Experts and Their Records - Web Appendix

Alexander Frankel Michael Schwarz

January 19, 2009

A Generalizations of Section 3 (Homogeneous Experts)

A.1 Finitely Many Experts

Our results so far have taken the set of experts to be infinite. This allows consumers to punish any number of experts by abandoning their services forever. Moving to a finite set of experts reduces the level of punishment that can be provided, which does not change our results qualitatively but restricts the range of parameter values for which a truthful equilibrium can be maintained. When there are fewer experts, it is harder to punish any one as strongly since experts will eventually be rehired. In this section we illustrate a version of the homogeneous expert truthful equilibrium when there are only μ experts: let everything else be as in Section 3, but let $E = \{e^1, e^2, ..., e^\mu\}$.

Proposition 1. *If there are μ experts, then a truthful equilibrium exists if*

$$r \geq \frac{1-\beta}{1-\beta\left(\frac{\beta-\beta p}{1-\beta p}\right)^{\mu-1}}$$

As μ increases towards infinity this cutoff falls to $1-\beta$, the level in the infinite case.

Customers can no longer keep moving to new experts after firing an old one. Instead they rotate through experts, returning to the first after the μ^{th} is fired.

The intuition here is the same as the Investor Game of Fudenberg and Levine (1994). In that paper, the long-run players are tempted to steal from the short-run players, and so the short-run players are hesitant to make efficient investments. An efficient equilibrium becomes possible as the number of long-run capitalists increases because it becomes easier for the short-run investors to punish any given capitalist by moving money elsewhere rather than withdrawing from the market entirely. This is in contrast to Park (2005), in which equilibria are more difficult to maintain with more experts. This is because in Park's model, each customer has a strict preference for some randomly chosen expert. When there are more experts, each gets less future business and therefore has a smaller promised payoff from following the equilibrium instead of deviating.

Notice that we can no longer ignore p. Individual experts are still indifferent across actions when they are chosen: M has a higher payoff today but also yields a higher probability of being fired. One expert prefers that others play M, however – this makes the other experts more likely to be fired, bringing this expert closer to being rehired. The lower p is, the more likely are major treatments, and the sooner an expert returns to work after getting fired. Hence, the lower p is, the more difficult it is to punish an expert and the smaller is the range of supportable parameter values.

Proof. We prove this by construction.

Let each expert's strategy be to play truthfully: for each $a \in A$, for each $e^i \in E$, and for each $H_t^i \subseteq \mathcal{H}^i$, let $\sigma^i(H_t^i, \theta^a) = a$.

Suppose the customer's strategy is the following. c^1 chooses $\rho^1(H_1) = e^1$. For $t > 1$, c^t chooses

$$\rho^t\big(H_{t-1} \oplus (e_{t-1} = e^i, a_{t-1})\big) = \begin{cases} e^i \text{ with Prob } 1 & \text{if } a_{t-1} = m \\ e^i \text{ with Prob } q, e^{[i+1]} \text{ with Prob } 1-q & \text{if } a_{t-1} = M \end{cases}$$

where $[x]$ is defined to be x if $x \leq \mu$ and 1 if $x = \mu + 1$.

For what value of q will the experts' strategies be optimal? We can solve this as a dynamic programming problem. Since experts are homogeneous, we can focus on expert e^1. Let V^i be expert e^1's expected lifetime payoff at a period in which e^i is chosen, prior to the realization of the state.

The strategy is optimal for e^1 (and thus for all experts) if, when chosen, her lifetime utility is equal across actions:

$$V^1 = 1 + \beta\left(qV^1 + (1-q)V^2\right) = r + \beta V^1. \tag{1}$$

Unconditional on the state there is a probability $(1-p)(1-q)$ of an expert being fired in any period, and $1-(1-p)(1-q) = p+q-pq$ probability of not being fired. So

$$\begin{aligned} V^2 &= \beta\left((p+q-pq)V^2 + (1-p)(1-q)V^3\right) \\ &\dots \\ V^{\mu-1} &= \beta\left((p+q-pq)V^{\mu-1} + (1-p)(1-q)V^{\mu}\right) \\ V^{\mu} &= \beta\left((p+q-pq)V^{\mu} + (1-p)(1-q)V^1\right) \end{aligned}$$

