This page is intentionally blank. Proper e-companion title page, with INFORMS branding and exact metadata of the main paper, will be produced by the INFORMS office when the issue is being assembled.

Online Supplement

Appendix A: Continuity of $\hat{u}$ and H

This section expands on the discussion of footnote 15 regarding the continuity of $\hat{u}$ (Section 3.1) and H (equation 3). Generally, these functions are not jointly continuous at points where one or more $\sigma_{sia} = 0$ and $(1-t)\eta = 0$, because it is not ensured that

$$\lim_{(1-t)\eta \searrow 0} (1-t)\eta \log(\sigma_{sia}) = 0$$

As an example, consider the sequence $\big((1-t)\eta, \sigma_{sia}\big)_n = \left(\frac{1}{n}, \frac{1}{e^n}\right)$, for which the above expression is constant and does not approach 0 as would be required for continuity. Intuitively, the example requires σ to decrease much faster than $(1-t)\eta$. However, the following will show this can not occur in the equilibrium sets (i.e. the zero set of H) used by the algorithm. The reason is that in equilibrium, σ_{sia} is bounded from below by a function of $(1-t)\eta$ which ensures convergence to 0. Thus, one could construct a domain that includes all equilibria while also ensuring joint continuity of $\hat{u}$ and H .

The following argument establishes existence of said bound. Clearly, only sequences where $\sigma_{sia} \to 0$ are potentially problematic. For each such action a, there must be another action $a' \in A_{si}$ of the same agent (s,i) that converges to a number strictly greater than 0. Consider then the equation

$$\frac{H^V_{sia'}(\boldsymbol{\sigma}, \boldsymbol{V}, t)}{\sigma_{sia'}} - \frac{H^V_{sia}(\boldsymbol{\sigma}, \boldsymbol{V}, t)}{\sigma_{sia}} = \bar{U}^t_{si}(a', \boldsymbol{\sigma}_{s,\text{-}i}, \boldsymbol{V}_i) - \bar{U}^t_{si}(a, \boldsymbol{\sigma}_{s,\text{-}i}, \boldsymbol{V}_i) + (1-t)\eta \left(\frac{\nu_{sia'}}{\sigma_{sia'}} - \frac{\nu_{sia}}{\sigma_{sia}} \right) = 0$$

derived from equation (3b) which must hold for all points in $H^{-1}(0)$. Rearranging terms yields

$$\sigma_{sia} = \frac{\nu_{sia}}{\frac{\bar{U}^t_{si}(a', \boldsymbol{\sigma}_{s,\text{-}i}, \boldsymbol{V}_i) - \bar{U}^t_{si}(a, \boldsymbol{\sigma}_{s,\text{-}i}, \boldsymbol{V}_i)}{(1-t)\eta} + \frac{\nu_{sia'}}{\sigma_{sia'}}}$$

Here, ν_{sia} and $\nu_{sia'}$ are positive constants. By assumption, $\sigma_{sia'} \not\to 0$ so that $\frac{\nu_{sia'}}{\sigma_{sia'}}$ is bounded. Because V_{si} are bounded in equilibrium (see Proposition 2), $\bar{U}^t_{si}$ are also bounded for all $\boldsymbol{\sigma}, \boldsymbol{V}, t$ in equilibrium. In addition, the fact that σ_{sia} goes to 0, while $\sigma_{sia'}$ does not, implies $\bar{U}^t_{si}(a', \boldsymbol{\sigma}_{s,\text{-}i}, \boldsymbol{V}_i) - \bar{U}^t_{si}(a, \boldsymbol{\sigma}_{s,\text{-}i}, \boldsymbol{V}_i) \geq 0$ for $(1-t)\eta$ sufficiently small (action a' must be at least as good as action a in equilibrium; compare proofs of Propositions 7 and 10). Therefore, for small $(1-t)\eta$ it holds that $M \geq \bar{U}^t_{si}(a', \boldsymbol{\sigma}_{s,\text{-}i}, \boldsymbol{V}_i) - \bar{U}^t_{si}(a, \boldsymbol{\sigma}_{s,\text{-}i}, \boldsymbol{V}_i) \geq 0$ for some $M > 0$.

Using all of the preceding, one can bound σ_{sia} by two expressions of the form $\frac{1}{\frac{q}{(1-t)\eta}+r}$ with $q \geq 0, r > 0$. Because

$$\lim_{(1-t)\eta \searrow 0} (1-t)\eta \log\left(\frac{1}{\frac{q}{(1-t)\eta}+r}\right) = 0$$

for all $q \geq 0, r > 0$, the sandwich theorem then gives

$$\lim_{(1-t)\eta \searrow 0} (1-t)\eta \log(\sigma_{sia}) = 0$$

as claimed.

Appendix B: Derivation of the Homotopy Function H

This section provides the proof of Proposition 1, which states that the zero set of $H(\boldsymbol{\sigma}, \boldsymbol{V}, t)$ coincides with the set of equilibria of the auxiliary stochastic games $\mathcal{G}^t$. To this end, we derive the homotopy function H from the maximization problems stated in equation (2). $H = 0$ corresponds to the problems' first order conditions, which are not only necessary, but due to concavity also sufficient for an equilibrium.

