6

The Mixing Problem: Purification and Conjectures

God does not play dice with the universe.

Albert Einsten

6.1 Why Play Mixed Strategies?

In Throwing Fingers (§2.7), there is a unique mixed strategy Nash equilibrium in which both players choose each of their pure strategies with probability 1/2. However, if both pure strategies have equal payoffs against the mixed strategy of the other player, why bother randomizing? Of course, this problem is perfectly general. By the Fundamental Theorem (§2.5), any mixed strategy best response consists of equal-payoff pure strategies, so why should a player bother randomizing? Moreover, this argument holds for all other players as well. Therefore, no player should expect any other player to randomize. This is the mixing problem.

We assume that the game is played only once (this is called a *one-shot* game, even though it could be an extensive form game with many moves by each player, as in chess), so there is no past history on which to base an inference of future play, and each decision node is visited at most once, so a statistical analysis during the course of the game cannot be carried out. If the stage game itself consists of a finite repetition of a smaller stage game, as in playing Throwing Fingers n times, one can make a good case for randomizing in the stage game. But, only a small fraction of games have such form.

One suggestion for randomizing is that perhaps it is easier for an opponent to discover a pure strategy choice on your part than a mixed strategy choice (Reny and Robson 2003). Indeed, when Von Neumann xxand Morgenstern (1944) introduced the concept of a mixed strategy in zero-sum games, they argued that a player would use this strategy to "protect himself against having his intentions found out by his opponent." (p. 146) This defense is weak. For one thing, it does not hold at all for many games, such as Battle of the Sexes (§2.8), where a player *gains* when his partner discovers his pure strategy move. More important, if there are informational processes, reputation effects, or other mechanisms whereby the agents' "types" can be known or discovered, this should be formally modeled in the specification of the game.

Finally, it is always *costly* to randomize, because one must have some sophisticated mental algorithm modeling the spinning of a roulette wheel or other randomizing device, so mixed strategies are in fact strictly more costly to implement than pure strategies. In general, we do not discuss implementation costs, but in this case they play a critical role in evaluating the relative costs of playing a mixed strategy best response or any of the pure strategies in its support. Hence the *mixing problem*: why bother randomizing?

There have been two major approaches to solving the mixing problem. The first approach, which we develop in §6.2, is due to Harsanyi (1973). Harsanyi treats mixed strategy equilibria as limit cases of slightly perturbed "purified" games with pure strategy equilibria. This remarkable approach handles many simple games very nicely, but fails to extend to more complex environments. The more recent approach, which uses interactive epistemology to define knowledge structures representing subjective degrees of uncertainty, is due to Robert Aumann and his coworkers. This approach, which we develop in §6.5, does not predict how agents will actually play, because it determines only the *conjectures* each player has of the *other* players' strategies. It follows that this approach does not solve the mixing problem. However, in §6.6, we show that a simple extension of the Aumann conjecture approach is valid under precisely the same conditions as Harsanyi purification, with the added attractions of handling pure strategy equilibria (§6.6) and applying to cases where payoffs are determinate.

Our general conclusion, which is fortified by two examples below, is that purification is possible in some simple games, but not in the sort of games that apply to complex social interaction, such as principal-agent models or repeated games. Yet, such complex models generally rely on mixed strategies. Thus, game theory alone is incapable of explaining such complex social interactions, even approximately, and even in principle. This provides one more nail in the coffin of methodological individualism (§8.8).

6.2 Harsanyi's Purification Theorem

To understand Harsanyi's (1973) defense of the mixed strategy Nash equilibrium in one-shot games, consider the game to the right, which has a mixed strategy equilibrium $(\sigma_1^*, \sigma_2^*) = (3U/4 + D/4, L/2 + R/2)$. Because player 1, Alice, is indifferent between U and D, if she had some personal, idiosyncratic, reason for slightly preferring U to D, she would play pure strategy U rather than the mixed strategy σ_1^*. Similarly, some slight preference for R would lead player 2, Bob, to play pure strategy R. The mixed strategy equilibrium would then disappear, thus solving the mixing problem.