For $i > 1$, V^i can be written as $V^i = \frac{\beta(1-p)(1-q)V^{[i+1]}}{1-\beta(p+q-pq)}$. Writing V^μ in terms of V^1 and working backwards, we see that

$$V^2 = \left(\frac{\beta(1-p)(1-q)}{1-\beta(p+q-pq)}\right)^{\mu-1} V^1$$

We can now plug this back into (1) to get

$$V^1 = 1 + \beta\left(qV^1 + (1-q)\left(\frac{\beta(1-p)(1-q)}{1-\beta(p+q-pq)}\right)^{\mu-1} V^1\right) = r + \beta V^1$$

and $V^1 = \frac{r}{1-\beta}$, so we have

$$1 + \beta\left(q\frac{r}{1-\beta} + (1-q)\left(\frac{\beta(1-p)(1-q)}{1-\beta(p+q-pq)}\right)^{\mu-1} \frac{r}{1-\beta}\right) = r + \beta\frac{r}{1-\beta}$$

If we can find a $q \in [0,1]$ which satisfies this equation, then the expert strategy will be a best response and we will have an equilibrium. Rearranging,

$$r = \frac{1-\beta}{1-\beta q - \beta(1-q)\left(\frac{\beta(1-p)(1-q)}{1-\beta(p+q-pq)}\right)^{\mu-1}}$$

For $q = 1$, the right-hand side is 1; for $q = 0$, the right-hand side is $\frac{1-\beta}{1-\beta^{\mu}\left(\frac{1-p}{1-\beta p}\right)^{\mu-1}}$. Because the expression is continuous in q, the intermediate value theorem tells us that we can find a q which solves the equation for any r in between $\frac{1-\beta}{1-\beta^{\mu}\left(\frac{1-p}{1-\beta p}\right)^{\mu-1}}$ and 1. $\square$

A.2 A Larger Action Set & Observable Heterogeneity

The truthful equilibrium of Section 3 naturally extends to a larger action set than the two considered in the body of the paper, and can also take into account observable payoff heterogeneity. The key is that after an expert takes an action, consumers can infer the expert's payoff relative to her other available actions.

Let the set of actions A be an arbitrary (possibly uncountably infinite) set with at least two elements and let $\Theta = \{\theta^a | a \in A\}$. $\theta_t \in \Theta$ is realized according to some stationary distribution. The stage payoff for an expert who plays action a is $R^i(a)$. Customers are indifferent across experts but prefer that action a be taken in state θ^a. We have an infinite set of experts as in Sections 3 and 4.

Normalize expert payoffs such that $\sup_a R^i(a) = 1$ for all e^i. Let $r = \inf_{a \in A, i \in \mathbb{N}} R^i(a)$, and suppose that $r > 0$. Let $r = \inf_{a \in A, i \in \mathbb{N}} R^i(a)$, and suppose that $0 < r \leq R^i(a) \leq 1$ for all $a \in A, i \in \mathbb{N}$.

As in the previous sections, customers observe the history of experts chosen and actions taken. Moreover, customers observe $R^i(a)$ after expert e^i takes action a.[1]

We can now restate the analog of Proposition 1 from Section 3:

Proposition 2. *A truthful equilibrium exists if* $\beta \geq 1 - r$.

A single expert with particularly low payoffs for a given action would not necessarily ruin the equilibrium. The logic goes through as long as customers can restrict their interactions to any infinite subset of experts for whom $R^i(a) \geq 1 - \beta$ for all a.

Proof. Suppose that $\beta \geq 1 - r$. We will construct a truthful equilibrium.

Let each expert's strategy be to play truthfully: for each $a \in A$, for each $e^i \in M$, and for each $H_t^i \subseteq \mathcal{H}^i$, let $\sigma^i(H_t^i, \theta^a) = a$.

Define the customer's strategy as the following. c^1 chooses $\rho^1(H_1) = e^1$. For $t > 1$, c^t chooses

$$\rho^t\big(H_{t-1} \oplus (e_{t-1} = e^i, a_{t-1})\big) = \begin{cases} e^i \text{ with Prob } q^i(a_{t-1}) \\ e^{i+1} \text{ with Prob } 1 - q^i(a_{t-1}) \end{cases}$$

with $q^i(a) = \frac{r + R^i(a)\beta - R^i(a)}{r\beta}$. Notice that $0 \leq q^i(a) \leq 1$ because $1 - \beta \leq r \leq R^i(a) \leq 1$.

