

Akdeniz, Aslihan; van Veelen, Matthijs

Working Paper

Mutation-selection equilibria for the ultimatum game

Tinbergen Institute Discussion Paper, No. TI 2021-074/I

Provided in Cooperation with:

Tinbergen Institute, Amsterdam and Rotterdam

Suggested Citation: Akdeniz, Aslihan; van Veelen, Matthijs (2021) : Mutation-selection equilibria for the ultimatum game, Tinbergen Institute Discussion Paper, No. TI 2021-074/I, Tinbergen Institute, Amsterdam and Rotterdam

This Version is available at:

<https://hdl.handle.net/10419/248758>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

TI 2021-074/I
Tinbergen Institute Discussion Paper

Mutation-selection Equilibria for the Ultimatum Game

Aslihan Akdeniz¹
Matthijs van Veelen¹

¹ University of Amsterdam

Tinbergen Institute is the graduate school and research institute in economics of Erasmus University Rotterdam, the University of Amsterdam and Vrije Universiteit Amsterdam.

Contact: discussionpapers@tinbergen.nl

More TI discussion papers can be downloaded at <https://www.tinbergen.nl>

Tinbergen Institute has two locations:

Tinbergen Institute Amsterdam
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 598 4580

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900

Mutation-selection equilibria for the ultimatum game

Aslıhan Akdeniz^{1,2} and Matthijs van Veelen^{1,2}

¹University of Amsterdam, The Netherlands.

²Tinbergen Institute, The Netherlands.

8th August 2021

Abstract

Rand et al. (2013) present a finite population model to explain the evolution of fair behaviour in the ultimatum game. They find that mutation and selection can balance at population states that resemble human behaviour, in that responders on average evolve sizable thresholds for rejection, and proposers make proposals that on average more than meet the threshold. Their mutation however is global, and therefore biased. We show that when mutations are local, thresholds as well as proposals are considerably lower. In order to still arrive at averages in the simulations that match the averages we observe in the lab, we would have to resort to intensities of selection that are so low, that a match between the averages over time in the simulations and the averages found in experiments stops being an indication that the data are in line with what the model predicts.

1 Introduction

The ultimatum game (Güth et al., 1982) is a classic. It is played between a proposer and a responder. The proposer makes a proposal how to distribute a given amount of money between herself and the responder. The responder then accepts or rejects the proposal. In case she rejects, neither player gets any money. For simplicity, we will begin by assuming that there is a minimum step size, like cents, and that all offers can only be made in euros (or some other currency) and cents.

A natural assumption could be that when they have to make a decision, both players only look at how much money they get themselves, and always choose the alternative in which they get the most. If the responder does that, then she will accept any proposal in which she is offered more than nothing, and she will be indifferent between accepting and rejecting the proposal in which she gets 0. The proposer, then, can be assumed to anticipate that, and offer the smallest possible amount that the responder accepts (either 0 or 1 cent).

If we imagine how natural selection would act on strategies for this game – in the absence of commitment, repetition, population structure, mutations, or noise – we would come to the same conclusion. For every strategy in which the responder rejects some proposal in which she was offered a positive amount, there is a strategy that does better by accepting it. Responders therefore will evolve to accept all proposals in which they get positive amounts, while there is no selection pressure for or against accepting proposals in which responders get nothing. In reaction to this, proposers evolve to offer to the responder the smallest positive amount, or zero, if the responders accept zero too.

Yet this is not what subjects in the lab do. Proposers on average make proposals in which the responder gets around 40% of the total money amount (Oosterbeek et al., 2004). Moreover, the way rejection rates depend on the offer suggests that most responders have thresholds, or minimal acceptable offers (MAO's), that on average are a bit below what proposers, on average, offer. There is a variety of possible explanations for these findings. Commitment to rejection may have evolved in order to influence the proposals that proposers make (Akdeniz and van Veelen, 2021; Nowak et al., 2000). Rejections can also be caused by inequity aversion (Bolton and Ockenfels, 2000; Fehr and Schmidt, 1999), the evolution of which would still require an ultimate explanation, for instance based on population structure. Another explanation could be that our behaviour has evolved for a repeated version of this game, as papers reviewed in

37 Debove et al. (2016) suggest. The explanation that we explore in this paper, is
38 one based on the asymmetry between proposers and responders regarding how
39 strong the selection pressure is against the different disadvantageous mutations
40 (Rand et al., 2013).

