Supplemental Appendices for "Direct Reciprocity Under Uncertainty Does Not Explain One-shot Cooperation, but Demonstrates the Benefits of a Norm Psychology"

Matthew R. Zefferman

A Probability of One-Shot Cooperation for All Parameters and Treatments

The figures over the next few pages show the results from all runs of the simulations as described in the main text.

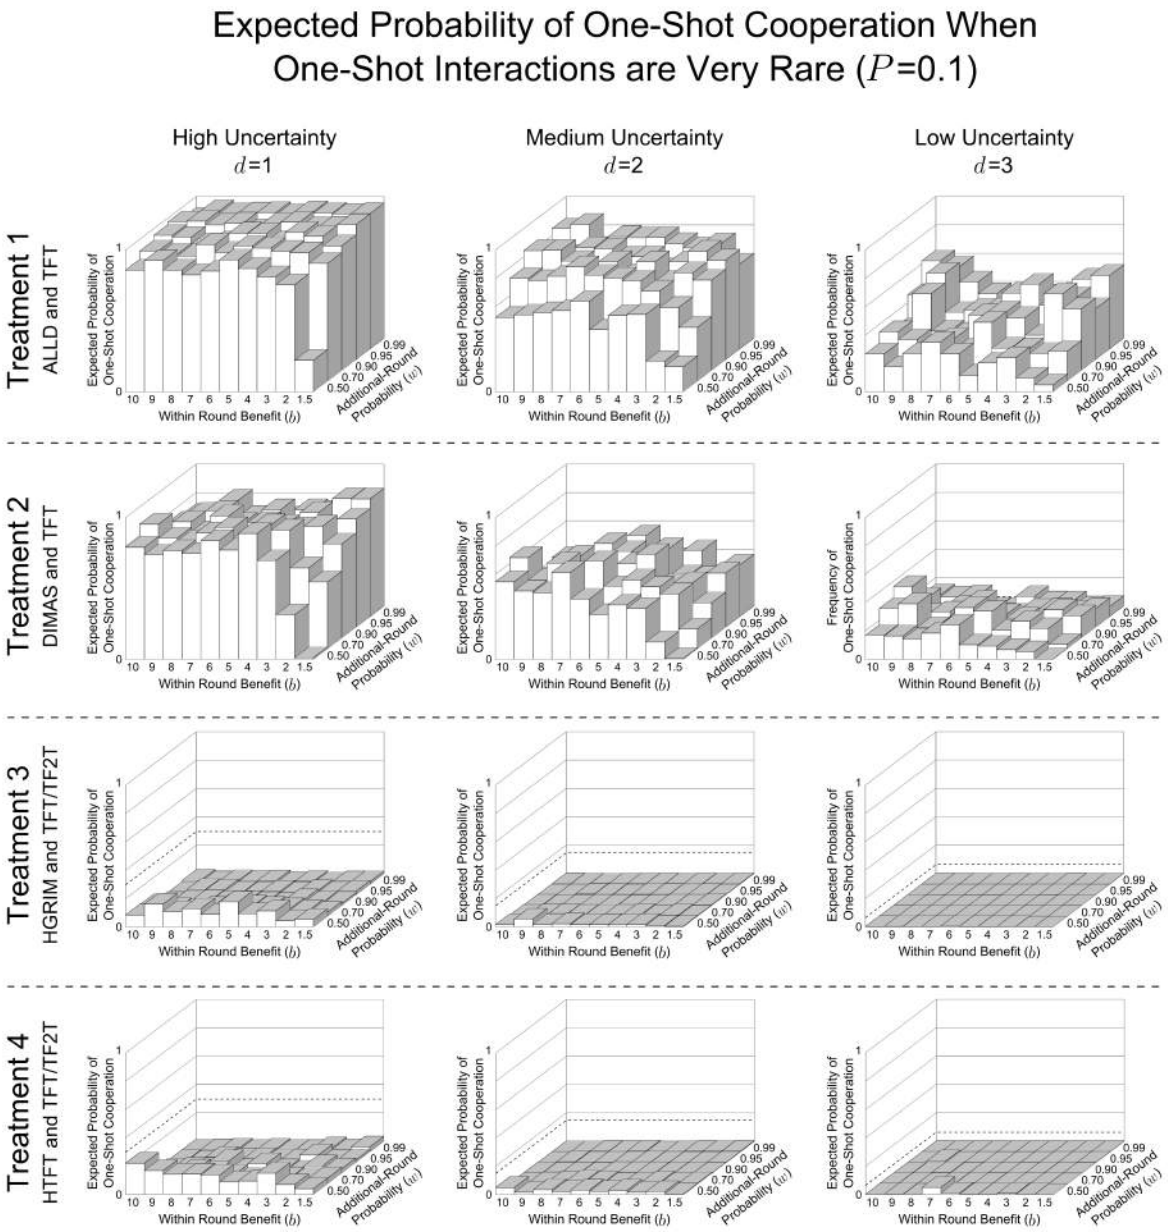

Figure A.1: **Repentant and forgiving strategies decrease one-shot cooperation when one-shot games are very rare ($P = 0.1$) and agents have 10 partners.** These show the expected probability of one-shot cooperation averaged over the last 500 generations of the 10,000 generation simulation for all values of d, b, and w. These are the same parameter combinations reported by Delton et al. (2011). Treatment 1, where agents, as in DKCT play only TFT or ALLD has the highest probability of one-shot cooperation. Treatment 2, where ALLD is replaced by a repentant strategy, DIMAS, has less one-shot cooperation. In Treatment 3, where a savvy repentant strategy, HGRIM, competes with TFT and TF2T, one-shot cooperation virtually disappears. And Treatment 4, where HGRIM is replaced by HTFT. This highlights the variability in outcomes for the same environmental parameters.

Figure A.2: **Repentant and forgiving strategies decrease one-shot cooperation when one-shot games are fairly rare ($P = 0.3$) and agents have 10 partners.** These show the expected probability of one-shot cooperation averaged over the last 500 generations of the 10,000 generation simulation for all values of d, b, and w. These are the same parameter combinations reported by Delton et al. (2011). Treatment 1, where agents, as in DKCT play only TFT or ALLD has the highest probability of one-shot cooperation. Treatment 2, where ALLD is replaced by a repentant strategy, DIMAS, has less one-shot cooperation. In Treatment 3, where a savvy repentant strategy, HGRIM, competes with TFT and TF2T, one-shot cooperation virtually disappears. And Treatment 4, where HGRIM is replaced by HTFT. This highlights the variability in outcomes for the same environmental parameters.