Because all experts play truthfully at every history, customers are indifferent across experts and any strategy is a best response.

To check that truthful play is a best response for the experts, consider the expected payoff V for an expert who follows the strategy, conditional on being chosen in a given period but unconditional on the realization of θ. By the one-shot deviation principle, the expert's strategy is optimal if she

[1] The crucial assumption is that customers condition their strategies on $R^i(a)$. They may be able to do this because $R^i(a)$ is ex post observable, because it is ex ante common knowledge, or simply because (as in Section 3) the customers somehow know the correct probabilities to play in equilibrium after expert e^i chooses action a. $R^i(a)$ might be ex post observable even if it is not known ex ante if $R^i(a)$ is proportional to (observable) time spent working.

always plays a maximizer of $R^i(a) + \beta q^i(a)V$, and if $V = \max_{a\in A}\{R^i(a) + \beta q^i(a)V\}$. This holds if we can find a V such that

$$\begin{aligned} V &= R^i(a) + \beta \frac{r + R^i(a)\beta - R^i(a)}{r\beta} V \quad \text{for all } a \\ \iff V &\left(1 - \frac{r + R^i(a)\beta - R^i(a)}{r}\right) = R^i(a) \quad \text{for all } a \\ \iff V &= \frac{r}{1-\beta} \quad \text{for all } a \end{aligned}$$

So the continuation value V is $\frac{r}{1-\beta}$, and there are no profitable deviations.

This shows that the above strategies are an equilibrium – the experts and customers are indifferent with respect to all actions at all histories. □

B Static Mechanisms and Prices

So far, an expert's utility for each action has been determined exogenously. Experts could only be incentivized to act truthfully by offering them the possibility of future business. There may be a way to induce experts to be truthful without having to invoke dynamic incentives: set prices to make experts indifferent between the two actions today. In the current section, we explore the extent to which a static mechanism with properly set prices can induce truthfulness. What we find is that static incentives are not enough. In a general environment of privately observed payoffs, no static incentives can ever support a positive probability of truthful play.

To formalize the new model, say that expert e^i has types drawn from some distribution over $\mathbb{R}^2$. The distribution may vary across experts, but each expert's the distribution is common knowledge. We will denote a representative type of an expert in $\mathbb{R}^2$ by $r = (r_m, r_M)$. An expert of type (r_m, r_M) who is selected, takes action a, and receives a transfer payment of λ gets utility $r_a + \lambda$; an expert who is not selected and receives a transfer payment of λ gets utility λ. The payoff of a major treatment is no longer normalized to 1 – because we have a numeraire good, the payoff levels matter and a normalization would not be without loss of generality. The types are naturally interpreted now as underlying costs, rather than profits net of cost.

There is a single period of play. As before, we leave customer behavior largely unmodeled and focus on the expert's side of the problem. Rather than searching for on an optimal contract that maximizes consumer welfare, we look for the maximum level of truthful play which can be supported by any type of contract.

If an expert's e^i's type (r^i_m, r^i_M) were common knowledge, then the customer could easily induce truthfulness with properly set prices. The customer could select e^i and pay her an amount of λ_a for playing a, where $r^i_m + \lambda_m = r^i_M + \lambda_M$. The expert's total payoff would be constant across actions and so she would be willing to be truthful.

When experts are homogeneous – and so, for all intents and purposes, any expert's type is common knowledge – the stories of repeat business and of prices give alternate avenues for supporting fully truthful play. We will show that in the heterogeneous setting we consider, static environments with endogenous prices will not be able to replicate the truthful play which we found in the equilibria of Section 4.