The Lagrangeans corresponding to the maximization problems stated in equation (2) are given by

$$\begin{aligned}\mathcal{L}^t_{si} \;=\; V^t_{si} \;+\; \alpha_{si}\Bigg[-V^t_{si} \;+\; \bar{u}^t_{si}(\boldsymbol{\sigma}_s) \;+\; \delta_i \sum_{s' \in S} \bar{\phi}^t_{s \to s'}(\boldsymbol{\sigma}_s)\, V^t_{s'i} \\ + (1-t)\eta \sum_{a \in A_{si}} \nu_{sia} \log(\sigma_{sia}) \Bigg] \;+\; \beta_{si}\Bigg[\sum_{a \in A_{si}} \sigma_{sia} - 1\Bigg]\end{aligned}$$

where $\alpha_{si}, \beta_{si} \neq 0$ are the Lagrange multipliers of the two constraints (2b) and (2c). Since the logarithmic penalty terms are strictly concave in $\boldsymbol{\sigma}_{si}$ and all other terms are linear in $\boldsymbol{\sigma}_{si}$, the Karush-Kuhn-Tucker conditions are both necessary and sufficient. For each agent $(s,i) \in S \times I$, they consist of the two constraints, as well as one equation for each action $a \in A_{si}$:

$$\frac{\partial \mathcal{L}^t_{si}}{\partial \sigma_{sia}} = \alpha_{si}\left[\bar{u}^t_{si}(a, \boldsymbol{\sigma}_{si}) + \delta_i \sum_{s' \in S} \bar{\phi}^t_{s \to s'}(a, \boldsymbol{\sigma}_{s,\text{-}i})\, V^t_{s'i} + (1-t)\frac{\eta \nu_{sia}}{\sigma_{sia}}\right] + \beta_{si} = 0$$

Multiplying each of these equations by the corresponding σ_{sia} and summing up over $a \in A_{si}$ yields

$$\alpha_{si} \underbrace{\left[\bar{u}^t_{si}(\boldsymbol{\sigma}_{si}, \boldsymbol{\sigma}_{s,\text{-}i}) + \delta_i \sum_{s' \in S} \bar{\phi}^t_{s \to s'}(\boldsymbol{\sigma}_{si}, \boldsymbol{\sigma}_{s,\text{-}i})\, V^t_{s'i} + (1-t)\eta \sum_{a \in A_{si}} \nu_{sia}\right]}_{= V^t_{si} - (1-t)\eta \sum_{a \in A_{si}} \nu_{sia}\left[\log(\sigma_{sia}) - 1\right]} + \beta_{si} \underbrace{\sum_{a \in A_{si}} \sigma_{sia}}_{=1} = 0$$

and thus

$$\frac{\beta_{si}}{\alpha_{si}} = -V^t_{si} + (1-t)\eta \sum_{a\in A_{si}} \nu_{sia}\big[\log(\sigma_{sia}) - 1\big]$$

Replacing $\frac{\beta_{si}}{\alpha_{si}}$ in the first order conditions $\frac{\partial \mathcal{L}^t_{si}}{\partial \sigma_{sia}} = 0$ gives the following necessary and sufficient conditions for all agents $(s,i)\in S\times I$:

$$\begin{aligned} 0 &= -V^t_{si} + \bar{u}^t_{si}(a, \boldsymbol{\sigma}_{s,\text{-}i}) + \delta_i \sum_{s'\in S} \bar{\phi}^t_{s\to s'}(a, \boldsymbol{\sigma}_{s,\text{-}i})\, V^t_{s'i} \\ &\qquad + (1-t)\eta\left(\frac{\nu_{sia}}{\sigma_{sia}} + \sum_{a'\in A_{si}} \nu_{sia'}\big[\log(\sigma_{sia'}) - 1\big]\right) \qquad \forall\ a\in A_{si} \\ 0 &= \sum_{a\in A_{si}} \sigma_{sia} - 1 \end{aligned}$$

which characterize the set of stationary equilibria of the game $\mathcal{G}^t$.

The homotopy function H is obtained by multiplying the right hand sides of the former set of equations by the corresponding σ_{sia}, and then collecting all conditions for all agents.

Appendix C: The Starting Point (σ^0, V^0)

C.1. Existence and Uniqueness

This section provides the proof of Proposition 3, which states that the auxiliary stochastic game $\mathcal{G}^0$ always has a unique equilibrium $(\boldsymbol{\sigma}^0, \boldsymbol{V}^0)$. At $t=0$, players maximize solely against their prior $\boldsymbol{\rho}_{\text{-}i}$, so that strategic interaction is completely absent. Formally, each player faces a discounted Markov decision problem with finite state space S, action spaces $\Delta(A_{si})$, and state transition functions $\bar{\phi}^0_{s\to s'}(\boldsymbol{\sigma}_{si}) = \phi_{s\to s'}(\boldsymbol{\sigma}_{si}, \boldsymbol{\rho}_{s,\text{-}i})$. Instantaneous utilities are given by:

$$\begin{aligned} \hat{u}^0_{si} : \Delta(A_{si}) &\quad\rightarrow\quad \{-\infty\}\cup\mathbb{R} \\ \boldsymbol{\sigma}_{si} &\quad\mapsto\quad u_{si}(\boldsymbol{\sigma}_{si}, \boldsymbol{\rho}_{s,\text{-}i}) + \eta \sum_{a\in A_{si}} \nu_{sia}\log(\sigma_{sia}) \end{aligned}$$