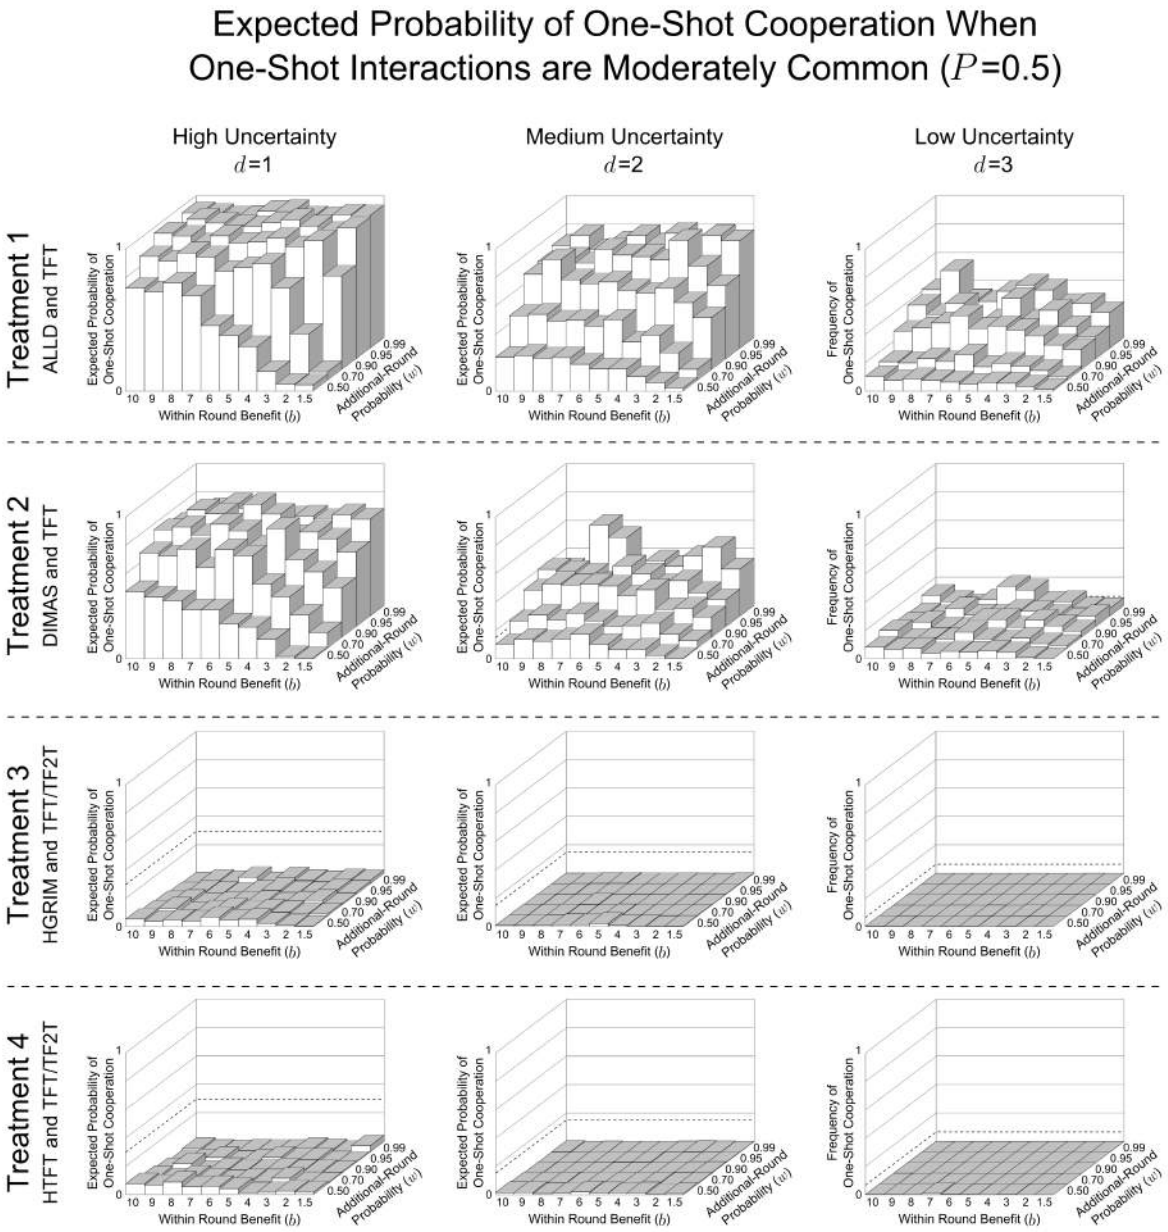

Figure A.3: **Repentant and forgiving strategies decrease one-shot cooperation when one-shot games are moderately rare ($P = 0.5$) and agents have 10 partners.** These show the expected probability of one-shot cooperation averaged over the last 500 generations of the 10,000 generation simulation for all values of d, b, and w. These are the same parameter combinations reported by Delton et al. (2011). Treatment 1, where agents, as in DKCT play only TFT or ALLD has the highest probability of one-shot cooperation. Treatment 2, where ALLD is replaced by a repentant strategy, DIMAS, has less one-shot cooperation. In Treatment 3, where a savvy repentant strategy, HGRIM, competes with TFT and TF2T, one-shot cooperation virtually disappears. And Treatment 4, where HGRIM is replaced by HTFT. This highlights the variability in outcomes for the same environmental parameters.