In the one-shot setting we consider here, the customer arrives on the market and then an expert is chosen and prices are set according to some process. For instance, experts might post prices which influence the customer's choice; or the customer might choose an expert randomly, at which point the two would haggle over a price. The revelation principle implies that, without loss of generality, we can model any such environment using a corresponding direct mechanism. Experts report their types simultaneously, and based on these reports the mechanism chooses an expert and determines prices. The chosen expert observes the state and the prices and then chooses an action. We consider mechanisms which set action-dependent payments that will be offered to the selected expert, and also provide a transfer payment to each expert that is independent of the expert selected or the action taken.[2]

Definition. Let a *mechanism* be a sequence of triplets $\left(Q^i(\vec{r}), \lambda^i(a|r), \lambda_0^i(r)\right)_{i\in\mathbb{N}}$ where $Q^i : \prod_{j=1}^{\infty} \mathbb{R}^2 \to [0,1]$, $\lambda^i : A \times \mathbb{R}^2 \to \mathbb{R}$, and $\lambda_0^i : \mathbb{R}^2 \to \mathbb{R}$. We require that $Q^i(\vec{r}) \geq 0$ for all $i, \vec{r}$ and $\sum_i Q^i(\vec{r}) = 1$ for each $\vec{r}$.

Each expert e^i simultaneously reports a value $r^i \in \mathbb{R}^2$. Expert e^i is then selected to serve the customer with probability $Q^i(\vec{r})$, where $\vec{r} \equiv (r^1, r^2, ...)$. The selected expert e^j observes the state $\theta \in \Theta$ and chooses an action $a \in A$. Each expert e^i then receives a transfer payment of $\lambda_0^i(r^i)$, and the chosen expert e^j further receives $\lambda^j(a|r^j)$.

Let $\tilde{Q}^i(r)$ be expert e^i's probability of being selected as a function of her report, unconditional on all other reports.

Say that a mechanism is *direct revelation incentive compatible* if there is an equilibrium of the corresponding game in which all experts of all types report their types honestly. We restrict ourselves to direct revelation incentive compatible mechanisms and focus on equilibria in which types are reported honestly. Because we are interested in truthful play, we will assume that any expert will be truthful if truthfulness is a best response.

Let (r^i DRIC) be the direct revelation incentive compatibility constraint for type r^i, and let (r^i Play a) be the condition under which action a is a best response for expert e^i of type r^i if chosen. We say that expert e^i of type r^i is "willing to play a" if the (r^i Play a) constraint is satisfied. An expert e^i of type r^i will act truthfully if and only if both (r^i Play m) and (r^i Play M) hold; I call this joint condition (r^i Truth). Writing these out,

$$r^i \in \underset{r\in\mathbb{R}^2}{\operatorname{argmax}} \left[\max_{a\in A} \left(\lambda_0^i(r) + \tilde{Q}(r)\big(r_a^i + \lambda(a|r)\big) \right) \right] \tag{r^i DRIC}$$

$$r_a^i + \lambda(a|r^i) \geq r_{\tilde{a}}^i + \lambda(\tilde{a}|r^i) \text{ for } \tilde{a} \neq a \text{ if } r^i \text{ is willing to play } a \tag{r^i Play a}$$

$$r_m^i + \lambda(m|r^i) = r_M^i + \lambda(M|r^i) \text{ if } r^i \text{ plays truthfully if chosen} \tag{r^i Truth}$$

We now state a lemma which will help to show limits on the level of truthfulness achievable by a mechanism.

Lemma 1.

i. If expert e^i has types $r' \neq r''$ which are both truthful, and if $r'_{a'} > r''_{a'}$ and $r'_{a''} < r''_{a''}$ for $a' \neq a'' \in A$, then $\tilde{Q}^i(r') = \tilde{Q}^i(r'') = 0$.

[2] We model a mechanism as determining the choice of expert and the conditional transfer payments (prices), after which the expert then chooses an action based on the state of the world. An alternate approach would be to incorporate the expert's choice of action in each state into the mechanism itself. The two approaches are equivalent.

It would also be possible to generalize the mechanism beyond what we consider, allowing transfer payments to each expert to be fully dependent on the vector of all reports as well as the selected expert and the action taken. This would complicate notation without changing any results.

ii. If expert e^i has types $r' \neq r''$ which are both truthful, and if $r'_{a'} = r''_{a'}$ and $r'_{a''} < r''_{a''}$ for $a' \neq a'' \in A$, then $\tilde{Q}^i(r') = 0$.