41 Before considering mutations, it is worthwhile to discuss Gale et al. (1995),
42 which is based on the same asymmetry, but has noise in how individuals execute
43 their strategy instead of mutations. Imagine a population where proposers make
44 proposals in which they get almost all of the money, while responders are offered
45 the remainder, and in which responders moreover accept those proposals, but
46 not less. Proposers now would stand to lose much more from overasking (or
47 underoffering) than responders do from rejecting. That implies that in such
48 a population, if there is a little bit of noise on both behaviours, this hurts
49 proposers that make proposals close to the threshold of responders much more
50 than it hurts responders with thresholds close to the offer of proposers. This
51 can push the offers made by proposers upwards, which then further reduces the
52 selection against strategies that reject small offers.

53 For the same reasons, selection against disadvantageous mutations can be
54 asymmetric. In a population where proposers earn a lot and responders earn
55 little, mutations that make proposers reduce their offer below the threshold of
56 responders are genuinely costly, and will be selected away pretty fast, or pretty
57 surely, while mutations that make responders shift their threshold above what
58 proposers offer will be less costly, and linger for much longer, or have a fair
59 chance of not being weeded out. The relative abundance of mildly disadvant-
60 ageous mutations can then change the selection pressure, and, in this case, move
61 offers to responders upwards.

62 In the remainder of the paper, we will see that this is indeed possible; muta-
63 tions can indeed create a gap between average offers and average MAO's, and
64 in combination with asymmetric selection pressure, this can push both of them
65 up from 0. We will however also see that in the simulations by Rand et al.
66 (2013), the average offer is pushed up mainly for a different reason. In their
67 model, mutations are global, which has two effects. On the one hand it re-
68 duces how much the proximity of average offers and average MAO's matters for
69 the asymmetry in selection pressures. If mutations are local, creating a bit of
70 space between your offer and the average MAO helps a lot to avoid rejection. If
71 mutations are global, however, creating the same amount of space increases the
72 probability of having the proposal accepted only marginally. That makes it less
73 worthwhile to create distance between your offer and the average MAO, which

74 reduces the effect of asymmetric selection pressure. On the other hand, global
75 mutations are biased. This bias pushes offers and MAO's towards the middle
76 of the strategy space – and given that selection brings offers and MAO's down,
77 the bias pushes them back up. This bias is relatively powerful when selection is
78 weak and mutations are frequent.

79 If we make mutations local, then the bias in mutations is reduced signific-
80 antly. Below, we will see that changing from global to local mutations reduces
81 the average offers and the average MAO's more or less across the board, which
82 implies that the reduction of the bias makes offers and MAO's come down much
83 more than the local concentration of the mutants, together with the asymmetry
84 in selection pressures, makes them go back up.

85 Both with global and with local mutations, high mutation rates and low
86 intensities of selection push the offers and MAO's up. In order to get to average
87 offers and average MAO's that match those found in experiments, with local
88 mutations one therefore would have to resort to higher mutation rates or lower
89 intensities of selection compared to global mutations – or a combination of both.
90 Here, that would mean that the mutation rate would have to be unrealistically
91 high, or the intensities of selection would have to be so low, that in the dynamical
92 model, the average offer and the average MAO in the population are too variable
93 across time to be predictive (even if, when also averaged over time, the average
94 of the average offers, and the average of the average MAO's match the averages
95 in empirical data).

96 Because we do not think that a bias in mutations constitutes a good basis for
97 an explanation for behaviour in the ultimatum game, we argue that the version
98 with local mutations is preferable.

99 **2 The model in Rand et al. (2013)**

100 Rand et al. (2013) consider a finite population model, in which 100 individuals
101 play ultimatum games in both roles. Every individual has a strategy that spe-
102 cifies the offer they make in the role of proposer, and their MAO in the role of
103 responder. These offers and thresholds range from 0 to 1, in some simulations
104 in steps of $1/12$, in other simulations (including the ones we will be focusing
105 on) continuously. Each generation, every individual plays the ultimatum game
106 with every other individual, once as a proposer and once as a responder. The
107 resulting payoff is the average of the payoffs over all 99 pairings.

108 The population is updated according to a Moran process. One agent is picked
109 at random to die, and individual $i \in \{1, \dots, 100\}$ is picked with probability
110 proportional to $\exp(w\pi_i)$ to reproduce, where w is the intensity of selection,
111 and π_i is the average payoff of individual i . Mutations happen at rate u at
112 reproduction; with probability $1 - u$, the new individual inherits the strategy
113 from the reproducing individual, and with probability u , the new individual
114 carries a randomly selected strategy. If offers and MAO's vary continuously,
115 both the new offer and the new MAO are drawn from a uniform distribution on
116 $[0, 1]$; if they vary in steps of $1/12$, all possible strategies have equal probability.