For the purpose of proving existence, it is helpful to set $\log(0) := -\infty$ and take the extended real line $\{-\infty\}\cup\mathbb{R}$ as range for $\hat{u}$, as indicated above. Note that still $\hat{u} < \infty$, so that total expected discounted utilities are always well-defined. Denote as $\hat{U}^0_{si}(\boldsymbol{\sigma}_i)$ total utility under $\boldsymbol{\sigma}_i$ and beginning in state s, so that one obtains in vector notation:

$$\hat{\boldsymbol{U}}_i^0(\boldsymbol{\sigma}_i) = \left(\hat{U}_{1i}^0(\boldsymbol{\sigma}_i), \hat{U}_{2i}^0(\boldsymbol{\sigma}_i), \dots\right)^\top = \left(I - \delta_i \Phi(\boldsymbol{\sigma}_i)\right)^{-1} \hat{\boldsymbol{u}}_i^0(\boldsymbol{\sigma}_i)$$

Any solution to the Markov decision problem of player i must then satisfy

$$V_{si} = \max_{\boldsymbol{\sigma}_i \in \times_{s\in S} \Delta(A_{si})} \hat{U}_{si}^0(\boldsymbol{\sigma}_i) \qquad \forall s \in S$$

Because $\hat{U}_{si}^0 : \times_{s\in S} \Delta(A_{si}) \to \{-\infty\} \cup \mathbb{R}$ is upper semi-continuous over a compact domain, the respective version of the extreme value theorem guarantees that this maximum exists (Bourbaki 1966, Ch. IV, § 6, Theorem 3, p. 361). Moreover it must be that $V_{si} > -\infty$, because an arbitrary interior strategy guarantees some finite total discounted utility in every state, giving a lower bound for the maximized values. Thus there exists a unique vector $\boldsymbol{V}_i^0 \in \mathbb{R}^{|S|}$ of state values for each player. Given this vector, optimal strategies $\boldsymbol{\sigma}^0$ are the solutions to

$$\boldsymbol{\sigma}_{si}^0 = \operatorname*{arg\,max}_{\boldsymbol{\sigma}_{si} \in \Delta(A_{si})} \left\{ u_{si}(\boldsymbol{\sigma}_{si}, \boldsymbol{\rho}_{s,\text{-}i}) + \delta_i \sum_{s'\in S} \phi_{s \to s'}(\boldsymbol{\sigma}_{si}, \boldsymbol{\rho}_{s,\text{-}i})\, V_{s'i}^0 + \eta \sum_{a\in A_{si}} \nu_{sia} \log(\sigma_{sia}) \right\}$$

Because the first two terms are linear in $\boldsymbol{\sigma}_{si}$, while the logarithmic term is strictly concave, optimal strategies are also unique.

C.2. Computation

To compute initial values $\boldsymbol{V}^0$ and strategies $\boldsymbol{\sigma}^0$, one can use standard value function iteration on the Bellman equation

$$\begin{aligned} V_{si}^{(k+1)} &= \max_{\boldsymbol{\sigma}_{si}} \left[u_{si}(\boldsymbol{\sigma}_{si}, \boldsymbol{\rho}_{s,\text{-}i}) + \delta_i \sum_{s'\in S} \phi_{s \to s'}(\boldsymbol{\sigma}_{si}, \boldsymbol{\rho}_{s,\text{-}i})\, V_{s'i}^{(k)} + \eta \sum_{a\in A_{si}} \nu_{sia} \log(\sigma_{sia}) \right] \\ &=: \max_{\boldsymbol{\sigma}_{si}} \left[U_{si}^k(\boldsymbol{\sigma}_{si}) + \eta \sum_{a\in A_{si}} \nu_{sia} \log(\sigma_{sia}) \right] \end{aligned}$$

where k counts iterations and U^k is simply a shorthand for the first two terms: The linear part of instantaneous utility plus expected discounted future value under the current estimate $V^{(k)}$. Introducing a multiplier γ_{si} for the constraint $\sum_a \sigma_{sia} = 1$, necessary and sufficient first order conditions for each (s, i) are then

$$\frac{\partial U_{si}^k(\boldsymbol{\sigma}_{si})}{\partial \sigma_{sia}} + \frac{\eta \nu_{sia}}{\sigma_{sia}} = U_{sia}^k + \frac{\eta \nu_{sia}}{\sigma_{sia}} = \gamma_{si} \qquad \forall a \in A_{si}$$

$$\sum_{a\in A_{si}} \sigma_{sia} = 1$$

Dropping indices s, i, and k, and labeling the strategies $\sigma_1, ..., \sigma_N$ with $N = |A_{si}|$, the above implies

$$U_1 + \frac{\eta\nu_1}{\sigma_1} = U_n + \frac{\eta\nu_n}{\sigma_n} \qquad \text{for } n = 1, \dots, N$$

and thus

$$\sigma_n = \frac{\nu_n}{\frac{U_1 - U_n}{\eta} + \frac{\nu_1}{\sigma_1}}$$

Plugging into the constraint gives an equation only in σ_1:

$$f(\sigma_1) := \sum_{n=1}^{N} \frac{\nu_n}{\frac{U_1 - U_n}{\eta} + \frac{\nu_1}{\sigma_1}} - 1 = 0$$