	L	R
U	0,0	0,-1
D	-1,0	1,3

To formalize this, let θ be a random variable uniformly distributed on the interval $[-1/2, 1/2]$ (in fact, θ could be any bounded random variable with a continuous density), and suppose that the distribution of payoffs to U for the population of player 1's is $2 + \epsilon\theta_1$, and the distribution of payoffs to L for the population of player 2's is $1 + \epsilon\theta_2$, where θ_1 and θ_2 are independently distributed as θ. Suppose Alice and Bob are chosen randomly from their respective populations of player 1's and player 2's to play the game, and each knows only the distribution of payoffs of their partners. Suppose Alice uses U with probability β and Bob uses L with probability α. You can check that Bob infers that the payoff to U for Alice is distributed as $\pi_U = \epsilon\theta_1$, and the payoff to D for Alice is $\pi_D = 1 - 2\alpha$. Similarly, Alice infers that the payoff to L for Bob is $\pi_L = \epsilon\theta_2$, and the payoff to R for Bob is $\pi_R = 3 - 4\beta$.

	L	R
U	$\epsilon\theta_1$,$\epsilon\theta_2$	$\epsilon\theta_1$,-1
D	-1,$\epsilon\theta_2$	1,3

Now, β is the probability $\pi_U > \pi_D$, which is the probability that $\theta_1 > (1 - 2\alpha)/\epsilon$, which gives

$$\alpha = \mathrm{P}[\pi_L > \pi_R] = \mathrm{P}\left[\theta_B > \frac{3 - 4\beta}{\epsilon}\right] = \frac{8\beta - 6 + \epsilon}{2\epsilon} \tag{6.1}$$

$$\beta = \mathrm{P}[\pi_U > \pi_D] = \mathrm{P}\left[\theta_A > \frac{1 - 2\alpha}{\epsilon}\right] = \frac{4\alpha - 2 + \epsilon}{2\epsilon}. \tag{6.2}$$

Solving these simultaneously for α and β, we find

$$\alpha = \frac{1}{2} - \frac{\epsilon}{8 - \epsilon^2} \qquad \beta = \frac{3}{4} - \frac{\epsilon^2}{4(8 - \epsilon^2)}, \tag{6.3}$$

This is our desired equilibrium. Now, this *looks* like a mixed strategy equilibrium, but it is not. The probability that an Alice randomly chosen from the population chooses pure strategy U is β, and the probability that a Bob randomly chosen from the population chooses pure strategy L is α. Thus, for instance, if an observer measured the frequency with which player 1's chose U, he would arrive at a number near α, despite the fact that no player 1 ever randomizes. Moreover, when ϵ is very small, $(\alpha,\beta) \approx (\sigma_1^*,\sigma_2^*)$. Thus, if we observed a large number of pairs of agents playing this game, the frequency of play of the various strategies would closely approximate their mixed strategy equilibrium values.

To familiarize himself with this analysis, the reader should derive the equilibrium values of α and β assuming that player i's idiosyncratic payoff is uniformly distributed on the interval $[a_i, b_i]$,$i = 1,2$, and show that α and β tend to the mixed strategy Nash equilibrium as $\epsilon \to 0$. Then, find the purified solution to Throwing Fingers (§2.7) assuming Player 1 favors H with probability $\epsilon\theta_1$, where θ_1 is uniformly distributed on [-0.5,0.5] and Player 2 favors H with probability $\epsilon\theta_2$, where θ_2 is uniformly distributed on [0,1], and show that as $\epsilon \to 0$, the strategies in the perturbed game move to the mixed strategy equilibrium.

Govindan, Reny and Robson (2003) present a very general statement and elegant proof of Harsanyi's purification theorem. They also correct an error in Harsanyi's original proof (see also van Damme, 1987, Ch. 5). The notion of a *regular equilibrium* used in the theorem is the same as that of a *hyperbolic fixed point* in dynamical systems theory (Gintis 2009), and is satisfied in many simple games with isolated and strictly perfect (meaning that if we add very small errors to each strategy, the equilibrium is displaced only a small amount) Nash equilibria.

The following theorem is weakened a bit from Govindan et al. (2003) to make it easier to follow. Let $\mathcal{G}$ be a finite normal form game with pure strategy set S_i for player i, $i = 1,\ldots,n$, and payoffs $u_i : S \to \mathbf{R}$, where $S = \prod_{i=1}^n S_i$ is the set of pure strategy profiles of the game.

A Nash equilibrium s is *strict* if there is a neighborhood of s (considered as a point in n-space) that contains no other Nash equilibrium of the game. The distance between equilibria is the Euclidean distance between the strategy profiles considered as points in $\mathbf{R}^{|S|}$ ($|S|$ is the number of elements in S). Another way of saying this is that a Nash equilibrium is strict if the connected component of Nash equilibria to which it belongs consists of a single point.