117 The setup implies that there are two variables one can vary: the mutation
118 rate u , and the intensity of selection w .

119 **3 Our version(s)**

120 There is one general, inconsequential difference between their simulations and
121 ours, and that is that we use a Wright-Fisher instead of a Moran process. The
122 Wright-Fisher process is computationally more efficient, but other than that, it
123 perfectly reproduces the findings in Rand et al. (2013) for global mutations.

124 The more important, and consequential difference is that in our version,
125 mutations are local. We study two local alternatives for the local mutation
126 process.

127 **3.1 Local, co-occurring mutations**

128 In the first one, mutations on both dimensions (offer and MAO) are co-occurring,
129 as they are in Rand et al. (2013). That means that if a mutation happens, then
130 both a new offer and a new MAO are drawn. The only difference with Rand
131 et al. (2013) is that they are drawn from a local distribution, instead of a global
132 one. If the old offer is p , then the new offer is $p + \Delta p$, where Δp is drawn from
133 a uniform distribution on $[-0.1, 0.1]$. There are two exceptions. The first is
134 a result of the fact that we do not allow for values below 0. Therefore, when
135 $p + \Delta p < 0$, the new offer is 0. Similarly, we also do not allow for values over
136 1, and therefore, if $p + \Delta p > 1$, the new offer is 1. This implies that mutations
137 are unbiased for trait values in $[0.1, 0.9]$, and become a little biased if they drop
138 below 0.1 or go over 0.9 (in which case the bias is still small compared to the
139 bias with uniform mutations in Rand et al., 2013). The same procedure applies
140 to the MAO.

141 **3.2 Local, independent mutations**

142 In the second version, mutations in the offer or the MAO happen independently.
143 At any reproduction event, the offer mutates with probability u , and so does
144 the MAO. That means that with probability u^2 mutations of the offer and of
145 the MAO co-occur, and with probability $2u(1-u)$ only one of them mutates.
146 Mutations still happen locally, as described above.

147 The differences between these two versions are relatively small (see the ap-
148 pendix). Because the second version is a bit more elegant, this is the one that
149 we will use below, when we compare global and local mutations.

150 **4 Global versus local mutation**

151 The way we compare global versus local mutation will be centered around the
152 simulations reported in Figure 2 in Rand et al. (2013). These figures give an
153 answer to the question which combinations of the intensity of selection and the
154 mutation rate put the average offer and the average MAO in the range of the
155 averages in empirical findings. There are two ways to rephrase that question,
156 or to visualize the answer. The first is: for a given mutation rate, one could
157 ask how low the intensity of selection would have to be in order to get the
158 offers and MAO's up to levels found in experiments. The second is: for a given
159 intensity of selection, one could ask how high the mutation rate would have to
160 be in order to get offers and MAO's up to the levels found in experiments. We
161 will compare the different mutation processes by looking at how the answers to
162 these questions change if we switch from global to local mutations.

163 **4.1 Fixed mutation rates, decreasing intensity of selection**

164 Figure 1 represents the first way to pose the question, or the first way to visualize
165 the answer. Here we fix mutation rates, and look at the average offers and the
166 average MAO's in the dynamics for a variety of intensities of selection. Thereby
167 we do the same as in Figure 2 in Rand et al. (2013), with one change, and
168 that is that we inverted the horizontal axis. In Rand et al. (2013), the figures
169 begin with low intensities on the left and high intensities on the right. We do
170 the opposite. The reason for that is that we want to make it clear that the
171 benchmark is the subgame perfect equilibrium, with offers and MAO's close to
172 0, and that we investigate ways to arrive at dynamics with substantial average