If $N = 1$, then $\sigma_1 = 1$. If $N > 1$, order strategies such that $U_1 - U_n \geq 0$ for all n, without loss of generality. The following shows that the equation then always has a unique solution $\sigma_1 \in (0, 1)$. First, f is continuous and monotonously increasing on the open unit interval:

$$f'(\sigma_1) = \sum_{n=1}^{N} \frac{\eta^2 \nu_1 \nu_n}{\left[(U_1 - U_n)\sigma_1 + \eta\nu_1\right]^2} > 0$$

Secondly, behavior at the boundaries of $(0, 1)$ is given by

$$\lim_{\sigma_1 \to 0} f(\sigma_1) = -1$$
$$\lim_{\sigma_1 \to 1} f(\sigma_1) = \sum_{n=1}^{N} \frac{\nu_n}{\frac{U_1 - U_n}{\eta} + \nu_1} - 1 = \sum_{n=2}^{N} \frac{\nu_n}{\frac{U_1 - U_n}{\eta} + \nu_1} > 0$$

By application of the intermediate value theorem, there exists a unique solution for σ_1 in the unit interval $(0, 1)$, which can be found by standard root-finding algorithms. All other σ_n are then also uniquely determined by the first order conditions. Plugging the optimal policy into the Bellman equation yields the next value iterate $\boldsymbol{V}^{(k+1)}$. Starting from an arbitrary $\boldsymbol{V}^{(0)}$, e.g. the zero vector, this process can be repeated until $\boldsymbol{V}$ has converged.

Appendix D: Jacobian of H

The Jacobian matrix of $H(\boldsymbol{\sigma}, \boldsymbol{V}, \boldsymbol{\nu}, \eta, t)$ is

$J : [0, 1]^{|A|} \times \mathbb{R}^{|S \times I|} \times \mathbb{R}_+^{|A|} \times [0, \infty) \times [0, 1] \to \mathbb{R}^{|A| + |S \times I|} \times \mathbb{R}^{|A| + |S \times I| + |A| + 2}$,

$$J(\boldsymbol{\sigma}, \boldsymbol{V}, \boldsymbol{\nu}, \eta, t) \quad = \quad \begin{pmatrix} \frac{\partial H^V_{sia}}{\partial \sigma_{s'i'a'}} & \frac{\partial H^V_{sia}}{\partial V_{s'i'}} & \frac{\partial H^V_{sia}}{\partial \nu_{s'i'a'}} & \frac{\partial H^V_{sia}}{\partial \eta} & \frac{\partial H^V_{sia}}{\partial t} \\ \frac{\partial H^\sigma_{si}}{\partial \sigma_{s'i'a'}} & \frac{\partial H^\sigma_{si}}{\partial V_{s'i'}} & \frac{\partial H^\sigma_{si}}{\partial \nu_{s'i'a'}} & \frac{\partial H^\sigma_{si}}{\partial \eta} & \frac{\partial H^\sigma_{si}}{\partial t} \end{pmatrix}$$

with

$$\frac{\partial H^V_{sia}(\boldsymbol{\sigma}, \boldsymbol{V}, \boldsymbol{\nu}, \eta, t)}{\partial \sigma_{s'i'a'}} = \begin{cases} -V_{si} + \bar{u}^t_{si}(a, \boldsymbol{\sigma}_{s,\text{-}i}) + \delta_i \sum_{s'' \in S} \bar{\phi}^t_{s \to s''}(a, \boldsymbol{\sigma}_{s,\text{-}i})\, V_{s''i} & \text{if } s' = s, i' = i \\ \qquad + (1-t)\eta\Big(\nu_{sia} + \sum_{a'' \in A_{si}} \nu_{sia''}\big[\log(\sigma_{sia''}) - 1\big]\Big) & \text{and } a' = a \\ (1-t)\eta\nu_{sia'}\dfrac{\sigma_{sia}}{\sigma_{sia'}} & \text{if } s' = s, i' = i \\ & \text{and } a' \neq a \\ t\sigma_{sia}\Big[u_{si}(a_{si}, a_{s,i'}, \boldsymbol{\sigma}_{s,\text{-}\{i,i'\}}) & \text{if } s' = s \\ \qquad + \delta_i \sum_{s'' \in S} \bar{\phi}^t_{s \to s''}(a_{si}, a_{s,i'}, \boldsymbol{\sigma}_{s,\text{-}\{i,i'\}})\, V_{s'',i}\Big] & \text{and } i' \neq i \\ 0 & \text{else} \end{cases}$$

$$\frac{\partial H^V_{sia}(\boldsymbol{\sigma}, \boldsymbol{V}, \boldsymbol{\nu}, \eta, t)}{\partial V_{s'i'}} = \begin{cases} \sigma_{sia}\Big(\delta_i \bar{\phi}^t_{s \to s'}(a, \boldsymbol{\sigma}_{s,\text{-}i}) - 1\Big) & \text{if } i' = i \text{ and } s' = s \\ \sigma_{sia}\, \delta_i \bar{\phi}^t_{s \to s'}(a, \boldsymbol{\sigma}_{s,\text{-}i}) & \text{if } i' = i \text{ and } s' \neq s \\ 0 & \text{if } i' \neq i \end{cases}$$

$$\frac{\partial H^V_{sia}(\boldsymbol{\sigma}, \boldsymbol{V}, \boldsymbol{\nu}, \eta, t)}{\partial \nu_{s'i'a'}} = \begin{cases} (1-t)\eta\Big(1 + \sigma_{sia}\big[\log(\sigma_{sia}) - 1\big]\Big) & \text{if } s' = s, i' = i \text{ and } a' = a \\ (1-t)\eta\sigma_{sia}\big[\log(\sigma_{sia}) - 1\big] & \text{if } s' = s, i' = i \text{ and } a' \neq a \\ 0 & \text{if } s' \neq s \text{ or } i' \neq i \end{cases}$$

