, the most likely replacement is the most enthusiastic of $j \neq i$. Hence, for vanishing noise, the probability of leaving z is determined by $\max[d_1^{\varepsilon}, b_2^{\varepsilon}]$ if $i = 1$ and $\max[d_i^{\varepsilon}, b_1^{\varepsilon}]$ otherwise. The active volunteers for all possible cases are identified below; the formal proof is in Appendix A.

Proposition 1. *Call player i an activist if and only if* $\lim_{\varepsilon \to 0}[\lim_{t \to \infty} \Pr[z_i^t = 1]] > 0$. *Recall that player r is the most reliable player; that is $d_r^{\varepsilon} \prec d_i^{\varepsilon}$ for all $i \neq r$.*

(i) *If $d_r^{\varepsilon} \succ b_1^{\varepsilon}$ then the unique activist is player r,*
(ii) *if $b_1^{\varepsilon} \succ d_1^{\varepsilon}$ then the unique activist is player* 1, *and*
(iii) *if $d_1^{\varepsilon} \succ b_1^{\varepsilon} \succ d_r^{\varepsilon}$ then the activists are all of the players in $M \equiv \{i\colon\ b_1^{\varepsilon} \succ d_i^{\varepsilon}\}$.*

Write $\bar{W}^{\varepsilon}$ for the expected wait until an activist volunteers. Then $\bar{W}^{\varepsilon} \preceq 1/\max[b_1^{\varepsilon}, d_r^{\varepsilon}]$.

For case (i), $d_r^{\varepsilon} \succ b_1^{\varepsilon}$ implies that deaths are always more likely than births. The quickest exit is always a death, and so the most reliable player is selected to volunteer.

For case (ii), $b_1^{\varepsilon} \succ d_1^{\varepsilon}$ implies that player 1 finds it easier to volunteer than to quit. An exit from an equilibrium state with a volunteer $j \neq 1$ must be at least as easy as the birth of player 1. Since $b_1^{\varepsilon} \succ \max[d_1^{\varepsilon}, b_2^{\varepsilon}]$, it is harder to escape from the equilibrium in which the most enthusiastic player is active. Under (homoskedastic) logit or probit specifications, case (ii) entails the efficient (lowest cost) provision of the public good.

For case (iii), there are equilibrium states from which the easiest exit is via a birth. These states involve reliable volunteers for whom any unreliability is overwhelmed by the enthusiasm of player 1. Such states are equally robust for vanishing noise and share positive weight in the limit; since $d_1^{\varepsilon} \succ b_1^{\varepsilon}$, player 1 is too unreliable to participate in the set of activists.

5. Enthusiasm and reliability

One natural configuration is when enthusiasts are reliable: $d_1^{\varepsilon} \prec \cdots \prec d_n^{\varepsilon}$. Under logit quantal-response (3) this happens when valuations are symmetric: $v_i = v$ for all i. The activist is the most enthusiastic player and the most efficient provider. On the other hand, enthusiasts might be unreliable: $d_1^{\varepsilon} \succ \cdots \succ d_n^{\varepsilon}$. Under random-utility specifications enthusiasm can stem from particularly high-variance payoff shocks; but high variances can also lessen reliability. A possibly high-cost but low-variance "plodder" may solve the volunteer's dilemma, since the relatively low-cost "star" is too unreliable to contribute consistently.

This discussion suggests that a positive association between enthusiasm and reliability favours the activism of enthusiasts. Since players are labelled via decreasing birth probability, this relationship is determined by the order of their death probabilities. To measure formally the degree to which death probabilities are ordered (and hence assess the enthusiasm-reliability association) write $d \equiv \{d_i^{\varepsilon}\}_{i=1}^{n}$ for a set of death probabilities, and $\hat{d} \equiv \{\hat{d}_i^{\varepsilon}\}_{i=1}^{n}$ for a permutation of this set:

for each i, $\hat{d}_i^\varepsilon = d_j^\varepsilon$ for some j, and vice versa. Write $d_{(y)}^\varepsilon$ for the yth smallest death probability, and note that $d_{(y)}^\varepsilon = \hat{d}_{(y)}^\varepsilon$ for all y. Now

$$O_{xy}(d) = \sum_{i=1}^{x} \mathcal{I}\left[d_i^\varepsilon \preceq d_{(y)}^\varepsilon\right] \tag{5}$$

is how many of the y most reliable players occupy one of the first x slots on the player list; it measures how well-ordered the death probabilities are. Equivalently, it is how many of the y most reliable players are also amongst the x most enthusiastic; this is the association of enthusiasm and reliability.[7] If $O_{xy}(\hat{d}) \geqslant O_{xy}(d)$ for all x and y then $\hat{d}$ is an order-improving permutation of d. Such a permutation "favours the activism of enthusiasts" if for any activist under d the number of less-enthusiastic activists is weakly reduced under $\hat{d}$.

Proposition 2. *An increase in the association of enthusiasm and reliability* (*formally, an order-improving permutation of death probabilities*) *favours the activism of enthusiasts.*

If low birth costs arise from low values of c_i, this result can be interpreted in terms of efficiency: an order-increasing permutation lowers the cost paid in the selected equilibrium.

To illustrate, adopt the logit specification (3) and consider "shifting value" from one player to another. Suppose that players i and j satisfy $c_i < c_j$ and $v_i - c_i < v_j - c_j$; hence i is more enthusiastic but less reliable than j. Let $\Delta v \equiv [v_j - v_i] - [c_j - c_i] > 0$, and note that $\Delta v < v_j$. Now move Δv utility (shifting value) from j to i so that new valuations are given by $\hat{v}_i \equiv v_i + \Delta v$ and $\hat{v}_j \equiv v_j - \Delta v$. This pairwise switch of the players' death probabilities is an order-improving permutation. By Proposition 2, the activist must be (at least weakly) more enthusiastic hence experience a weakly lower provision cost. So, if utility is transferable in this way, it is efficiency enhancing to shift value from high-cost players to low-cost players.

6. Implications

The volunteer's dilemma is a simple yet important game. The voluntary development of open-source software provides a contemporary application. Johnson (2002) models this problem as an incomplete-information volunteer's dilemma. Rather than address the equilibrium-selection problem, he instead provided a careful characterisation of the symmetric Bayesian–Nash equilibrium in which the probability that the players volunteer is decreasing in the correlation between their cost and value parameters.

The focus here is on selection between the pure-strategy Nash equilibria. It might be expected that, owing to asymmetries in payoffs, the activist would be a player with relatively low costs.

[7] To measure formally the enthusiasm–reliability association write y for the rank of a player's reliability and x for the rank of the player's enthusiasm. Thus $b_{(x)}^\varepsilon$ is the xth highest birth probability, and $d_{(y)}^\varepsilon$ is the yth lowest death probability. The joint distribution of x and y is an empirical copula:

$$C_{xy}(b,d) \equiv \sum\nolimits_{i=1}^{n} \mathcal{I}\left[b_i^\varepsilon \succeq b_{(x)}^\varepsilon\right] \times \mathcal{I}\left[d_i^\varepsilon \preceq d_{(y)}^\varepsilon\right] = \sum\nolimits_{i=1}^{x} \mathcal{I}\left[d_i^\varepsilon \preceq d_{(y)}^\varepsilon\right] = O_{xy}(d).$$

The penultimate equality follows because players are ordered via decreasing birth probability. The copula captures any association between birth and death probabilities. A measure of this association is concordance: $(\hat{b},\hat{d})$ is more concordant than (b,d) if and only if $C_{xy}(\hat{b},\hat{d}) \geqslant C_{xy}(b,d)$ for all x and y, which implies an increase in the standard empirical measures of association, such as Spearman's ρ and Kendall's τ.

However, the result is more subtle: the activist is either the most "enthusiastic" player (case (ii) in Proposition 1), or the most "reliable" player(s) (cases (i) and (iii) in Proposition 1).