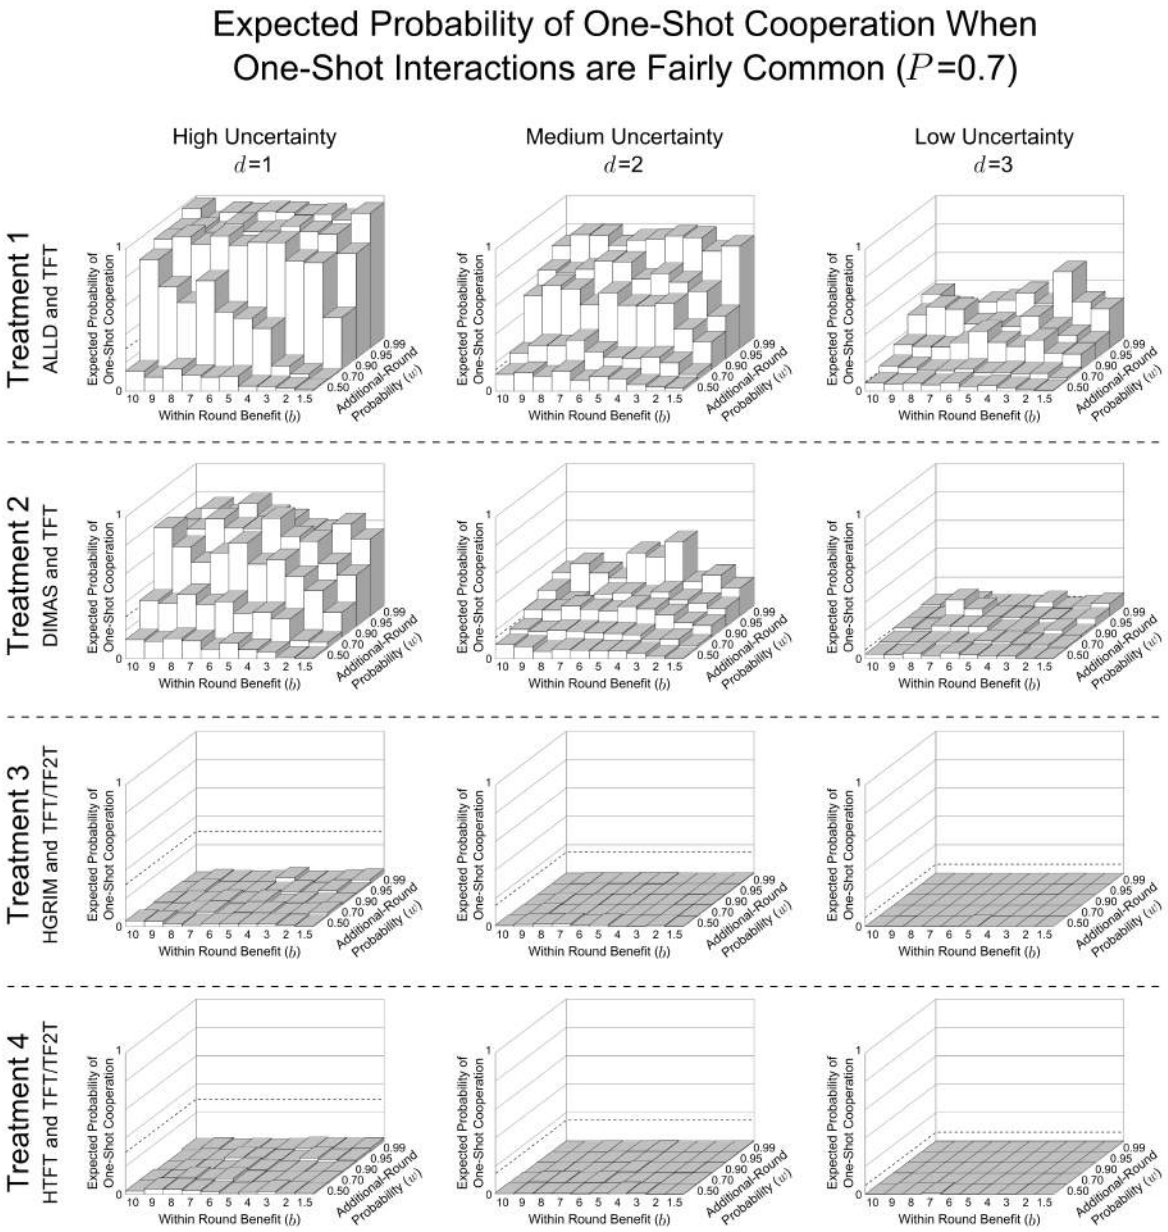

Figure A.4: **Repentant and forgiving strategies decrease one-shot cooperation when one-shot games are fairly common ($P = 0.7$) and agents have 10 partners.** These show the expected probability of one-shot cooperation averaged over the last 500 generations of the 10,000 generation simulation for all values of d, b, and w. These are the same parameter combinations reported by Delton et al. (2011). Treatment 1, where agents, as in DKCT play only TFT or ALLD has the highest probability of one-shot cooperation. Treatment 2, where ALLD is replaced by a repentant strategy, DIMAS, has less one-shot cooperation. In Treatment 3, where a savvy repentant strategy, HGRIM, competes with TFT and TF2T, one-shot cooperation virtually disappears. And Treatment 4, where HGRIM is replaced by HTFT. This highlights the variability in outcomes for the same environmental parameters.

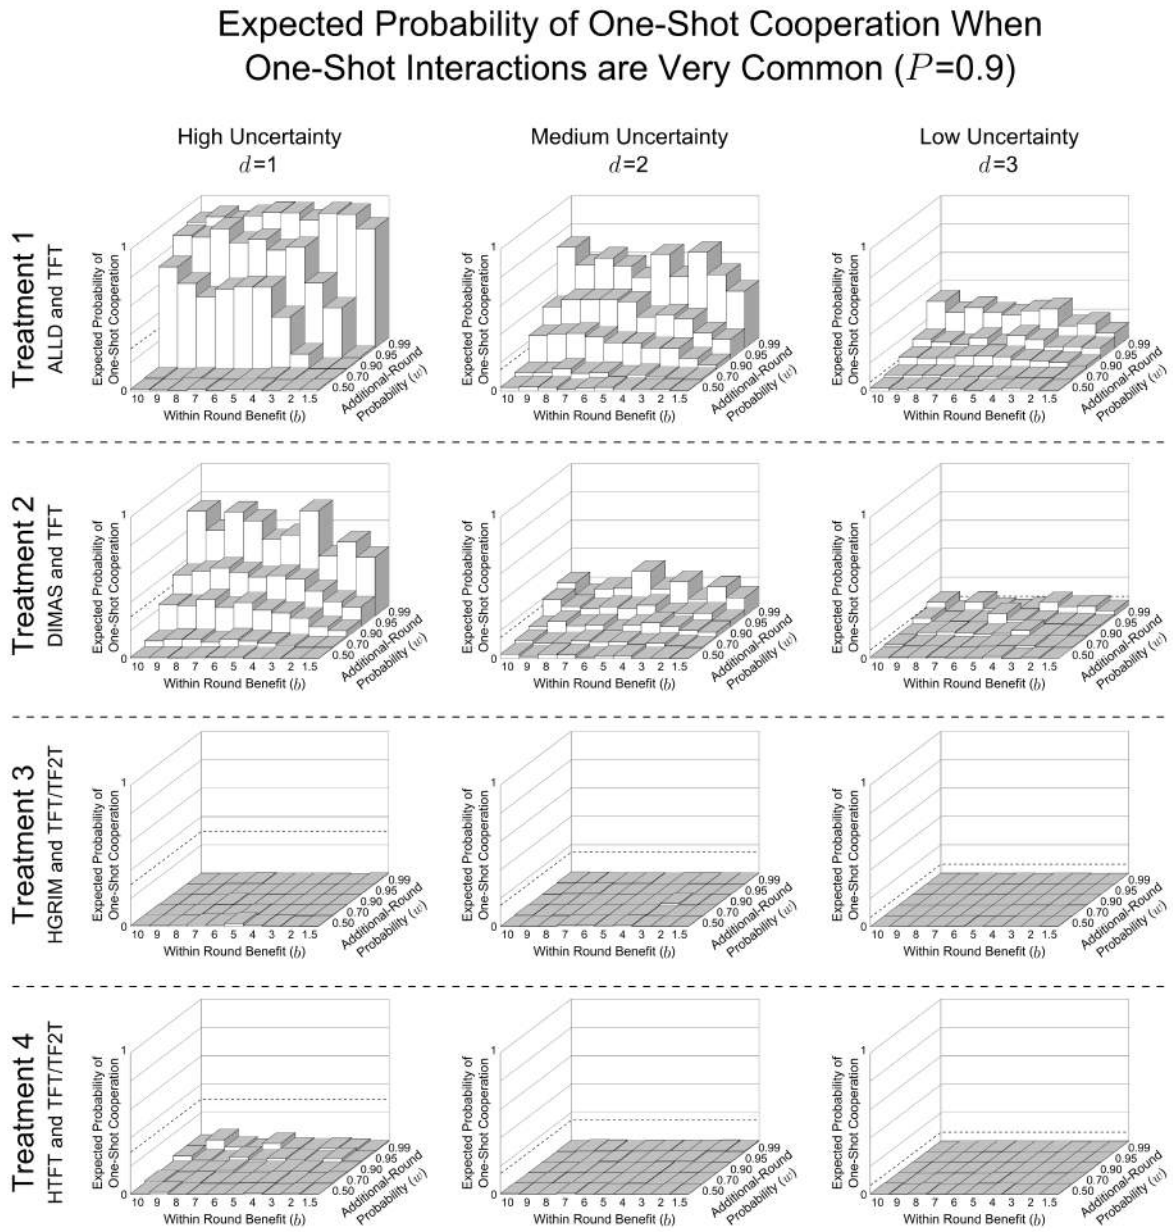

Figure A.5: **Repentant and forgiving strategies decrease one-shot cooperation when one-shot games are very common ($P = 0.9$) and agents have 10 partners.** These show the expected probability of one-shot cooperation averaged over the last 500 generations of the 10,000 generation simulation for all values of d, b, and w. These are the same parameter combinations reported by Delton et al. (2011). Treatment 1, where agents, as in DKCT play only TFT or ALLD has the highest probability of one-shot cooperation. Treatment 2, where ALLD is replaced by a repentant strategy, DIMAS, has less one-shot cooperation. In Treatment 3, where a savvy repentant strategy, HGRIM, competes with TFT and TF2T, one-shot cooperation virtually disappears. And Treatment 4, where HGRIM is replaced by HTFT. This highlights the variability in outcomes for the same environmental parameters.

Figure A.6: **Repentant and forgiving strategies decrease one-shot cooperation when one-shot games are very rare ($P = 0.1$) and agents have one partner.** These show the expected probability of one-shot cooperation averaged over the last 500 generations of the 10,000 generation simulation for all values of d, b, and w. These are the same parameter combinations reported by Delton et al. (2011). Treatment 1, where agents, as in DKCT play only TFT or ALLD has the highest probability of one-shot cooperation. Treatment 2, where ALLD is replaced by a repentant strategy, DIMAS, has less one-shot cooperation. In Treatment 3, where a savvy repentant strategy, HGRIM, competes with TFT and TF2T, one-shot cooperation virtually disappears. And Treatment 4, where HGRIM is replaced by HTFT. This highlights the variability in outcomes for the same environmental parameters.