$$\frac{\partial H^V_{sia}(\boldsymbol{\sigma}, \boldsymbol{V}, \boldsymbol{\nu}, \eta, t)}{\partial \eta} = (1-t)\Big(\nu_{sia} + \sigma_{sia} \sum_{a' \in A_{si}} \nu_{sia'}\big[\log(\sigma_{sia'}) - 1\big]\Big)$$

$$\begin{aligned} \frac{\partial H^V_{sia}(\boldsymbol{\sigma}, \boldsymbol{V}, \boldsymbol{\nu}, \eta, t)}{\partial t} = \sigma_{sia}\Big(& u_{si}(a, \boldsymbol{\sigma}_{s,\text{-}i}) - u_{si}(a, \boldsymbol{\rho}_{s,\text{-}i}) \\ & + \delta_i \sum_{s' \in S} \big[\phi_{s \to s'}(a, \boldsymbol{\sigma}_{s,\text{-}i}) - \phi_{s \to s'}(a, \boldsymbol{\rho}_{s,\text{-}i})\big] V_{s'i}\Big) \\ & - \eta\Big(\nu_{sia} + \sigma_{sia} \sum_{a' \in A_{si}} \nu_{sia'}\big[\log(\sigma_{sia'}) - 1\big]\Big) \end{aligned}$$

$$\frac{\partial H^{\sigma}_{si}(\boldsymbol{\sigma},\boldsymbol{V},\boldsymbol{\nu},\eta,t)}{\partial \sigma_{s'i'a'}} = \begin{cases} 1 & \text{if } s'=s \text{ and } i'=i \\ 0 & \text{else} \end{cases}$$

$$\frac{\partial H^{\sigma}_{si}(\boldsymbol{\sigma},\boldsymbol{V},\boldsymbol{\nu},\eta,t)}{\partial V_{s'i'}} = \frac{\partial H^{\sigma}_{si}(\boldsymbol{\sigma},\boldsymbol{V},\boldsymbol{\nu},\eta,t)}{\partial \nu_{s'i'a'}} = \frac{\partial H^{\sigma}_{si}(\boldsymbol{\sigma},\boldsymbol{V},\boldsymbol{\nu},\eta,t)}{\partial \eta} = \frac{\partial H^{\sigma}_{si}(\boldsymbol{\sigma},\boldsymbol{V},\boldsymbol{\nu},\eta,t)}{\partial t} = 0$$

Appendix E: Regularity: Parametrized Sard's Theorem, Full Rank of Jacobian

This section provides the proofs for Proposition 4 and Lemma 8.1. To that end, we will show that a generalization of Sard's theorem, known as parametrized Sard's theorem, applies to H. This proves that for generic $\boldsymbol{\nu}$, $\mathbf{0}$ is a regular value of H.

To give an intuitive idea of this result, suppose one picked $\boldsymbol{\nu}$ such that the solution set to $H|_{\boldsymbol{\nu}} = \mathbf{0}$ contained singularities; then these singularities must be unstable and disappear with probability 1 when using a slightly perturbed vector $\boldsymbol{\nu}+\boldsymbol{\epsilon}$ instead. In consequence, it is sufficient to pick an appropriately randomized $\boldsymbol{\nu}$ to ensure regularity. (For the purpose of numerical computation of equilibria, this may not even be necessary. In our experience, simply setting all ν_{sia} to 1 poses no problem for numerical continuation. The only singularities that we then encountered in extensive testing were transversal bifurcations. These are unproblematic from a numerical perspective, as the path can simply be continued across such points. We failed to create genuinely problematic singularities such as higher-dimensional subsets contained in the solution set.)

We now turn to proving regularity. Parametrized Sard's theorem (Chow et al. 1978, Theorem 2.1, p. 891) reads:

Let $\mathcal{Y} \subset \mathbb{R}^m$, $\mathcal{V} \subset \mathbb{R}^q$ be open and let $H : \mathcal{Y} \times \mathcal{V} \to \mathbb{R}^p$ be C^r, $r > \max\{0, m-p\}$. If $\mathbf{0} \in \mathbb{R}^p$ is a regular value of H, i.e. if the Jacobian J satisfies $\operatorname{rank}(J(\boldsymbol{y},\boldsymbol{\nu})) = p$ for all $(\boldsymbol{y},\boldsymbol{\nu}) \in H^{-1}(\mathbf{0})$, then for almost every $\boldsymbol{\nu} \in \mathcal{V}$, $\mathbf{0}$ is a regular value of $H_{\boldsymbol{\nu}}(\cdot) = H(\boldsymbol{\nu},\cdot)$.

The theorem applies to H on Y. First, H is smooth (C^∞), so that the differentiability requirement is met. Second, take $\boldsymbol{\nu} \in \mathcal{V} := \mathbb{R}^{|A|}_{>0}$ (or an arbitrary open subset thereof). Third, while Y is itself not open as required, one can easily extend it to an open domain $\mathcal{Y}$ for $H(\boldsymbol{\sigma},\boldsymbol{V},t)$, for example by setting

$$\mathcal{Y} := (0, 1+\varepsilon)^{|A|} \times \mathbb{R}^{|S\times I|} \times (-\varepsilon, 1)$$

for some $\varepsilon > 0$. Finally, the following will show that the Jacobian J has full rank on $(\mathcal{Y} \times \mathcal{V}) \cap H^{-1}(\mathbf{0})$.