Interpreting strategy revisions as quantal responses, an enthusiast has a relatively low cost parameter. Enthusiasm is associated with efficiency, and it is socially optimal for the activist to be the most enthusiastic player. If $b_i^\varepsilon \succ d_j^\varepsilon$ for all i and j then enthusiasm "overwhelms" reliability, and this is indeed true; case (ii) applies here. Alternatively, efficiency is attained when the relatively enthusiastic players are also relatively reliable (Proposition 2).

Negative association between b_i^ε and d_i^ε might be understood as negative correlation between v_i and c_i "across" the players, and efficiency is enhanced if those who value the good highly find it cheapest to provide: shifting value to lower-cost players reduces the cost paid in the selected equilibrium. Returning to the open-source software example, networking utilities might exhibit negative correlation between v_i and c_i since skilled programmers find them useful. On the other hand, word processors might exhibit positive correlation since often programmers are not end users. These observations resonate with the aforementioned comparative static of Johnson (2002), who argued that value–cost correlation provides a resolution to the "puzzle in the open source community [...] why some obviously useful software does not get written [...] while open source word processors and spreadsheets do exist, it is fair to say that only recently have they begun to be comparable in quality to, for example, Microsoft Office. On the other hand, hundreds of other free utilities and applications exist."

It is interesting to distinguish this effect from correlation "within" payoffs. Consider the specification in (4): high ρ_i ceteris paribus implies low d_i^ε and favours the overwhelming effect described above (a high draw of c_i, which might have led to a death, is likely contemporaneous to a counteractively high draw of v_i, implying improved reliability). Efficiency is favoured by negative correlation "across" the panel of players and positive correlation "within."

Finally, consider "sharing costs" when more than one player volunteers (perhaps, if k players volunteer, each player pays their cost with probability $1/k$). Since this has no effect upon the best-reply structure of the game, the same analysis applies. The difference is that the cost of player i contributing when a single other player j already provides the good drops from c_i to $c_i/2$. Under a quantal-response interpretation, this has no impact on the death probabilities, but simply raises all of the birth probabilities. This provides an alternative source for the overwhelming effect described above, and thus favours efficiency.

Another interesting alteration that might be made to the game involves requiring the contributions of many individuals to provide the public good successfully. In this case, the game becomes a threshold public-good provision game (Palfrey and Rosenthal, 1984) in which m out of n players must volunteer to provide the good. The best-reply structure is now changed; further analysis is in a companion paper (Myatt and Wallace, 2006).

Acknowledgments

The authors thank an associate editor and two referees for their valuable comments.

Appendix A

The following notation is employed: (beginning from z) $W(z)$ is the expected wait before leaving z, $Q(z, z')$ is the probability of escaping from z and reaching z' before returning to z, $N(z, z')$ is the expected number of occurrences of z before z' is reached, and $p(z) \equiv \lim_{t\to\infty} \Pr[z^t = z]$. Then

$$N(z', z) = \frac{1}{Q(z', z)} \quad \text{and so} \quad \frac{p(z)}{p(z')} = \frac{N(z, z')}{N(z', z)} = Q(z', z)N(z, z') \geqslant Q(z', z)W(z). \tag{A.6}$$

The various equalities are from Kemeny and Snell (1960, Ch. 6) and are used in the proof of Ellison (2000, Lemma 1). The inequality is obvious. All terms are functions of ε, suppressed for simplicity.

Lemma 1. $\lim_{\varepsilon\to\infty} p(z') = 0$ *for all* $z' \notin Z_1$, *and hence* $\lim_{\varepsilon\to 0}[\lim_{t\to\infty} \Pr[z^t \in Z_1]] = 1$.