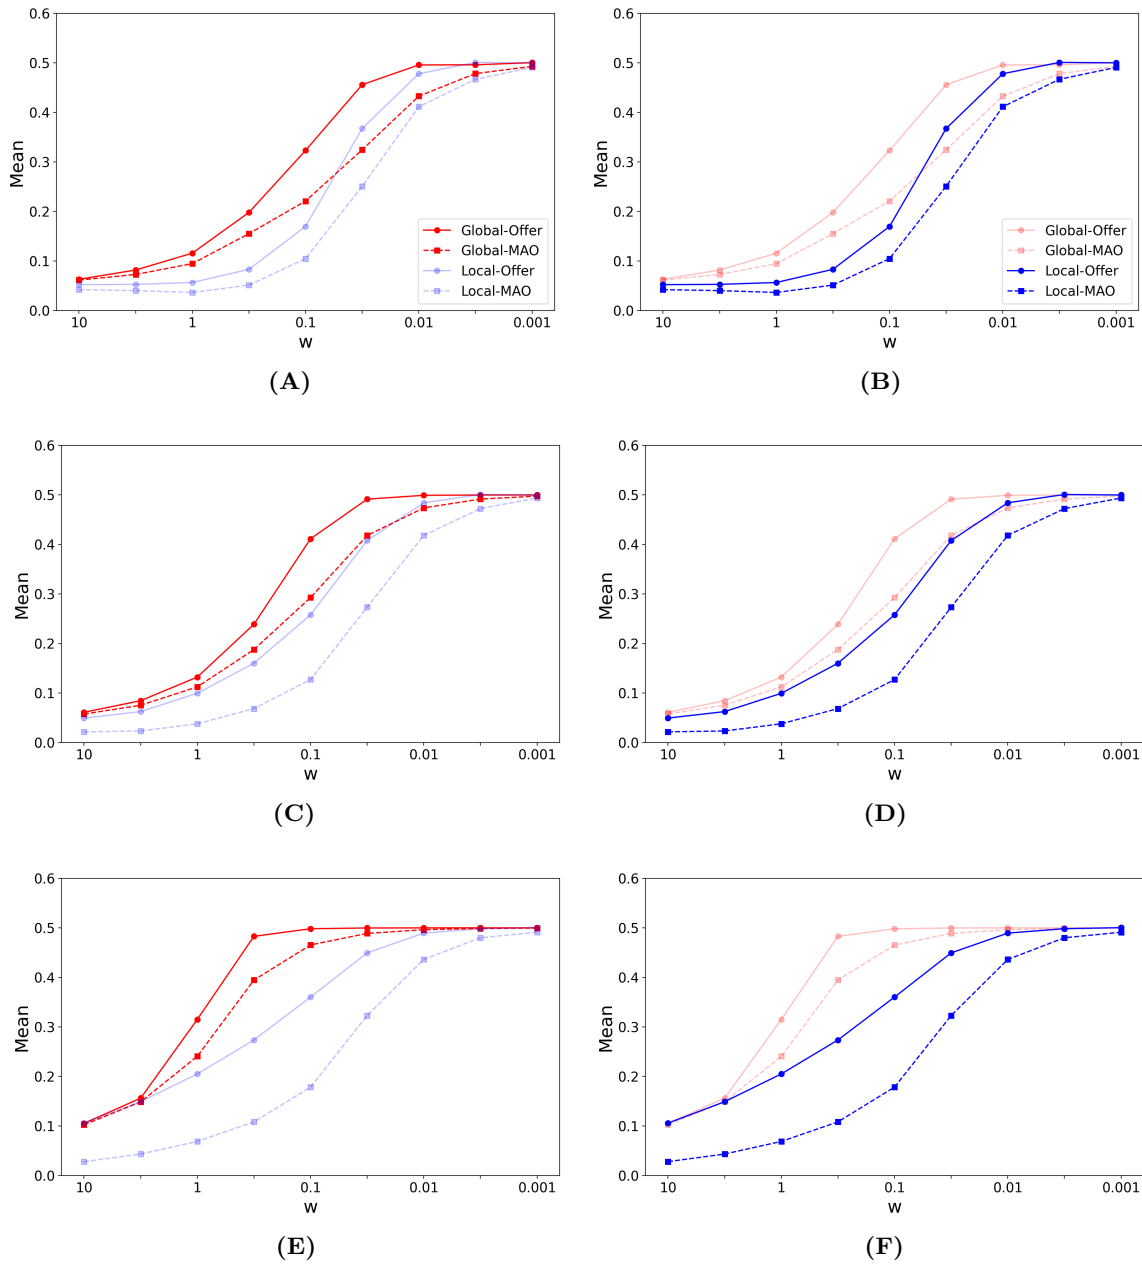


Figure 1: Average offers and MAO's at $u = 0.001$, $u = 0.01$, and $u = 0.1$. In red the average offers and MAO's with global mutation, as in Rand et al. (2013), and in blue the same, but for local, independent mutations. Both the average offers and the average MAO's are lower with local mutations. In order to get offers, or MAO's, up to similar levels with local and global mutations, one would have to move to lower intensities of selection with local mutations.