The Jacobian J of H is written out in Appendix D. It has the following block structure:

$$J(\boldsymbol{\sigma}, \boldsymbol{V}, \boldsymbol{\nu}, t) \quad = \quad \begin{pmatrix} \frac{\partial H^V_{sia}}{\partial \sigma_{s'i'a'}} & \frac{\partial H^V_{sia}}{\partial V_{s'i'}} & \frac{\partial H^V_{sia}}{\partial \nu_{s'i'a'}} & \frac{\partial H^V_{sia}}{\partial t} \\ \frac{\partial H^\sigma_{si}}{\partial \sigma_{s'i'a'}} & \mathbf{0} & \mathbf{0} & \mathbf{0} \end{pmatrix}$$

For detailed contents of the blocks, again refer to Appendix D. To establish full row rank, first consider the block

$$\frac{\partial H^\sigma_{si}(\boldsymbol{\sigma}, \boldsymbol{V}, \boldsymbol{\nu}, t)}{\partial \sigma_{s'i'a'}} = \begin{pmatrix} 1\dots 1 & & \mathbf{0} \\ & \ddots & \\ \mathbf{0} & & 1\dots 1 \end{pmatrix} \in \mathbb{R}^{|S\times I|\times|A|}$$

whose rows clearly are independent. Next, consider the block

$$\frac{\partial H^V_{sia}(\boldsymbol{\sigma}, \boldsymbol{V}, \boldsymbol{\nu}, t)}{\partial \nu_{s'i'a'}} = \begin{pmatrix} B_{1,1} & & \mathbf{0} \\ & \ddots & \\ \mathbf{0} & & B_{|S|,|I|} \end{pmatrix} \in \mathbb{R}^{|A|\times|A|}$$

which is itself comprised of quadratic blocks B_{si}. If an agent (s,i) has only a single action, then $B_{si} = \begin{pmatrix} 0 \end{pmatrix}$: These cases will be covered later. If otherwise $|A_{si}| \geq 2$, then

$$B_{si} = (1-t)\eta \begin{pmatrix} 1+\sigma_{si1}(\log\sigma_{si1}-1) & \sigma_{si2}(\log\sigma_{si2}-1) & \dots & \sigma_{si|A_{si}|}(\log\sigma_{si|A_{si}|}-1) \\ \sigma_{si1}(\log\sigma_{si1}-1) & 1+\sigma_{si2}(\log\sigma_{si2}-1) & \dots & \sigma_{si|A_{si}|}(\log\sigma_{si|A_{si}|}-1) \\ \vdots & & \ddots & \vdots \\ \sigma_{si1}(\log\sigma_{si1}-1) & \sigma_{si2}(\log\sigma_{si2}-1) & \dots & 1+\sigma_{si|A_{si}|}(\log\sigma_{si|A_{si}|}-1) \end{pmatrix}$$

After subtracting the first row from each other row, one obtains the following arrowhead matrix:

$$B_{si} = (1-t)\eta \left(\begin{array}{c|c} D & E \\ \hline F & G \end{array}\right)$$

$$
= (1-t)\eta \left(\begin{array}{c|ccc} 1+\sigma_{si1}(\log\sigma_{si1}-1) & \sigma_{si2}(\log\sigma_{si2}-1) & \dots & \sigma_{si|A_{si}|}(\log\sigma_{si|A_{si}|}-1) \\ \hline -1 & 1 & & \mathbf{0} \\ \dots & & \ddots & \\ -1 & \mathbf{0} & & 1 \end{array} \right)
$$

Because G is invertible, the determinant of each such block can be computed using the Schur complement:

$$
\begin{aligned}
\det(B_{si}) &= (1-t)\eta \ \det\big(G\big) \ \det\big(D - EG^{-1}F\big) \\
&= (1-t)\eta \left(1+\sigma_{si1}(\log\sigma_{si1}-1) - \sum_{k=2}^{|A_{si}|} -\sigma_{sik}\left(\log\sigma_{sik}-1\right) \right) \\
&= (1-t)\eta \sum_{k=1}^{|A_{si}|} \sigma_{sik}\log\sigma_{sik} \ < \ 0
\end{aligned}
$$

The inequality follows from $t \in [0,1)$, $\eta > 0$, and $\sigma_{sia} \in (0,1)$. The blocks B_{si} taken together thus provide a basis for all rows that do *not* correspond to singleton actions.