Proof. There are transitions from $z' \notin Z_1$ into some $z \in Z_1$ with probability bounded away from zero as $\varepsilon \to 0$, and hence $Q(z', z)$ is bounded away from zero. As $\varepsilon \to 0$ the probability of escape from z vanishes and so $W(z) \to \infty$. Employing (A.6), $p(z)/p(z') \to \infty$ and so $p(z') \to 0$. □

For some $z \in Z_1$ take i such that $z_i = 1$. Write $W_i \equiv W(z)$ and $p_i \equiv p(z)$. Given $z' \in Z_1$ satisfying $z'_j = 1$ for $j \neq i$ write $Q_{ij} \equiv Q(z, z')$ and $N_{ij} \equiv N(z, z')$. (A.6) yields $(p_i/p_j) \geqslant Q_{ji} W_i$.

Lemma 2. *For players* $i \neq j$, $Q_{ji} \succeq \max[d_j^\varepsilon, b_i^\varepsilon]$. *For player* i, *the expected wait to escape satisfies*

$$\frac{1}{W_i} \simeq \begin{cases} \max[d_1^\varepsilon, b_2^\varepsilon] & i = 1, \\ \max[d_i^\varepsilon, b_1^\varepsilon] & i \neq 1. \end{cases}$$

Proof. Fix $z \in Z_1$ and $z' \in Z_1$ where $z_i = z'_j = 1$. From z', the birth of i occurs with probability b_i^ε/n. In the next period j dies with probability $(1 - b_j^\varepsilon)/n$. These two transitions lead to z. Hence,

$$Q_{ji} \geqslant \frac{b_i^\varepsilon}{n} \times \frac{1 - b_j^\varepsilon}{n} \simeq b_i^\varepsilon \quad \Rightarrow \quad Q_{ji} \succeq b_i^\varepsilon.$$

An alternative path is the death of j with probability d_j^ε/n followed by the birth of i with probability $(1 - d_i^\varepsilon)/n$, yielding $Q_{ji} \succeq d_j^\varepsilon$. Hence $Q_{ji} \succeq \max[d_j^\varepsilon, b_i^\varepsilon]$. Now, an escape from z involves the death of i with probability d_i^ε/n or the birth of $k \neq i$ with probability b_k^ε/n. The escape probability is

$$\frac{1}{n}\left[d_i^\varepsilon + \sum_{k\neq i} b_k^\varepsilon\right] \simeq d_i^\varepsilon + \max_{k\neq i} b_k^\varepsilon \simeq \begin{cases} \max[d_1^\varepsilon, b_2^\varepsilon] & i = 1, \\ \max[d_i^\varepsilon, b_1^\varepsilon] & i \neq 1. \end{cases}$$

The expected wait until an escape from z is the inverse of the escape probability. □

Proof of Proposition 1. The stochastically stable set is in Z_1 (Lemma 1). For case (i), $d_r^\varepsilon \succ b_1^\varepsilon$ implies that $d_i^\varepsilon \succ b_j^\varepsilon$ for all i and j. From Lemma 2, $W_i \simeq 1/d_i^\varepsilon$. Compare r and $j \neq r$:

$$\frac{p_r}{p_j} \geqslant Q_{jr} W_r \succeq \frac{\max[d_j^\varepsilon, b_r^\varepsilon]}{d_r^\varepsilon} \geqslant \frac{d_j^\varepsilon}{d_r^\varepsilon} \to \infty \quad \text{as } \varepsilon \to 0,$$

since $d_r^\varepsilon \prec d_j^\varepsilon$ for $j \neq r$. For case (ii), suppose that $b_2^\varepsilon \succ d_1^\varepsilon$. Compare p_1 to p_j for $j \neq 1$:

$$\frac{p_1}{p_j} \geqslant Q_{j1} W_1 \succeq \frac{\max[d_j^\varepsilon, b_1^\varepsilon]}{\max[d_1^\varepsilon, b_2^\varepsilon]} \simeq \frac{\max[d_j^\varepsilon, b_1^\varepsilon]}{b_2^\varepsilon} \geqslant \frac{b_1^\varepsilon}{b_2^\varepsilon} \to \infty,$$