Suppose for each player i and each pure strategy profile $s \in S$, there is a random perturbation $\nu_i(s)$ with probability distribution μ_i, such that the actual payoff to player i to $s \in S$ is $u_i(s) + \epsilon\nu_i(s)$, where $\epsilon > 0$ is a small number. We assume the ν_i are independent, and each player knows only his own outcomes $\{\nu_i(s) | s \in S_i\}$. We have

THEOREM 6.1 *Suppose σ^* is a regular mixed strategy Nash equilibrium of $\mathcal{G}$. Then for every $\delta > 0$ there is an $\epsilon > 0$ such that the perturbed game with payoffs $\{u_i(s) + \epsilon\nu_i(s) | s \in S\}$ has a strict Nash equilibrium $\hat{\sigma}$ within ϵ of σ^*.*

6.3 A Reputational Model of Honesty and Corruption

Consider a society with many individuals. In the first period, a pair is selected randomly, one being designated "Needy" and the other "Giver." Giver and Needy then play a stage game $\mathcal{G}$ in which if Giver Helps, a benefit b is conferred on Needy at a cost c to Giver, where $0 < c < b$; or, if Giver Defects, both players receive zero. In each succeeding period, the Needy from the previous period becomes the Giver of the current period, is paired with a random Needy, and the game $\mathcal{G}$ is played by the new pair. If we assume that helping behavior is public information, there is a Nash equilibrium of the following form. At the start of the game, each player is labeled "in good standing." In every period, a Giver Helps if and only if his Needy partner is in good standing. Failure to do so puts a player "in bad standing," where he remains for the rest of the game. It is clear that this strategy implements a Nash equilibrium in which every Giver Helps in every period.

Suppose, however the informational assumption is that each new Giver knows only whether his Needy partner did or did not help his partner in the previous period. Now, if Giver Alice's Needy partner Bob did not Help when he was Giver, it could be either because when Bob was Giver, his Needy partner Carole had Defected when she was Giver, or because Bob failed to Help Carole even though she had Helped her previous Needy partner Donald when she was Giver. Because Alice cannot condition her action on Bob's previous action, Bob's best response is to Defect on Carole, no matter what she did. Therefore, Carole will Defect on Donald, no matter what he did. Thus, there can be no Nash equilibrium with the pure strategy Help.

This argument extends to the richer informational structure where a Giver knows the previous k actions, for any finite k. Here is the argument for $k=2$, which the reader is encouraged to generalize. Suppose the last five players are Alice, Bob, Carole, Donald, and Eloise, in that order. Alice can condition her choice on the actions taken by Bob, Carole, and Donald, but not on Eloise's action. Therefore, Bob's best response to Carole will not be conditioned on Eloise's action, and hence Carole's response to Donald will not be conditioned on Eloise's action so, finally Donald's response to Eloise will not be conditioned her action, so her best response is to Defect when she is Giver. Thus, there is no Helping Nash equilibrium.

Suppose, however, back in the $k=1$ case, instead of Defecting unconditionally when facing a Needy partner who had Defected improperly, a Giver Helps with probability $p = 1 - c/b$ and Defects with probability $1 - p$. The gain from Helping unconditionally is then $b - c$, while the gain from following this new strategy is $p(b - c) + (1 - p)pb$, where the first term is the probability p of Helping times the reward b in the next period if one Helps minus the cost c of Helping in the current period, and the second term is the probability $1 - p$ of Defecting times the probability p that you will be Helped anyway when your are Needy, times the benefit b. Equating this expression with $b - c$, the cost of Helping unconditionally, we get $p = 1 - c/b$, which is a number strictly between zero and one, and hence a valid probability.

Consider the following strategy. On each round, the Giver Helps if his partner Helped in the previous period, and otherwise Helps with probability p and Defects with probability $1 - p$. With this strategy each Giver i is indifferent to Helping or Defecting, because Helping costs i the amount c when Giver but i gains b when Needy, for a net gain of $b - c$. However, Defecting costs zero when Giver, but gives $bp = b - c$ when Needy. Because the two actions have the same payoff, it is incentive compatible for each Giver to Help when his Needy partner Helped, and Defect with probability p otherwise. This strategy thus gives rise to a Nash equilibrium with Helping in every period.