173 offers and substantial average MAO's. Reducing the intensity of selection is one
174 such way.

175 From the simulations, we learn that there are two important differences
176 between global and local mutations. The first is that with the bias significantly
177 reduced, the average offers and MAO's stay low for longer, and start getting into
178 the range of the experiments later on in the sequence of ever lower intensities of
179 selection. Here it is important to note that on the right hand side of the graph,
180 with low intensities of selection, average offers and average MAO's end up at
181 0.5 in both versions. The reason why this eventually happens for any possible
182 modeling choice, is that 0.5 is halfway the parameter space, and therefore this
183 must be the average over time in the limit of weak selection, where relative
184 payoffs stop mattering. In Section 5 we discuss this in more detail, including
185 reasons why the dynamics lose predictive power as selection becomes weaker.

186 The second remarkable difference is especially visible in Fig. 1E and 1F for
187 $u = 0.1$. Here we see that on the left side of the graph, at high intensities of
188 selection, there is almost no gap between offer and MAO for global mutation,
189 while there is a visible gap there for local mutations. The latter is consistent
190 with the idea that, given that there is a consistent inflow of mutations, proposers
191 benefit from creating some space between their offer and the average MAO in
192 the population, and responders benefit from creating some space between their
193 MAO and the average offer in the population. Both of these pressures contribute
194 to the gap, but asymmetrically, because at low average MAO's, proposers have
195 more to lose from rejections than responders do from rejecting. All else equal,
196 this makes proposers want to create more distance than responders do, and that
197 can result in both averages moving up a little.

198 If we then start on the left hand side of the graph, and move a little to the
199 right, then first the effect of reducing the intensity of selection is that this keeps
200 mutants around for longer, creating a wider distribution of offers and MAO's,
201 thereby selecting strategies that on average keep more distance – which in turn
202 leads to higher offers and MAO's due to the asymmetry in selection pressures.
203 Later on, when selection gets even weaker, the gaps gets smaller again. Why
204 this happens, is discussed in Section 5.

205 With global mutations, on the other hand, there is hardly any gap at first.
206 Here, the moving up of the offers and MAO's as selection gets weaker is pre-
207 dominantly the result of the bias getting countered less by selection. Only if the
208 bias has a sufficiently large effect, because of further reduced intensities, do we
209 see a divergence between offers and MAO's.

210 Both differences indicate that the dynamics for local mutations capture the
211 selection effect that we are interested in. This selection effect is driven by
212 asymmetry in fitness effects, and not by bias in mutations.

213 4.2 Fixed intensity of selection, increasing mutation rates

214 In Figure 2 we fix intensities of selection, and look at the average offers and
215 MAO's in the dynamics for a variety of mutation rates. The simulations suggest
216 the same; if we go from global to local mutations, and thereby reduce the bias,
217 the average offers and the average MAO's go down. In Fig. 2E and 2F we
218 see that they stay low for longer, and that even at the maximum mutation
219 rate, where everybody mutates to a meaningful element of the trait space, only
220 the average offer and the average MAO with global mutations, and not local
221 mutations, reach the averages from experiments. In Fig. 2C and 2D they do,
222 but with local mutations, it requires unreasonably large mutation rates. In Fig.
223 2A and 2B all averages are smaller deviations from 0.5. In Section 5 we will
224 describe why the dynamics for such low intensities of selection lose predictive
225 power.

226 With global mutations, both average offers and average MAO's inevitably get
227 to 0.5 as mutation rates increase. The reason is that when mutations are global,
228 then at $u = 1$, when both the offer and the MAO mutate at every reproduction
229 event, it becomes irrelevant who is reproducing. The parents therefore stop
230 passing on any genetic information; every new individual is a mutant, and all
231 mutants are drawn from the same distribution, regardless of what the parents
232 are. Therefore, at $u = 1$, on both dimensions, the population at any point in
233 time just becomes a collection of independent random draws from $[0, 1]$.

234 With local mutations, average offers and average MAO's do not necessarily
235 get to 0.5 as the mutation rate increases. In this case, parents still pass on ge-
236 netic information, because even if everyone mutates, these mutations are drawn
237 from a distribution that is centered around the trait value of the parent. That
238 makes it possible for the average offer and the average MAO to stay below 0.5,
239 even if the mutation rate is 1.