To complete the proof for these, we use $\frac{\partial H^V_{sia}(\boldsymbol{\sigma},\boldsymbol{V},\boldsymbol{\nu},t)}{\partial V_{s'i'}}$. Note that all entries are zero for $i \neq i'$, so that one can treat players separately. Fix any i and consider only those rows and columns corresponding to a state in which i has only a single action, enumerated as $s_1, s_2, \dots, s_n$. In these cases, $\sigma_{sia} = 1$, so that the resulting submatrix is

$$
\delta_i\bar{\boldsymbol{\phi}}_i - \boldsymbol{I}_n = \begin{pmatrix} \delta_i\bar{\phi}_{s_1 \to s_1} - 1 & \delta_i\bar{\phi}_{s_1 \to s_2} & \dots & \delta_i\bar{\phi}_{s_1 \to s_n} \\ \delta_i\bar{\phi}_{s_2 \to s_1} & \delta_i\bar{\phi}_{s_2 \to s_2} - 1 & & \dots \\ \dots & & \ddots & \\ \delta_i\bar{\phi}_{s_n \to s_1} & \dots & & \delta_i\bar{\phi}_{s_n \to s_n} - 1 \end{pmatrix}
$$

We have $\delta_i < 1$, and because $\bar{\boldsymbol{\phi}}_i$ is part of a transition matrix, all entries are between 0 and 1, with row sums less or equal to 1. This implies $||\delta_i\bar{\boldsymbol{\phi}}_i||_\infty < 1$ in maximum absolute row sum norm. Therefore, the above matrix is invertible, with

$$
\big(\delta_i\bar{\boldsymbol{\phi}}_i - \boldsymbol{I}_n\big)^{-1} = -\sum_{m=0}^{\infty} \big(\delta_i\bar{\boldsymbol{\phi}}_i\big)^m
$$

For each player, this gives a basis for the rows corresponding to singleton action sets. Since all rows of J are covered now, the proof is complete.

Note that all preceding arguments apply just as well if one considers η as an argument of H, rather than a fixed parameter, provided $\eta > 0$. The Jacobian $J(\boldsymbol{\sigma}, \boldsymbol{V}, \boldsymbol{\nu}, \eta, t)$ likewise has full row rank for all $(\boldsymbol{\sigma}, \boldsymbol{V}, t, \eta) \in \mathcal{Y} \times (0, \infty)$. Application of parametrized Sard's theorem implies that for generic $\boldsymbol{\nu}$, the zero set of $H(\boldsymbol{\sigma}, \boldsymbol{V}, \eta, t)$ is a smooth, 2-dimensional manifold in the interior of $Y \times (0, \infty)$. This proves Lemma 8.2, which in turn is part of the proof of Proposition 8.

Appendix F: ODE Representation of L^η

Here we give an explicit representation of the curve L^η, parametrized in arc length. This is done in the form of an ordinary differential equation (ODE); for a general discussion of the results used here, see e.g. Zangwill and Garcia (1981).

To simplify notation for this purpose, we collect all the variables σ_{sia}, V_{si}, and t in a vector $\boldsymbol{x} \in \mathbb{R}^{N+1}$, where $N = |A| + |S \times I|$. Similarly, enumerate the components of H as $H_1, ..., H_N$. The Jacobian of H with respect to $\boldsymbol{x}$ is then a matrix J of dimensions $N \times (N+1)$. We denote by $J_{\text{-}n}$ the square submatrix obtained by dropping the nth column from J.

By a well known application of the implicit function theorem, one can traverse any solution path contained in $H^{-1}(\mathbf{0})$ by means of a system of ordinary differential equations constructed from J, provided that J is of full rank N along this path. This system is given by $N+1$ equations

$$\dot{x}_n = (-1)^n \det\Big(J_{\text{-}n}(\boldsymbol{x})\Big)$$

Note that $\dot{\boldsymbol{x}}$ is simply a tangent vector of the path in point $\boldsymbol{x}$. For a detailed exposition and deductions, see Zangwill and Garcia (1981, equation 2.1.2, p. 26).

Since $\mathbf{0}$ is a regular value of $H|_Y$, this applies to all paths contained in Z, and in particular to the distinguished path starting at $(\boldsymbol{\sigma}^0, V^0, 0) =: \boldsymbol{x}^0$. All points on this path can be represented as the solution $\boldsymbol{x}(S)$ to an initial value problem, given by $\boldsymbol{x}(0) := \boldsymbol{x}^0$, and

$$x_n(S) = x_n^0 \pm \int_0^S \frac{\dot{x}_n(s)}{\|\dot{\boldsymbol{x}}(s)\|} ds$$

where the sign before the integral is chosen such that t initially increases. The vector field is normalized, so that the curve is parametrized in arc length, and S simply represents distance traveled along the curve.

This ODE can be used to compute the distinguished equilibrium numerically. Finding a stationary equilibrium of some game $\mathcal{G}$ is generally a hard task, as it involves solving a high-dimensional system of nonlinear equations. Standard methods to do so require a good initial guess, which will usually not be available. Homotopy methods such as the present one circumvent this problem: By construction, the starting point is easy to compute. All that is left to do then is to calculate a sequence of points along the distinguished path, a task that is much easier, because each such step can start from a solution that is very close by. In effect, the global task of finding an equilibrium of $\mathcal{G}$ is transformed into a sequence of local tasks.

Note that the prior vector can be chosen freely, for example as the centroid or some other focal strategy. Alternatively, one can perform the procedure on a collection of priors, e.g. by constructing a grid over the prior space. The procedure will potentially find different equilibria for different priors, which allows to assess the respective sizes of the basins of attraction of different equilibria.

Appendix G: Timings

G.1. Timings for Non-Generic Randomized Games

Here, we report timings for randomized non-generic games, to complement the timings of similar, but generic games in Section 6. General procedures and the computer used where as described there, except for the following changes. Values of u were drawn from a discrete uniform distribution with support $\{0, 0.1, \ldots, 1\}$. For each action profile a_s in any state, the vector of probabilities for the resulting states, $\left(\phi_{s \to s_1}(a_s), \phi_{s \to s_2}(a_s), ..., \phi_{s \to s_{|S|}}(a_s)\right)$, was created using a uniform multinomial distribution with $2|S|$ trials; this vector was then normalized so that the result sums to 1. To obtain generic (ν_{sia}), values were drawn from a uniform distribution over the interval $[0.75, 1.25]$.