where the limit follows from $b_1^\varepsilon \succ b_2^\varepsilon$. Hence $p_j \to 0$ for $j \neq 1$. If instead $b_2^\varepsilon \prec d_1^\varepsilon$, then

$$\frac{p_1}{p_j} \geqslant Q_{j1} W_1 \succeq \frac{\max[d_j^\varepsilon, b_1^\varepsilon]}{\max[d_1^\varepsilon, b_2^\varepsilon]} \simeq \frac{\max[d_j^\varepsilon, b_1^\varepsilon]}{d_1^\varepsilon} \geqslant \frac{b_1^\varepsilon}{d_1^\varepsilon} \to \infty,$$

since $b_1^\varepsilon \succ d_1^\varepsilon$. Once again, $p_j \to 0$ for $j \neq 1$, and hence $p_1 \to 1$. For case (iii), $d_1^\varepsilon \succ b_1^\varepsilon \succ d_r^\varepsilon$ and $M \equiv \{i\colon b_1^\varepsilon \succ d_i^\varepsilon\}$. Hence $1 \notin M$. First compare $i \in M$ and $j \notin M$:

$$\frac{p_i}{p_j} \geqslant Q_{ji} W_i \succeq \frac{\max[d_j^\varepsilon, b_i^\varepsilon]}{\max[d_i^\varepsilon, b_1^\varepsilon]} \simeq \frac{\max[d_j^\varepsilon, b_r^\varepsilon]}{b_1^\varepsilon} \geqslant \frac{d_j^\varepsilon}{b_1^\varepsilon} \to \infty,$$

where the limit follows because $j \notin M$ and hence $d_j^\varepsilon \succ b_1^\varepsilon$. This ensures that $p_j \to 0$ for any $j \notin M$. Now, however, compare $i \in M$ with $j \in M$, noting that $i \neq 1$ and $j \neq 1$. Applying Lemma 2,

$$\frac{p_i}{p_j} = \frac{p_i}{p_1} \times \frac{p_1}{p_j} \geqslant Q_{1i} W_i \times Q_{j1} W_1 \succeq \frac{\max[d_1^\varepsilon, b_i^\varepsilon]}{\max[d_i^\varepsilon, b_1^\varepsilon]} \times \frac{\max[d_j^\varepsilon, b_1^\varepsilon]}{\max[d_1^\varepsilon, b_2^\varepsilon]} \simeq \frac{d_1^\varepsilon}{b_1^\varepsilon} \times \frac{b_1^\varepsilon}{d_1^\varepsilon} = 1.$$

The "$\simeq$" step follows from the conditions of case (iii) and the definition of the set M. This ensures that (p_i/p_j) is bounded away from zero as $\varepsilon \to 0$. A symmetric argument ensures that (p_j/p_i) is bounded away from zero, and hence $\lim_{\varepsilon\to 0} p_i \in (0,1)$ for all $i \in M$.

Finally, consider the expected wait $\bar{W}^\varepsilon$ until $Z^\dagger$ is reached. Beginning from outside Z_1, the expected wait until Z_1 is reached is bounded above as $\varepsilon \to 0$. Suppose that the process begins from $z \in Z_1$ where $z \notin Z^\dagger$ and $z_j = 1$ for some j. In case (i) the death of j followed by the birth of r has probability of at least $d_j^\varepsilon(1-d_r^\varepsilon)/n^2 \simeq d_j^\varepsilon \succ d_r^\varepsilon \succeq \max[d_r^\varepsilon, b_1^\varepsilon]$, and hence $\bar{W}^\varepsilon \prec 1/\max[d_r^\varepsilon, b_1^\varepsilon]$. For case (ii) the death of j followed by the birth of 1 has probability of at least $d_j^\varepsilon(1-d_1^\varepsilon)/n^2 \simeq d_j^\varepsilon \succ d_r^\varepsilon$ and so $\bar{W}^\varepsilon \preceq 1/d_r^\varepsilon$. Alternatively, the birth of 1 followed by the death of j leads to $\bar{W}^\varepsilon \prec 1/b_1^\varepsilon$. Combining these two possibilities, $\bar{W}^\varepsilon \preceq 1/\max[d_r^\varepsilon, b_1^\varepsilon]$. Finally, for case (iii), the same logic leads to $\bar{W}^\varepsilon \preceq 1/d_j^\varepsilon$ for $j \notin M$. However, $d_j^\varepsilon \succ d_r^\varepsilon$ and $d_j^\varepsilon \succ b_1^\varepsilon$ for $j \notin M$, and hence $\bar{W}^\varepsilon \prec 1/\max[d_r^\varepsilon, b_1^\varepsilon]$. □