Proof. We first show helpful result. For a given expert e^i, take any $r', r'' \in \mathbb{R}^2$. Suppose that, for $a', a'' \in A$ (possibly identical), the expert of type r' is willing to play a' and r'' is willing to play a''. Then

$$\tilde{Q}^i(r')(r'_{a'} - r''_{a'}) \geq \tilde{Q}^i(r'')(r'_{a''} - r''_{a''}) \tag{2}$$

This holds because

$$\begin{aligned} \lambda_0(r') + \tilde{Q}(r')(r'_{a'} + \lambda(a|r')) &\geq \lambda_0^i(r'') + \tilde{Q}^i(r'')(r'_{a''} + \lambda(a''|r'')) \qquad \text{(By } r' \text{ DRIC + } r' \text{ Play } a'\text{)} \\ &= \left[\lambda_0^i(r'') + \tilde{Q}^i(r'')(r''_{a''} + \lambda^i(a''|r''))\right] + \tilde{Q}^i(r'')(r'_{a''} - r''_{a''}) \\ &\geq \left[\lambda_0^i(r') + \tilde{Q}^i(r')(r''_{a'} + \lambda^i(a'|r'))\right] + \tilde{Q}^i(r'')(r'_{a''} - r''_{a''}) \\ &\qquad \text{(By } r'' \text{ DRIC + } r'' \text{ Play } a''\text{)} \end{aligned}$$

$$\Longrightarrow$$

$$\tilde{Q}^i(r')(r'_{a'} - r''_{a'}) \geq \tilde{Q}^i(r'')(r'_{a''} - r''_{a''})$$

Now, to prove part (i), suppose that $r' \neq r''$ are both truthful if chosen, where $r'_{a'} > r''_{a'}$ and $r'_{a''} < r''_{a''}$. Then by Equation (2), $\tilde{Q}^i(r')(r'_{a''} - r''_{a''}) \geq \tilde{Q}^i(r'')(r'_{a'} - r''_{a'})$ because r' is willing to play a'' and r'' is willing to play a'. The left-hand side is nonpositive and the right-hand side is nonnegative, so both sides must be equal to 0. This is only possible if $\tilde{Q}^i(r')$ and $\tilde{Q}^i(r'')$ are 0.

To prove part (ii), if r' and r'' are truthful if chosen then Equation (2) implies that

$$\tilde{Q}^i(r')(r'_{a''} - r''_{a''}) \geq \tilde{Q}^i(r'')(r'_{a'} - r''_{a'}).$$

If $r'_{a'} = r''_{a'}$ and $r'_{a''} < r''_{a''}$ then the parenthetical term on the left-hand side is negative and the parenthetical term on the right-hand side is 0. Therefore $\tilde{Q}^i(r') \leq 0$, which can only hold if $\tilde{Q}^i(r') = 0$. □

This lemma implies that the set of types of expert e^i which are picked with positive probability and act truthfully if picked must have a strict ordering in efficiency. Any one such point is either larger in both r_a dimensions or smaller in both than another such point. In other words, all truthful types of a given expert which may be picked with positive probability lie on a strictly increasing function in (r_m, r_M)-space.

For the following Proposition, we will say that a distribution is *nondegenerate* if it places a probability of 0 on any set of measure 0.

Proposition 3. *Suppose that for each e^i, the ex ante distribution of r^i over $\mathbb{R}^2$ is nondegenerate. Then in any mechanism, the ex ante probability of the chosen mechanic playing truthfully is 0.*

Proof. This is an immediate corollary of Lemma 1. For any r^i_m, there is at most a single value of r^i_M for which e^i will be truthful if selected. Therefore the set of truthful types has measure 0 in $\mathbb{R}^2$, and the probability that e^i realizes a truthful type from a nondegenerate distribution is 0. There are countably many experts and each is truthful if selected with probability 0, so the probabilty that any is of a type which is truthful if selected is 0. □

References

Drew Fudenberg and David K. Levine. Efficiency and observability in games with long-run and short-run players. *Journal of Economic Theory*, 62:103–135, 1994.

In-Uck Park. Cheap talk referrals of differentiated experts in repeated relationships. *RAND Journal of Economics*, 36(2), Summer 2005.