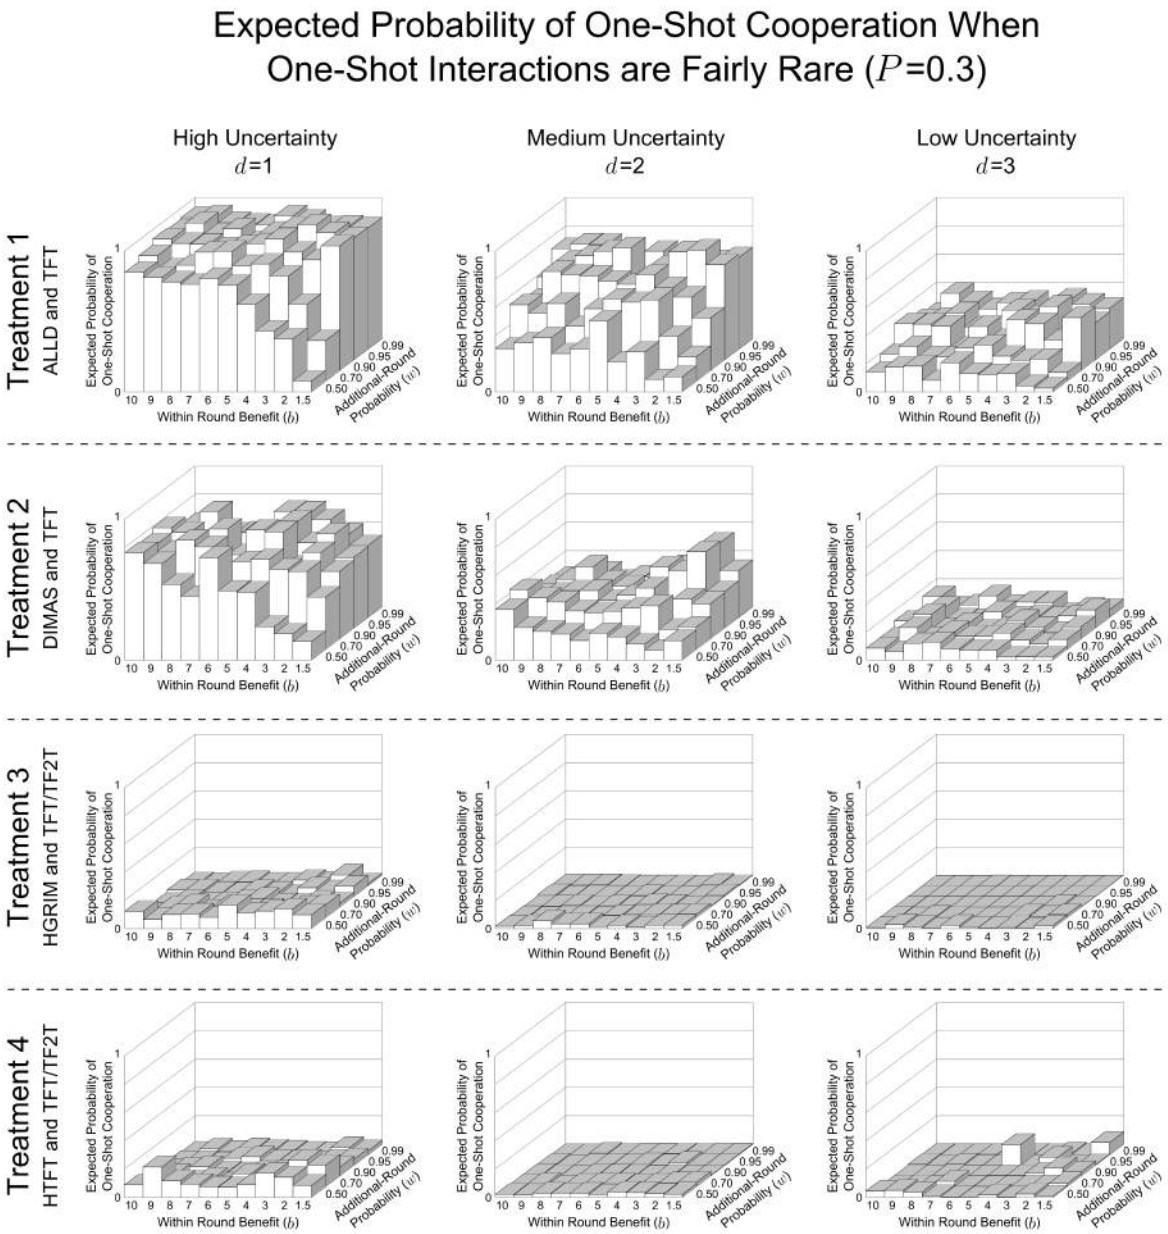

Figure A.7: **Repentant and forgiving strategies decrease one-shot cooperation when one-shot games are fairly rare ($P = 0.3$) and agents have one partner.** These show the expected probability of one-shot cooperation averaged over the last 500 generations of the 10,000 generation simulation for all values of d, b, and w. These are the same parameter combinations reported by Delton et al. (2011). Treatment 1, where agents, as in DKCT play only TFT or ALLD has the highest probability of one-shot cooperation. Treatment 2, where ALLD is replaced by a repentant strategy, DIMAS, has less one-shot cooperation. In Treatment 3, where a savvy repentant strategy, HGRIM, competes with TFT and TF2T, one-shot cooperation virtually disappears. And Treatment 4, where HGRIM is replaced by HTFT. This highlights the variability in outcomes for the same environmental parameters.