The bizarre nature of this equilibrium is clear from the fact that there is no reason for any player to follow this strategy as opposed to any other, since all strategies have the same payoff. So, for instance, if you slightly favor some players (e.g., your friends or co-religionists) over others (e.g., your enemies and religious infidels), then you will Help the former and Defect on the latter. But then, if this is generally true, each player Bob knows that he will be Helped or not by Alice independent of whether Bob Helps the Needy Carole, so Bob has no incentive to Help Carole. In short, if we add a small amount of "noise" to the payoffs in the form of slight preference for some potential partners over others, there is no longer a Nash equilibrium with Helping. Thus, this repeated game model with private signals cannot be purified (Bhaskar 1998b).

However, if players receive a subjective payoff from "following the rules" (we term this a *normative predisposition* in chapter 7) greater than the largest subjective gain from Helping a friend or loss from Defecting on an enemy, complete cooperation can be reestablished even with private signals. Indeed, even more sophisticated models can be constructed in which Alice can calculate the probability that the Giver with whom she is paired when Needy will condition his action on hers, a calculation that depends on the statistical distribution of her friends and enemies, and the statistical distribution of the strength of the predisposition to observe social norms. For certain parameter ranges of these variables, Alice will behave "meritocratically," and for others, she will act "corruptly" by favoring friends over enemies.

6.4 Purifying Honesty and Corruption

Suppose police are hired to apprehend criminals, but only the word of the police officer who witnessed the transgression is used to punish the offender—there is no forensic evidence involved in the judgment, and the accused have no means of self-defense. Moreover, it costs the police officer a fixed amount f to file a criminal report. How can this society erect incentives to induce the police to act honestly?

Let us assume that the society's Elders set up criminal penalties so that it is never profitable to commit a crime, provided the police are honest. The police, however, are self-regarding, and so have no incentive to report a crime, which costs them f. If the Elders offer the police an incentive w per criminal report, the police will file zero reports for $w < f$ and as many as possible for $w > f$. However, if the Elders set $w = f$, the police will be indifferent between reporting and not reporting a crime, and there is a Nash equilibrium in which the officers report all crimes they observe, and none others.

This Nash equilibrium cannot be purified. If there are small differences in the cost of filing a report, or if police officers derive small differences in utility from reporting crimes, depending on their relationship to the perpetrator, the Nash equilibrium disappears. We can foresee how this model could be transformed into a full-fledged model of police honesty and corruption, by adding effective monitoring devices, keeping tabs of the reporting rate of different officers, and the like. We could also add a in the form of a police culture favoring honesty or condemning corruption, and explore the interaction of moral and material incentives in controlling crime.

6.5 Epistemic Games: Mixed Strategies as Conjectures

Let $\mathcal{G}$ be an epistemic game, where each player i has a subjective prior $p_i(\cdot;\omega)$. We say $p \in \Delta\Omega$ is a *common prior* for $\mathcal{G}$ if, for each player i, and for each $P \in \mathcal{P}_i$, $p(P) > 0$ and i's subjective prior $p_i(\cdot|P)$ satisfies $p_i(\omega|P)p(P) = p(\omega)$ for $\omega \in P$,

The sense in which conjectures solve the problem of why agents play mixed strategies is given by the following theorem due to Aumann and Brandenburger (1995), which we will prove in §8.7.

THEOREM 6.2 *Let $\mathcal{G}$ be an epistemic game with $n > 2$ players. Suppose the players have a common prior p, and it is commonly known at $\omega \in \Omega$ that ϕ^ω is the set of conjectures for $\mathcal{G}$. Then for each $j = 1,\ldots,n$, all $i \neq j$ induce the same conjecture $\sigma_j(\omega) = \phi_j^\omega$ about j's conjectured mixed strategy, and $(\sigma_1(\omega),\ldots,\sigma_n(\omega))$ form a Nash equilibrium of $\mathcal{G}$.*

Several game theorists have suggested that this theorem resolves the problem of mixed strategy Nash equilibria. In their view, each player chooses a pure strategy, but there is a Nash equilibrium in player *conjectures* (see, for instance, §4.3). However, the fact that player conjectures are mutual best responses does not permit us to deduce anything concerning the relative frequency of player pure strategy choices, except that pure strategies not in the support of the equilibrium mixed strategy will have frequency zero. This suggested solution to the mixing problem is thus incorrect, assuming one cares about explaining behavior.