240 5 Predictions for weak selection

241 In the previous section, we have seen that both with global and with local
242 mutations, high mutation rates and low intensities of selection push average

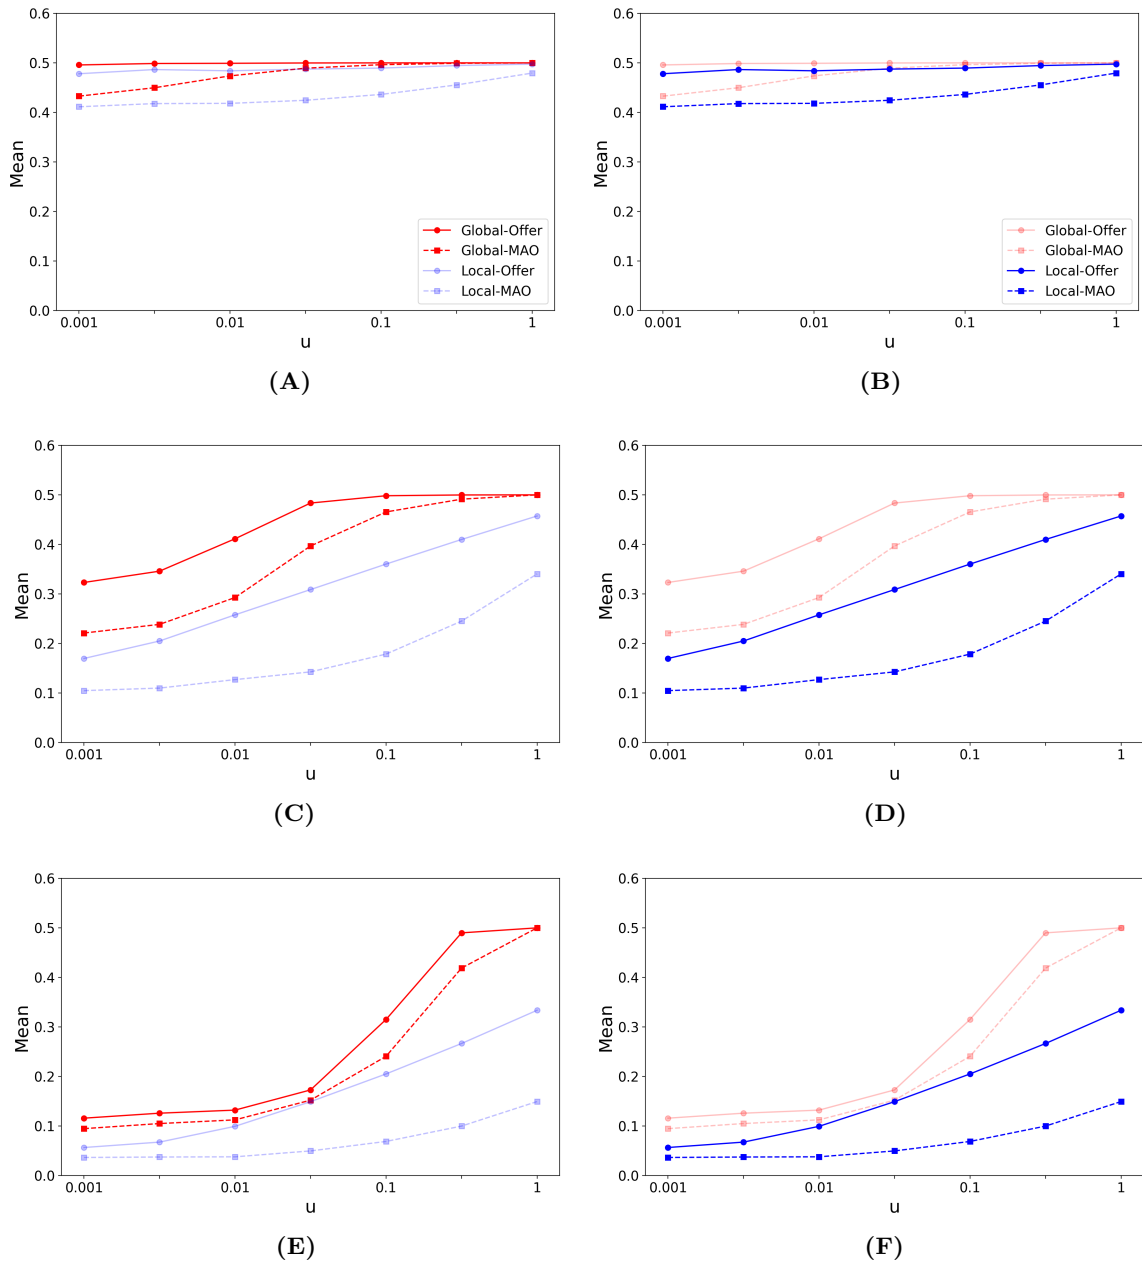


Figure 2: Average offers and MAO's at $w = 1$, $w = 0.1$, and $w = 0.01$. In red the average offers and MAO's with global mutation, as in Rand et al. (2013), and in blue the same, but for local, independent mutations. For $w = 0.01$ everything gets close to 0.5. For $w = 0.1$ one can get to the averages observed in experiments, but with local mutations it requires very high mutation rates. For $w = 1$ and local mutations, even a mutation rate of $u = 1$ is not high enough.

243 offers and the average MAO's up from 0. There are however limitations to how
244 far these parameters can be pushed, and still constitute a meaningful prediction.
245 To see why, we look at the dynamics for low intensity of selection.

246 When the intensity of selection is 0, the dynamics in the model by Rand et al.
247 (2013) become symmetric, in the sense that any transition from one population
248 state to the other is equally likely as its mirror image. More precisely, if p_i
249 denotes the offer of player i in the role of proposer, and q_i is the MOA of
250 player i in the role of responder, then a population state is characterized by
251 vectors $\mathbf{p} = [p_1, \dots, p_N]$ and $\mathbf{q} = [q_1, \dots, q_N]$, where N is the population size.
252 Symmetry means that a transition from population state (\mathbf{p}, \mathbf{q}) to population
253 state $(\mathbf{p}', \mathbf{q}')$ is equally likely as its mirror image, going from population state
254 $(\mathbf{1} - \mathbf{p}, \mathbf{1} - \mathbf{q})$ to population state $(\mathbf{1} - \mathbf{p}', \mathbf{1} - \mathbf{q}')$, where $\mathbf{1}$ is a vector of 1's.