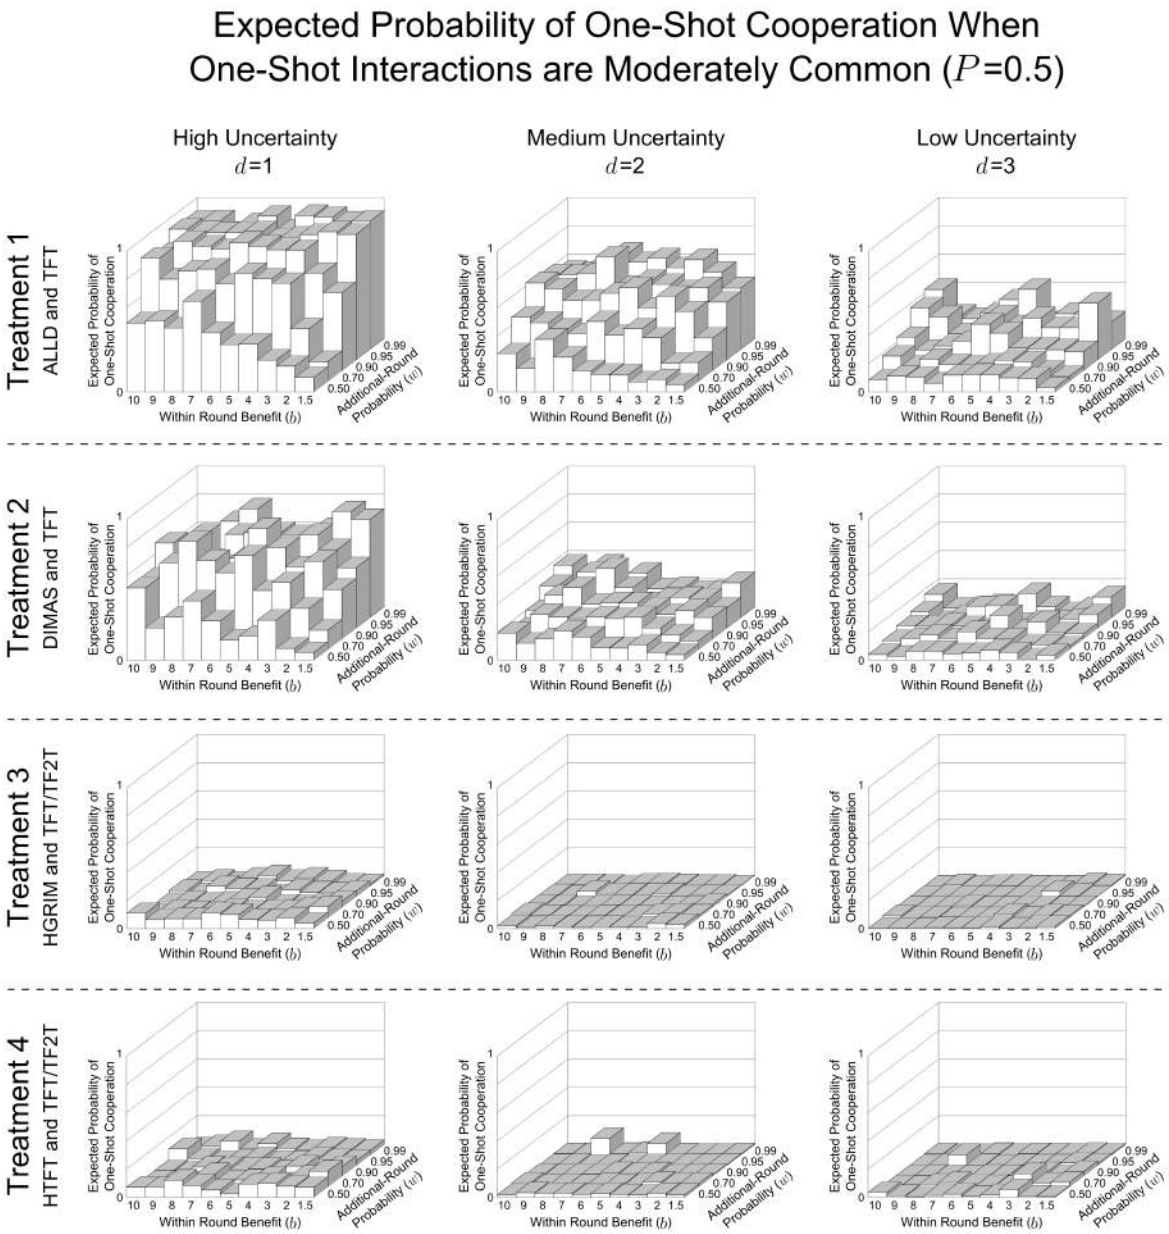

Figure A.8: **Repentant and forgiving strategies decrease one-shot cooperation when one-shot games are moderately rare ($P = 0.5$) and agents have one partner.** These show the expected probability of one-shot cooperation averaged over the last 500 generations of the 10,000 generation simulation for all values of d, b, and w. These are the same parameter combinations reported by Delton et al. (2011). Treatment 1, where agents, as in DKCT play only TFT or ALLD has the highest probability of one-shot cooperation. Treatment 2, where ALLD is replaced by a repentant strategy, DIMAS, has less one-shot cooperation. In Treatment 3, where a savvy repentant strategy, HGRIM, competes with TFT and TF2T, one-shot cooperation virtually disappears. And Treatment 4, where HGRIM is replaced by HTFT. This highlights the variability in outcomes for the same environmental parameters.

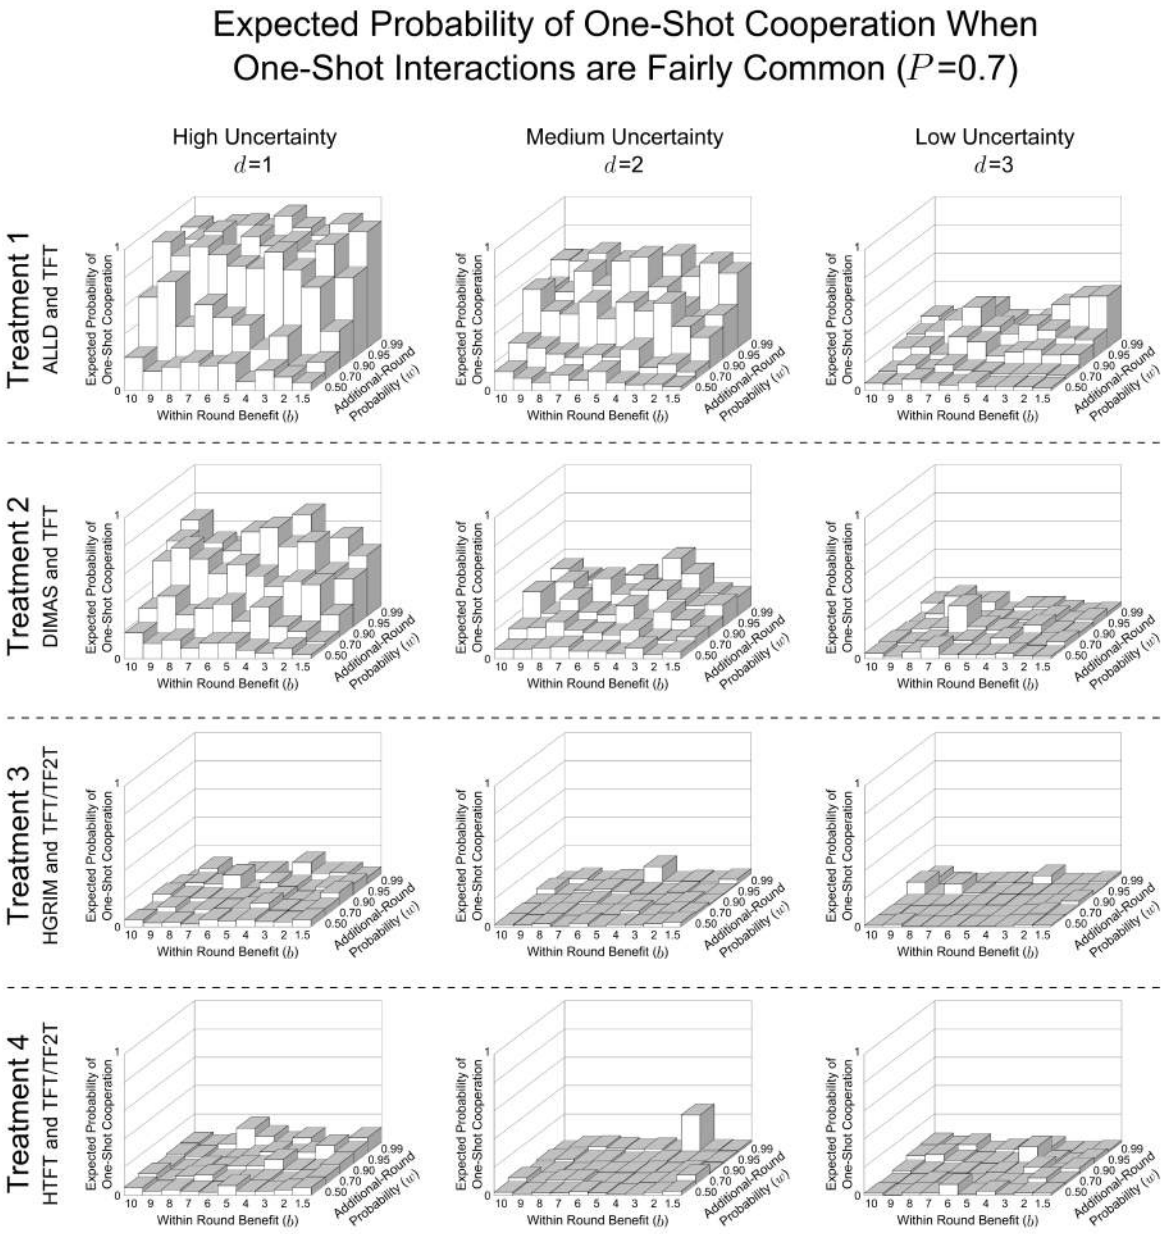

Figure A.9: **Repentant and forgiving strategies decrease one-shot cooperation when one-shot games are fairly common ($P = 0.7$) and agents have one partner.** These show the expected probability of one-shot cooperation averaged over the last 500 generations of the 10,000 generation simulation for all values of d, b, and w. These are the same parameter combinations reported by Delton et al. (2011). Treatment 1, where agents, as in DKCT play only TFT or ALLD has the highest probability of one-shot cooperation. Treatment 2, where ALLD is replaced by a repentant strategy, DIMAS, has less one-shot cooperation. In Treatment 3, where a savvy repentant strategy, HGRIM, competes with TFT and TF2T, one-shot cooperation virtually disappears. And Treatment 4, where HGRIM is replaced by HTFT. This highlights the variability in outcomes for the same environmental parameters. Some populations show increased one-shot cooperation. This is due to stochastic shocks as explained in C.