Given the death probabilities $d \equiv \{d_i^\varepsilon\}_{i=1}^n$, the most reliable player r satisfies $d_i^\varepsilon \succ d_r^\varepsilon$ for $i \neq r$. For a new permutation $\hat{d} \equiv \{\hat{d}_i^\varepsilon\}_{i=1}^n$ write $\hat{r}$ for the most reliable player, and note that $d_r^\varepsilon = d_{\hat{r}}^\varepsilon$.

Proof of Proposition 2. If $d_r^\varepsilon \succ b_1^\varepsilon$ then (Proposition 1) the unique activist is r. Player r is the most reliable and so $O_{r1}(d) = 1$. Under the permutation $\hat{d}$, $d_{\hat{r}}^\varepsilon = d_r^\varepsilon \succ b_1^\varepsilon = \hat{b}_1^\varepsilon$, and so the unique activist is $\hat{r}$. If $r < \hat{r}$ then $O_{r1}(\hat{d}) = 0 < O_{r1}(d)$ contradicts the assumption that $\hat{d}$ is an order-improving permutation of d. Hence $\hat{r} \leqslant r$, and the unique activist is (weakly) more enthusiastic.

If $b_1^\varepsilon \succ d_1^\varepsilon$ then (Proposition 1) player 1 is the unique activist. Suppose that player 1 is the yth most reliable, so that $O_{1y}(d) = 1$. $\hat{d}$ is an order-improving permutation of d and so $O_{1y}(\hat{d}) \geqslant 1$. Hence in this new permutation player 1 must be one of the y most reliable players. Hence $\hat{d}_1^\varepsilon \preceq d_1^\varepsilon \prec b_1^\varepsilon = \hat{b}_1^\varepsilon$, and player 1 remains the unique activist; trivially, the activist is weakly more enthusiastic.

If $d_1^\varepsilon \succ b_1^\varepsilon \succ d_r^\varepsilon$ then (Proposition 1) all members of $M \equiv \{i\colon b_1^\varepsilon \succ d_i^\varepsilon\}$ are activists. Note that $1 \notin M$. Under the new permutation $\hat{d}_{\hat{r}}^\varepsilon = d_r^\varepsilon \prec b_1^\varepsilon = \hat{b}_1^\varepsilon$ and hence case (i) of Proposition 1 cannot apply. If case (ii) applies, then the most enthusiastic player is the new unique activist, and the claim of the proposition holds. The remaining situation is when case (iii) continues to apply. The new activists are members of $\hat{M} \equiv \{i \mid b_1^\varepsilon \succ \hat{d}_i^\varepsilon\}$; note that M and $\hat{M}$ have the same cardinality y. They consist of the y most reliable players; hence $M = \{i\colon d_i^\varepsilon \preceq d_{(y)}^\varepsilon\}$ and $\hat{M} = \{i\colon \hat{d}_i^\varepsilon \preceq d_{(y)}^\varepsilon\}$.