Results are listed in Table G.1. In comparison to the timings for generic games in Table 1, computation times are 11% slower on average (the difference is statistically significant with $p < 0.001$, in a regression of logarithmized running time on game type, using fixed effects for the respective game sizes). As for the generic games, less than 1% of all runs are initially not successful when using a set of default path tracking parameters.

Table G.1 Timings for Non-Generic Games

$\lvert S\rvert$	$\lvert A_{si}\rvert$	$\lvert I\rvert$ = 2	3	4	5
1	2	0:00 0:00	0:00 0:00	0:00 0:00	0:00 0:00
	4	0:00 0:00	0:00 0:00	0:00 0:00	0:00 0:00
	8	0:00 0:00	0:00 0:00	0:00 0:00	0:02 0:00
2	2	0:00 0:00	0:00 0:00	0:00 0:00	0:00 0:00
	4	0:00 0:00	0:00 0:00	0:00 0:00	0:00 0:00
	8	0:00 0:00	0:00 0:00	0:01 0:00	0:05 0:02
5	2	0:00 0:00	0:00 0:00	0:00 0:00	0:00 0:00
	4	0:00 0:00	0:00 0:00	0:00 0:00	0:02 0:01
	8	0:00 0:00	0:01 0:01	0:04 0:02	0:53 0:29
10	2	0:00 0:00	0:00 0:00	0:00 0:00	0:00 0:00
	4	0:00 0:00	0:01 0:00	0:02 0:01	0:06 0:07
	8	0:01 0:00	0:04 0:02	0:21 0:11	4:04 2:45
20	2	0:00 0:00	0:00 0:00	0:01 0:01	0:02 0:01
	4	0:00 0:00	0:03 0:03	0:15 0:11	0:59 0:48
	8	0:04 0:03	0:47 0:43	5:53 2:54	51:24 35:21
50	2	0:01 0:00	0:04 0:02	0:11 0:04	0:39 0:24
	4	0:10 0:07	1:45 1:35	12:30 9:57	57:50 29:02
	8	2:42 2:28	1:03:38 33:57		
100	2	0:05 0:01	0:29 0:19	3:00 4:09	11:53 8:42
	4	1:57 1:42	33:17 26:03		
	8	55:40 48:32			
200	2	0:34 0:11	6:47 6:21	52:59 41:24	6:40:27 4:01:59
	4	50:39 35:59			
400	2	5:06 1:59			
800	2	49:14 22:34			

Computation times to solve random non-generic games with $|S|$ states, $|I|$ players and $|A_{si}|$ actions for each player in each state. Listed are average times as well as standard deviations (in small print) in *m:ss* or *h:mm:ss*. All timings under 15:00 are based on 100 independently drawn games of the respective size; all others on 10 runs per size.

G.2. Instructions for Replication

Both sets of timings can be replicated by following these steps:

1. Install the anaconda python distribution (`anaconda.com`), if you haven't yet. Then set up a virtual environment and install the required packages by running in a system terminal:

```
conda create --name logtracing-env python=3.9
conda activate logtracing-env
```

```
conda install scipy numpy=1.23.5 cython=0.29.33 openpyxl matplotlib pandas
pip install sgamesolver==1.0.2
```

2. Download the files `LogTracing_Generic.xlsx` and `LogTracing_NonGeneric.xslx` from the online supplement for this article. These files contain all the raw data from the runs performed by us, for generic and non-generic games respectively. To repeat all runs from either file, open it in Microsoft Excel or a compatible program and delete all rows but the first from the sheet named *Runs*. (Each set of timings took about 2 weeks on the computer used by us.) To repeat computation only for specific games, delete only the according rows from the sheet. Save the file and close Excel.
3. To start computation, open a system terminal and navigate to the location of the files. Then run

```
conda activate logtracing-env
sgamesolver-timings FILENAME.xlsx
```

replacing the filename accordingly. The program will begin solving the games and keep you updated about its progress. Every 5 minutes, all finished runs will be saved to the Excel file; make sure that it is not opened in Excel while computation is in progress, as this will lock writing access. Computations can be canceled by pressing `CTRL+C` at any time. Restarting from the last save is possible by running the commands again.
4. You can check the new results in the *Summary* sheet, which is updated whenever the file is saved. Note that the summary always includes all rows that appear in the *Runs* sheet (e.g. those rows you left from our runs, or from your own previous sessions). If you are interested in individual games only, check the according rows in the *Runs* sheet directly. To obtain a latex table as used in this article, run

```
conda activate logtracing-env
sgamesolver-timings -l FILENAME.xlsx
```

which will create or update a `FILENAME.tex` file in the current folder.

References

Bourbaki N (1966) Elements of Mathematics. General Topology. Part 1 (Paris and Reading: Hermann and Addison-Wesley).

Chow SN, Mallet-Paret J, Yorke JA (1978) Finding zeroes of maps: Homotopy methods that are constructive with probability one. Mathematics of Computation 32(143):887–899.

Zangwill WI, Garcia CB (1981) Pathways to Solutions, Fixed Points, and Equilibria (Upper Saddle River, New Jersey: Prentice-Hall).