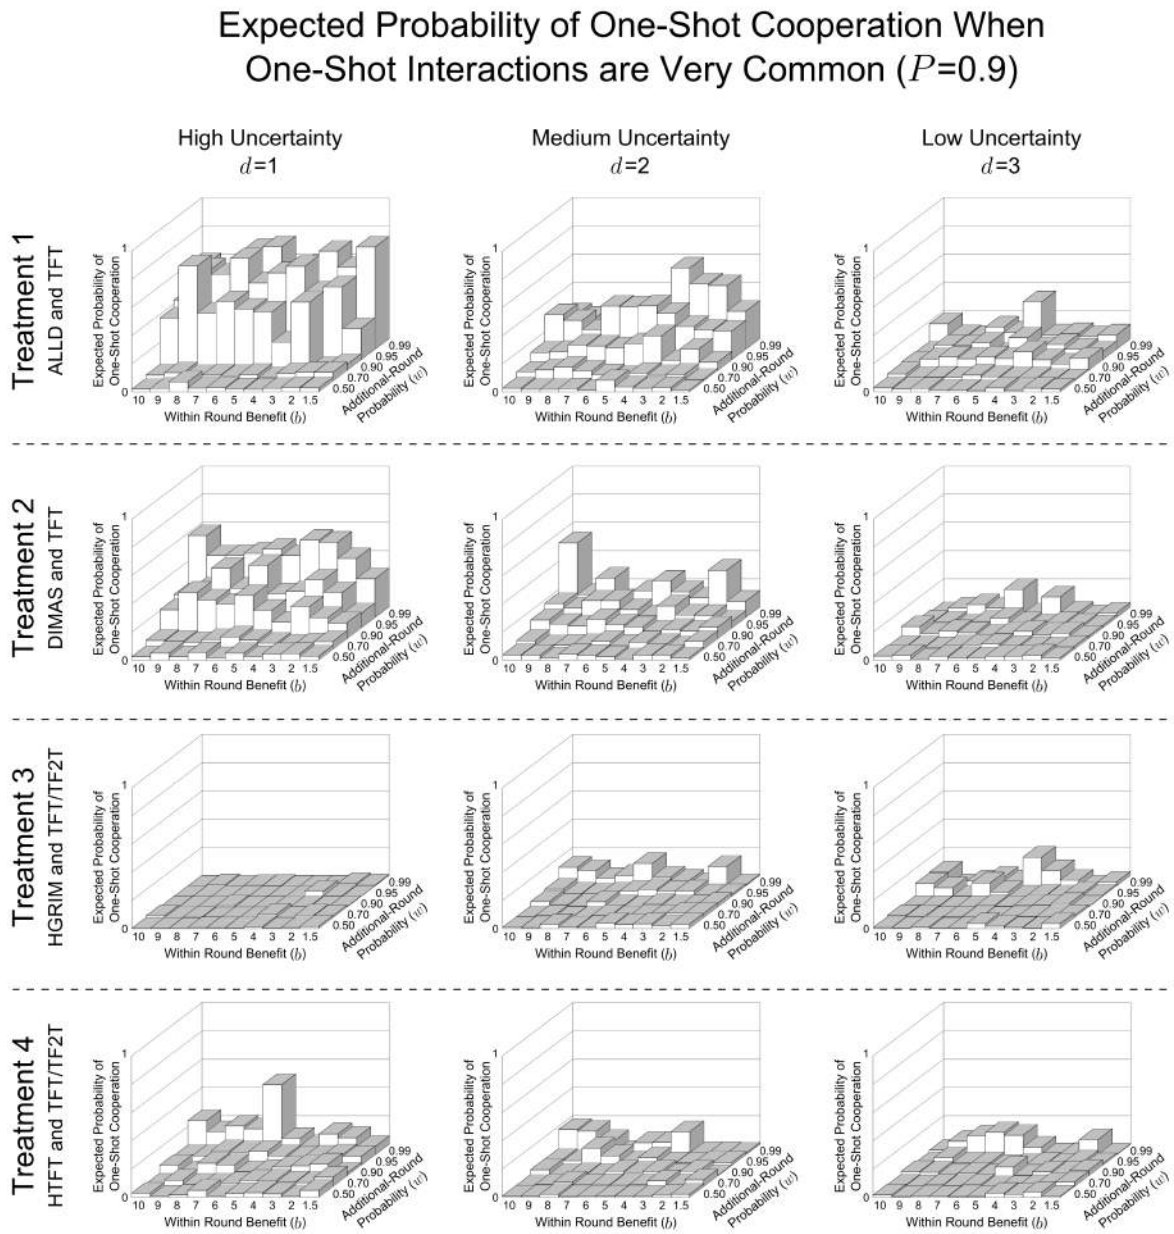

Figure A.10: **Repentant and forgiving strategies decrease one-shot cooperation when one-shot games are very common ($P = 0.9$) and agents have one partner.** These show the expected probability of one-shot cooperation averaged over the last 500 generations of the 10,000 generation simulation for all values of d, b, and w. These are the same parameter combinations reported by Delton et al. (2011). Treatment 1, where agents, as in DKCT play only TFT or ALLD has the highest probability of one-shot cooperation. Treatment 2, where ALLD is replaced by a repentant strategy, DIMAS, has less one-shot cooperation. In Treatment 3, where a savvy repentant strategy, HGRIM, competes with TFT and TF2T, one-shot cooperation virtually disappears. And Treatment 4, where HGRIM is replaced by HTFT. This highlights the variability in outcomes for the same environmental parameters. Some populations show increased one-shot cooperation. This is due to stochastic shocks as explained in C.

B Complete Uncertainty Calculations

The best-case scenario for the evolution of one-shot cooperation in the DKCT model is under the condition of complete uncertainty, where there is no signal concerning whether a game is one-shot or repeated. Under complete uncertainty, DKCT's model reduces to a standard repeated Prisoner's Dilemma where the probability of transitioning from the first to second round is $(1-P)w$, which is lower than the transition probability for every other round, w. Fig. B.1 gives the expected payoffs for the strategy combinations in all four treatments under complete uncertainty. In this appendix, I show the calculations used to derive the stability and risk-dominance conditions in Table 3.

Treatments 1 and 2

	TFT	ALLD (Treatment 1)	DIMAS (Treatment 2)
TFT	$(b-c)\left(\frac{1-wP}{1-w}\right)$	$-c$	$-c+(b-wc)\frac{(1-P)w}{1-w}$
ALLD	b	0	
DIMAS	$b+(wb-c)\frac{(1-P)w}{1-w}$		$(b-c)\frac{(1-P)w}{1-w}$

Treatments 3 and 4

	TFT	TF2T	HGRIM (Treatment 3)	HTFT (Treatment 4)
TFT	$(b-c)\left(\frac{1-wP}{1-w}\right)$	$(b-c)\left(\frac{1-wP}{1-w}\right)$	$-c+w(1-P)(b-wc)$	$-c+w(1-P)\frac{b-wc}{1-w^2}$
TF2T	$(b-c)\left(\frac{1-wP}{1-w}\right)$	$(b-c)\left(\frac{1-wP}{1-w}\right)$	$-c+(b-c)\frac{(1-P)w}{1-w}$	$-c+w(1-P)\frac{b-c}{1-w}$
HGRIM	$b+w(1-P)(wb-c)$	$b+(b-c)\frac{(1-P)w}{1-w}$	$(b-c)\frac{(1-P)w}{1-w}$	
HTFT	$b+w(1-P)\frac{wb-c}{1-w^2}$	$b+w(1-P)\frac{b-c}{1-w}$		$w(1-P)\frac{b-c}{1-w}$

Figure B.1: **Expected payoffs for strategies in the four treatments when agents are completely uncertain whether a game is one-shot or repeated.** b is the benefits to cooperation, c is the cost of cooperating, P is the probability that an interaction is one-shot, and w is the probability that, after each round in a repeated game, there is another round. TFT occurs in all treatments. ALLD occurs in Treatment 1. DIMAS occurs in Treatment 2. HGRIM occurs in Treatment 3. HTFT occurs in Treatment 4. TF2T occurs in Treatments 3 and 4.