The proposition's claim holds if $\sum_{i=j}^n \mathcal{I}[i \in M] \geqslant \sum_{i=j}^n \mathcal{I}[i \in \hat{M}]$ for all j. Since M and $\hat{M}$ have the same cardinality, this inequality automatically holds (as an equality) for $j = 1$. For $j > 1$, it holds if and only if $\sum_{i=1}^{j-1} \mathcal{I}[i \in M] \leqslant \sum_{i=1}^{j-1} \mathcal{I}[i \in \hat{M}]$, or equivalently

$$\sum_{i=1}^{j-1} \mathcal{I}\big[d_i^\varepsilon \preceq d_{(y)}^\varepsilon\big] \leqslant \sum_{i=1}^{j-1} \mathcal{I}\big[\hat{d}_i^\varepsilon \preceq d_{(y)}^\varepsilon\big].$$

This last inequality is equivalent to $O_{xy}(d) \leqslant O_{xy}(\hat{d})$ for $x = j - 1$. This holds by assumption. $\square$

References

Bergin, J., Lipman, B., 1996. Evolution with state-dependent mutations. Econometrica 64, 943–956.

Bliss, C., Nalebuff, B., 1984. Dragon slaying and ballroom dancing: The private supply of a public good. J. Public Econ. 25, 1–12.

Blume, L.E., 1995. The statistical mechanics of best-response strategy revision. Games Econ. Behav. 11, 111–145.

Blume, L.E., 1997. Population games. In: Arthur, W.B., Durlauf, S.N., Lane, D.A. (Eds.), The Economy as an Evolving Complex System II. Westview Press, Boulder, CO.

Blume, L.E., 2003. How noise matters. Games Econ. Behav. 44, 251–271.

Blume, L.E., Durlauf, S.N., 2001. The interactions-based approach to socioeconomic behaviour. In: Durlauf, S.N., Young, H.P. (Eds.), Social Dynamics. MIT Press, Cambridge, MA.

Carlsson, H., van Damme, E., 1993. Global games and equilibrium selection. Econometrica 61, 989–1018.

Diekmann, A., 1985. Volunteer's dilemma. J. Conflict Resolution 29, 605–610.

Ellison, G., 2000. Basins of attraction, long run stochastic stability and the speed of step-by-step evolution. Rev. Econ. Stud. 67, 17–45.

Foster, D., Young, H.P., 1990. Stochastic evolutionary game dynamics. Theoret. Pop. Bio. 38, 219–232.

Gradstein, M., 1992. Time dynamics and incomplete information in the private provision of public goods. J. Polit. Economy 100, 581–597.

Gradstein, M., 1994. Efficient provision of a discrete public good. Int. Econ. Rev. 35, 877–897.

Johnson, J.P., 2002. Open source software: Private provision of a public good. J. Econ. Manage. Strategy 11, 637–662.

Kandori, M., Mailath, G.J., Rob, R., 1993. Learning, mutation and long-run equilibria in games. Econometrica 61, 29–56.

Kemeny, J.G., Snell, J.L., 1960. Finite Markov Chains. Van Nostrand, Princeton, NJ.

Kornhauser, L., Rubinstein, A., Wilson, C., 1989. Reputation and patience in the war of attrition. Economica 56, 15–24.

McKelvey, R.D., Palfrey, T.R., 1995. Quantal response equilibria for normal form games. Games Econ. Behav. 10, 6–38.

Myatt, D.P., 2003. Instant exit from the asymmetric war of attrition. Department of Economics Discussion Paper 160, Oxford University.

Myatt, D.P., Wallace, C., 2006. When does one bad apple spoil the barrel? An evolutionary analysis of collective action. Department of Economics Discussion Paper 269, Oxford University.

Palfrey, T.R., Rosenthal, H., 1984. Participation and the provision of discrete public goods: A strategic analysis. J. Public Econ. 24, 171–193.

Riley, J.G., 1999. Asymmetric contests: A resolution of the tullock paradox. In: Howitt, P., De Antoni, E., Leijonhufvud, A. (Eds.), Money, Markets and Method: Essays in Honour of Robert W. Clower. Edward Elgar, Cheltenham, UK.

Weesie, J., 1994. Incomplete information and timing in the volunteer's dilemma: A comparison of four models. J. Conflict Resolution 38, 557–585.

Young, H.P., 1993. The evolution of conventions. Econometrica 61, 57–84.