255 This symmetry implies that if we average the population states over time, we will
256 find a symmetric distribution. The average offer over this distribution therefore
257 will be 0.5, and the average MAO will also be 0.5, and both of these are a
258 consequence of the fact that 0.5 is halfway the strategy set that the population
259 moves around in – with probabilistic symmetry.

260 All of this implies that the fact that it is possible to get the average offer or
261 the average MAO up to any value between 0 and 0.5 by choosing a sufficiently
262 low intensity of selection is not necessarily something that reflects anything to do
263 with selection. Selection pulls both of them down, and if one reduces selection
264 ever more, one can reduce how much both are dragged down. The fact that one
265 can get them to average at a values arbitrarily close to 0.5 by almost eliminating
266 selection, however, is a somewhat arbitrary result of the shape of the strategy
267 set, and not of what selection does to the strategies in it.

268 In order to see other reasons why the model with weak selection poorly
269 predicts the empirical observations, we will look in some detail simulation runs
270 for s few selected parameter combinations.

271 **5.1 Weak selection, low mutation rate, global mutation**

272 The left hand side of Fig. 3 depicts a few aspects of a run with global mutation,
273 a low intensity of selection ($w = 0.001$), and a low mutation rate ($u = 0.001$).
274 Panel A shows how the average offer and the average MAO change over time
275 for a part of a simulation run. Here it is important to note that the average
276 offer and the average MAO move around quite a bit, and not in overwhelming
277 synchrony (in fact, there is hardly any correlation). Panel C gives a snapshot