There are many stunning indications of contemporary game theorists' disregard for explaining behavior, but perhaps none more stunning that the complacency surrounding the acceptance of this argument. The methodological commitment behind this complacency was eloquently expressed by Ariel Rubinstein in his presidential address to the Econometric Society (Rubinstein 2006). "As in the case of fables, models in economic theory…are not meant to be testable.…a good model can have an enormous influence on the real world, not by providing advice or by predicting the future, but rather by influencing culture." It is hard not be sympathetic with Rubinstein's disarming frankness, despite his being dead wrong: the value of model is its contribution to explaning reality, not its contribution to society's stock of pithy aphorisms.

6.6 Resurrecting the Conjecture Approach to Purification

Harsanyi purification is motivated by the notion that payoffs may have a statistical distribution, rather than being the determinate values assumed in classical game theory. Suppose, however, that payoffs are indeed determinate, but the conjectures (§6.5) of individual players have a statistical distribution around the game's mixed strategy equilibrium values. In this case, the epistemic solution to the mixing problem might be cogent. Indeed, as we shall see, the assumption of stochastic conjectures has advantages over the Harsanyi assumption of stochastic payoffs. This avenue of research has not been studied in the literature, to my knowledge, but it clearly deserves to be explored.

Consider the Battle of the Sexes (§2.8), and suppose that Alfredo's in the population play opera with mean probability α, but the conjecture as to α by the population of Violetta's is distributed as $\alpha + \epsilon\theta_V$. Similarly, suppose that Violetta's in the population play opera with mean probability β, but the conjecture as to β by the population of Alfredo's is distributed as $\beta + \epsilon\theta_A$.

Let π_o^A and π_g^A be the expected payoffs to a random Alfredo chosen from the population from playing opera and gambling, respectively, and let π_o^V and π_g^V be the expected payoff to a random Violetta chosen from the population from playing opera and gambling, respectively. An easy calculation shows that

$$\pi_o^A - \pi_g^A = 3\beta - 2 + 3\epsilon\theta_A$$
$$\pi_o^V - \pi_g^V = 3\alpha - 1 + 3\epsilon\theta_V.$$

Therefore

$$\alpha = \mathrm{P}[\pi_o^A > \pi_g^A] = \mathrm{P}\left[\theta_A > \frac{2-3\beta}{3\epsilon}\right] = \frac{6\beta - 4 + 3\epsilon}{6\epsilon} \tag{6.4}$$

$$\beta = \mathrm{P}[\pi_o^V > \pi_g^V] = \mathrm{P}\left[\theta_V > \frac{1-3\alpha}{3\epsilon}\right] = \frac{6\alpha - 2 + 3\epsilon}{6\epsilon}. \tag{6.5}$$

If we assume that the beliefs of the agents reflect the actual state of the two populations, we may solve these equations simultaneously, getting

$$\alpha^* = \frac{1}{3} + \frac{\epsilon}{6(1+\epsilon)}, \qquad \beta^* = \frac{2}{3} - \frac{\epsilon}{6(1+\epsilon)} \tag{6.6}$$

Clearly, as $\epsilon \to 0$, this pure strategy equilibrium tends to the mixed strategy equilibrium of the stage game (2/3,1/3), as prescribed by the Purification Theorem.

However, note that in our calculations in arriving at (6.6) assumed that $\alpha, \beta \in (0, 1)$. This is true, however, only when

$$\frac{1}{3} - \frac{\epsilon}{2} < \alpha < \frac{1}{3} + \frac{\epsilon}{2}, \qquad \frac{2}{3} - \frac{\epsilon}{2} < \beta < \frac{2}{3} + \frac{\epsilon}{2}. \tag{6.7}$$

Suppose, however, that $\alpha < 1/3 - \epsilon/2$. Then, all Violettas choose gambling, to which gambling is any Alfredo's best response. In this case, the only equilibrium is $\alpha = \beta = 0$. Similarly, if $\alpha > 1/3 + \epsilon/6$, then all Violettas choose opera, so all Alfredos choose opera, and we have the pure strategy Nash equilibrium $\alpha = \beta = 1$. This approach to solving the mixing problem has the added attraction that it yields not only an approximation to the mixed strategy equilibrium for some statistical distribution of beliefs, but to one or another of the two pure strategy equilibria with other distributions of beliefs.