B.1 Treatment 1: When is TFT stable against ALLD?

TFT is stable against ALLD when the payoff to TFT given TFT is greater than the payoff to ALLD given TFT:

$$\Pi(TFT|TFT) > \Pi(ALLD|TFT)$$

Substituting from Fig. B.1 and simplifying yields the condition where TFT, and thus one-shot cooperation is stable:

$$\frac{c}{b} < \frac{w(1-P)}{1-wP}$$

B.2 Treatment 2: When is TFT stable against DIMAS?

TFT is stable against DIMAS when the payoff to TFT given TFT is greater than the payoff to DIMAS given TFT:

$$\Pi(TFT|TFT) > \Pi(DIMAS|TFT)$$

Substituting from Fig. B.1 and simplifying yields the condition where TFT, and thus one-shot cooperation is stable:

$$\frac{c}{b} < w(1-P)$$

B.3 Treatment 3: When is TFT stable against *direct* invasion by HGRIM?

TFT is stable against HGRIM when the payoff to TFT given TFT is greater than the payoff to HGRIM given TFT:

$$\Pi(TFT|TFT) > \Pi(HGRIM|TFT)$$

Substituting and Simplifying:

$$\frac{c}{b} < \frac{(1-P)w(w^2-w+1)}{(1-P)w^2-w+1}$$

B.4 When is TFT stable against indirect invasion by HGRIM via TF2T?

One of two conditions must be met for TFT to be stable against indirect invasion by HGRIM via TF2T. First the payoff to TFT given HGRIM could be greater than the payoff to TF2T given HGRIM:

$$\Pi(TFT|HGRIM) > \Pi(TF2T|HGRIM)$$

Substituting from Fig. B.1 and simplifying yields the condition where TFT, and thus one-shot cooperation is stable:

$$\frac{c}{b} > \frac{w}{w^2-w+1}$$

Second the payoff to TF2T given HGRIM may be greater than the payoff to TF2T given TF2T.

$$\Pi(TF2T|HGRIM) > \Pi(TF2T|TF2T)$$

By inspection of Fig. B.1 this can never be true because these payoffs are equivalent in all rounds except the first where TF2T gains a benefit against itself, but not against HGRIM..

B.5 When is TFT stable against both *direct* and *indirect* invasion by HGRIM?

From B.3 and B.4, the condition where TFT is stable against both direct and indirect invasion is:

$$\frac{w}{w^2 - w + 1} < \frac{c}{b} < \frac{(1-P)w(w^2 - w + 1)}{(1-P)w^2 - w + 1}$$

It is easy to show that $\frac{w}{w^2-w+1} > \frac{(1-P)w(w^2-w+1)}{(1-P)w^2-w+1}$ over all possible values of P, w, and $\frac{c}{b}$, therefore TFT is never stable against co-invasion by TFT and HGRIM.

B.6 Treatment 4: When is TFT stable against *direct* invasion by HTFT?

The payoff matrix for this HTFT, TFT and TF2T under complete uncertainty is shown in Fig. B.1.

TFT is stable against direct invasion by HTFT when the payoff to TFT given TFT is greater than the payoff to HTFT given TFT:

$$\Pi(TFT|TFT) > \Pi(HTFT|TFT)$$

Substituting from the matrix in Fig. B.1:

$$(b-c)\frac{1-wP}{1-w} > b + w(1-P)\frac{wb-c}{1-w^2}$$

After algebraic manipulation, TFT is stable against direct invasion by HTFT when:

$$\frac{c}{b} < \frac{w(1-P)}{1-w^2P}$$

B.7 When is TFT stable against indirect invasion by HTFT via TF2T?

Either of two conditions must be met for TFT to be stable against indirect invasion by HTFT via TF2T. First, the payoff to TFT given HTFT could be higher than the payoff to TF2T given HTFT:

$$\Pi(TFT|HTFT) > \Pi(TF2T|HTFT)$$

Substituting from the matrix in Fig. B.1:

$$-c + w(1-P)\frac{b-wc}{1-w^2} > -c + w(1-P)\frac{b-c}{1-w}$$

After algebraic manipulation, TFT is stable to invasion by TF2T in the presence of HTFT when:

$$\frac{c}{b} > w$$

Second, the payoff to TF2T given TF2T could be greater than the payoff to HTFT given TF2T.

$$\Pi(HTFT|TF2T) < \Pi(TF2T|TF2T)$$

Substituting from the matrix in Fig. B.1:

$$b + w(1-P)\frac{b-c}{1-w} < b - c + w(1-P)\frac{b-c}{1-w}$$

After algebraic manipulation, TF2T invades TFT in the presence of HTFT when:

$$-c > 1$$

This is never true, so TFT is only stable to invasion by HTFT in the presence of TF2T when:

$$w < \frac{c}{b} < \frac{wP}{1 - w^2(1 - P)}$$

It is easy to show that $w \geq \frac{wP}{1-w^2(1-P)}$ for all $P \in [0 : 1]$ and all $w \in [0 : 1)$. Therefore TFT can always be invaded by HTFT in the presence of TF2T as long as the number of rounds is not infinite.

C Effect of Increasing Partners from One to Ten

The analytic analysis presented in Table 3 indicates that, in Treatments 3 and 4, one-shot cooperation should decrease under all parameter conditions. However, Figures A.9 and A.10 show some instances of increased one-shot cooperation in these treatments, especially when the probability of a one-shot interaction (P) and the probability of an additional round given a repeated interaction (w) are high. It turns out that this cooperation is maintained as a by-product of a large amount of stochasticity in agents' payoffs due to a highly skewed distribution in rounds played. This effect was not apparent in DKCT's original simulations because, when agents can only play TFT and ALLD, the amount of repeated interaction only matters for agents' payoffs when both are playing TFT. But when there are patterns of repentant and forgiving strategies, the distribution of rounds in repeated games becomes more important (Fig. 3).

Fig. C.1A shows the probability density function (generated numerically from 10 million samples) for the number rounds per partner an agent plays when agents have only one partner and P and w are high ($P = 0.9$ and $w = 0.99$). This distribution has a very long tail. In expectation, more than 90% of agents agents play only one round in their lifetime, but one dyad per generation is expected to play more than 318 rounds - over 3.5 times as many rounds as the bottom 90% of dyads *combined*. If this dyad happens to be a pair of agents with under-performing, but cooperative strategies, like TFT in Treatments 3 and 4, the strategy would increase dramatically in the population, not because it is a better strategy, but because it was lucky enough to receive a large number of rounds. Taking the whole 10,000 generation simulation, where one dyad is expected to play more than 1225 rounds, into account makes the potential for a stochastic shock even more apparent.

This effect would be interesting if it reflected real world distributions in patterns of interaction. However, this extreme skew is an artifact of agents having only one interaction partner in their lifetime. The amount of skew (as one would expect from the Central Limit Theorem) is drastically reduced when agents interact with multiple other agents. Fig. C.1B shows the probability density function under the same parameters when when agents play with ten partners. The distribution of the number of rounds, normalized to rounds per partner, is much less skewed. Here only 34.9% of agents play one round per partner. In a given generation one agent is expected to play greater than 48 rounds/partner. Out of 10 million samples used to generate Fig. C.1B, *no* agents played over 92 rounds per partner.

Due to the Central Limit Theorem, the expectation of number of rounds per partner in each distribution is the same (i.e., $\frac{1-Pw}{1-w}$), but playing with multiple partners is more realistic, less skewed and limits the effect of stochastic shocks, and still maintains the logical structure of the game itself.

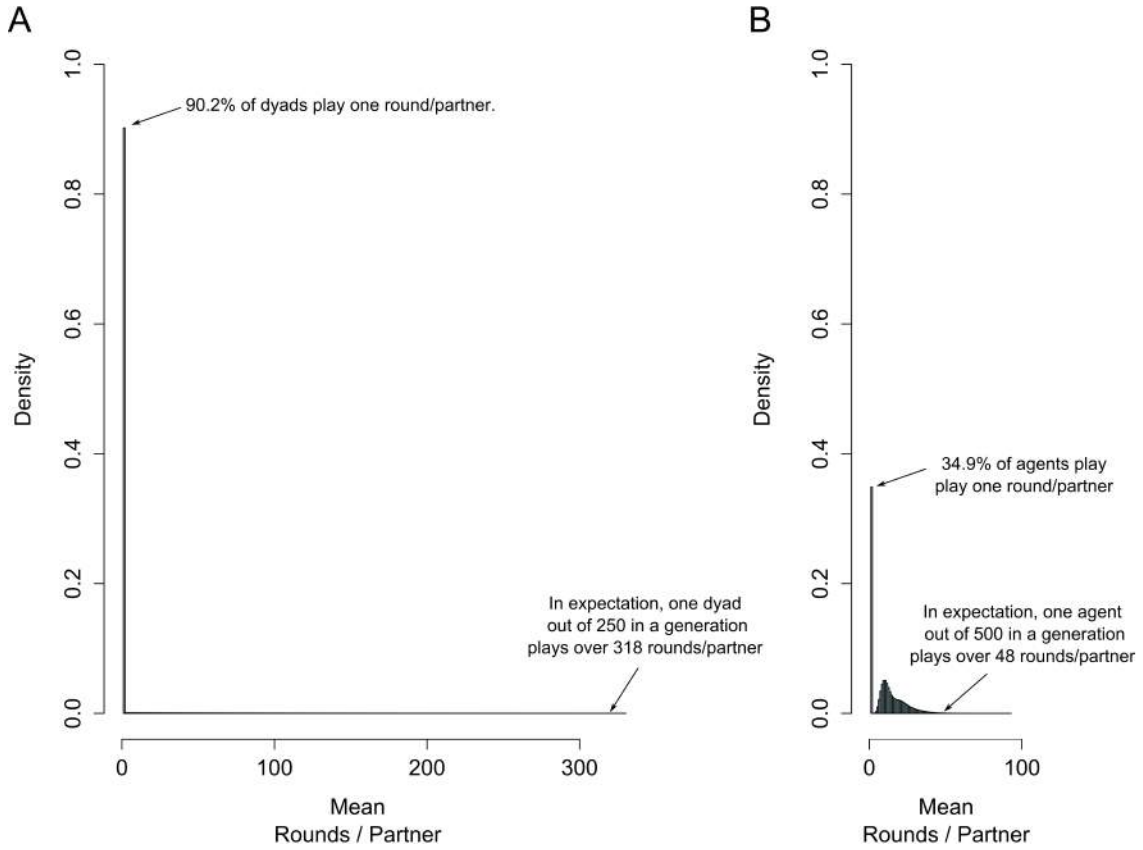

Figure C.1: **Increasing the number of partners per agent decreases the magnitude of outliers in the number of rounds played per partner. A** shows the probability density function for number of rounds played by dyads under the conditions $P = 0.1$ and $w = 0.99$. **B** shows the probability density function for number rounds of played per partner, if each agent has ten partners, under the same conditions

D Finite State Machine Representations of Strategies

Strategies from all four treatments are represented in Fig. D.1 as Finite State Automata. Initial plays of the strategy are represented by the state in the double circle (nasty strategies start with Defect and nice strategies start with Cooperate). Transition rules are represented by arrows. For example, HGRIM starts with Defect, transitions to Cooperate, and stays at Cooperate until its opponent defects. After defection by an opponent, DIMAS transitions to Defect where it stays. HTFT is similar, except that once it reaches Defect, it can transition back to Cooperate if its opponent cooperates.

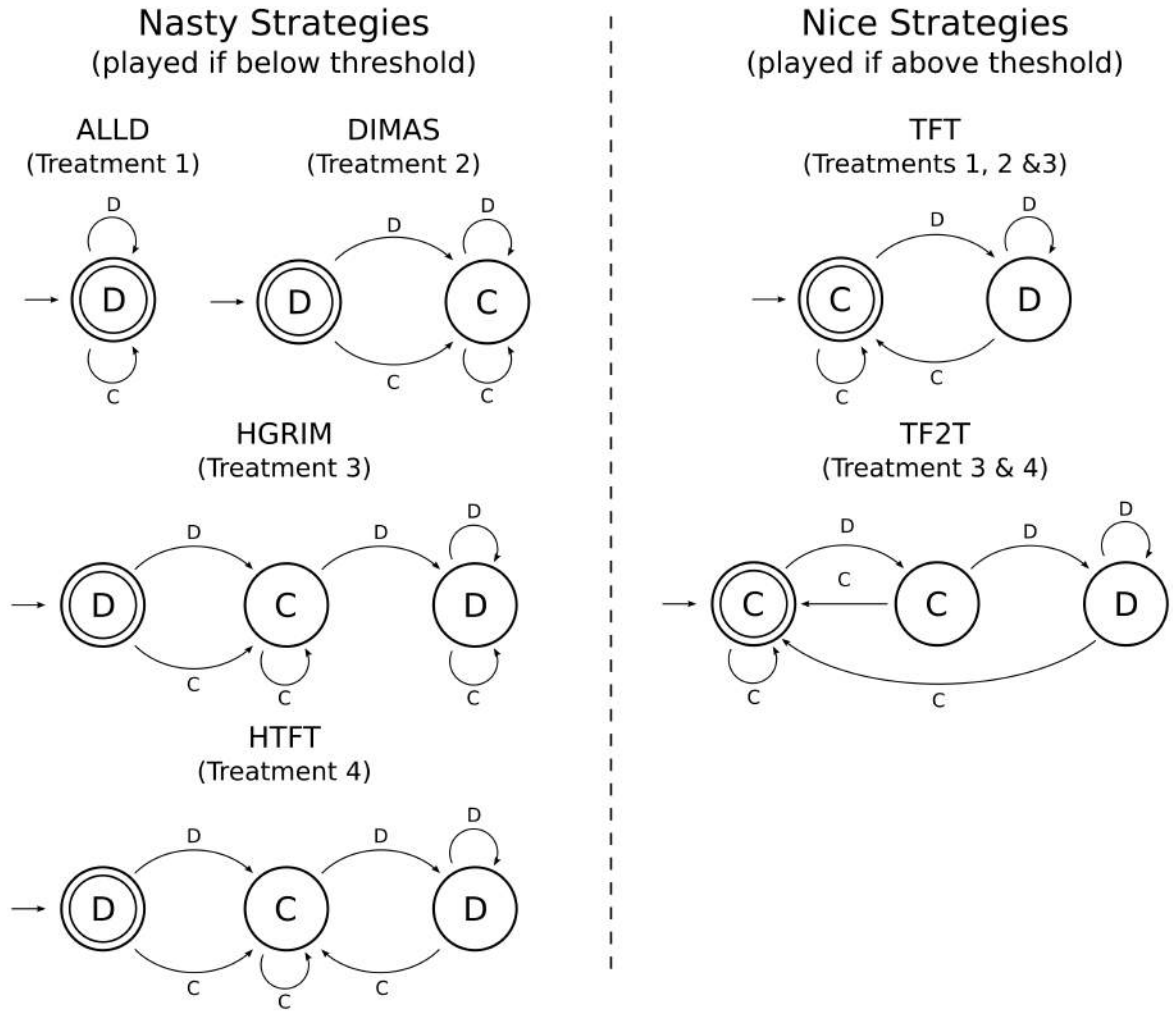

Figure D.1: **The strategies included in all four treatments represented as Moore Machines, a class of Finite State Automata.** The number of states in the minimal FSA is a measure of the complexity of a strategy. ALLD is a one-state strategy. TFT and DIMAS are two-state strategies. TF2T and HGRIM are three-state strategies.

References

Delton, A.W., Krasnow, M.M., Cosmides, L., Tooby, J., 2011. Evolution of direct reciprocity under uncertainty can explain human generosity in one-shot encounters. Proceedings of the National Academy of Sciences 108, 13335–13340.