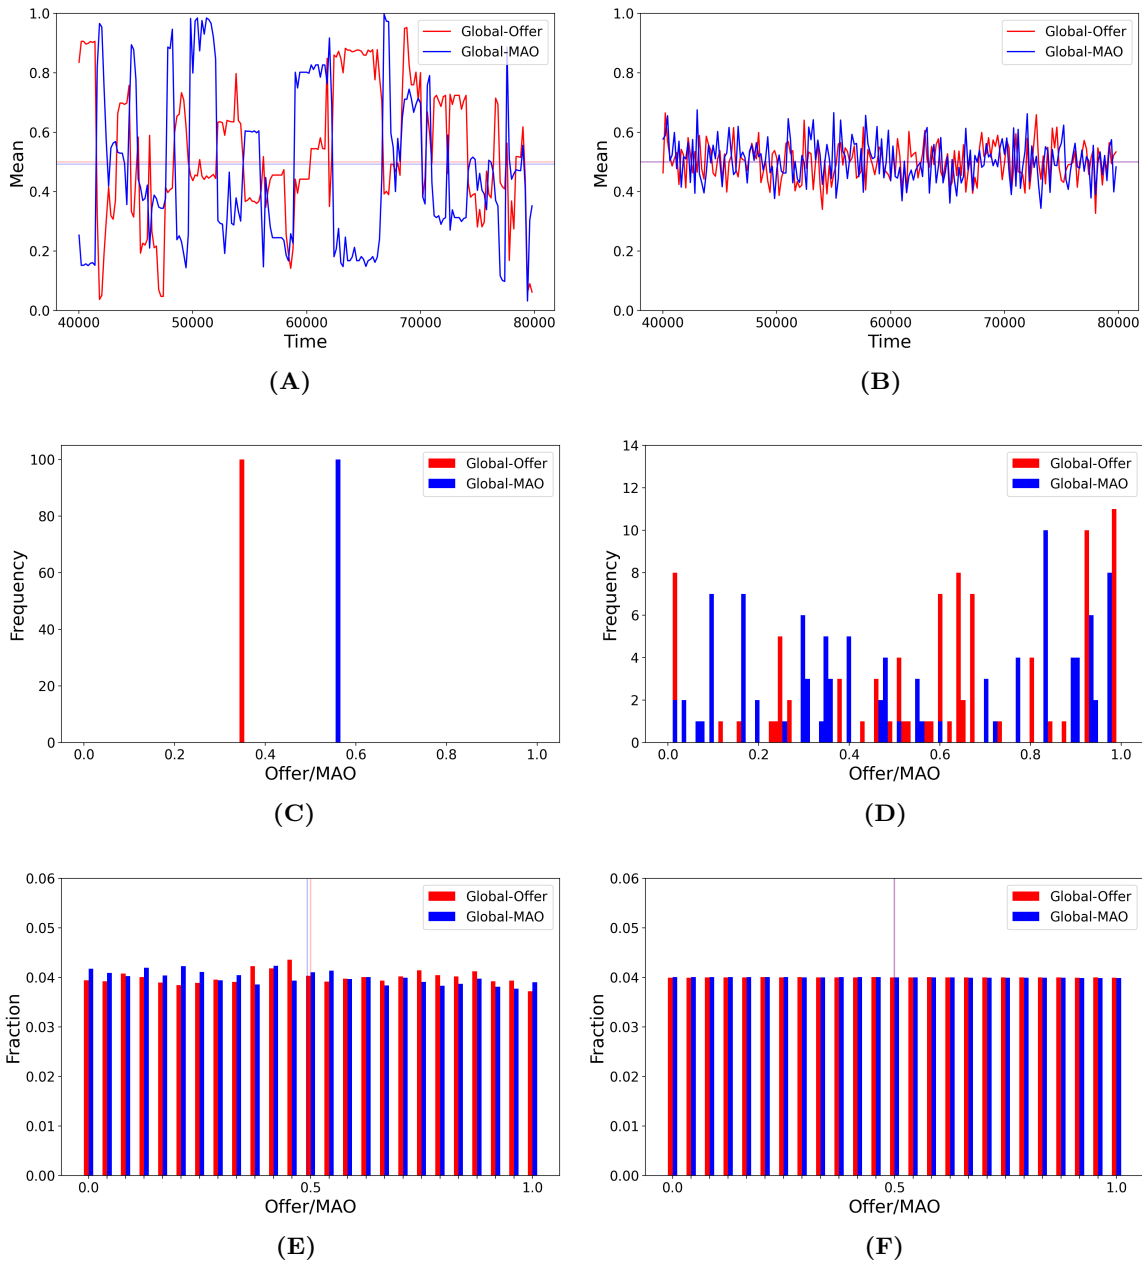


Figure 3: Global mutation, $w = 0.001$, and $u = 0.001$ (left) and $u = 0.1$ (right). The top panels give the average offer and MAO over time for a part of the run. The middle panels give the distribution of strategies at some random moment in the simulation run. The bottom panels give the average distribution over time, where we collected strategies within intervals of length 0.05.

278 of the distribution at some moment in time, and here we find both traits to
279 be at fixation, as is expected to be the case for most of the time with such a
280 low mutation rate. Panel E averages these distributions (like the one given in
281 panel C) across time, producing the average distribution over time. As is to
282 be expected, this is quite close to the uniform distribution on $[0, 1]$, which is
283 the distribution that all mutants come from. The average offer and the average
284 MAO of the average distribution (both are averages of averages) are horizontal
285 lines in panel A, and vertical lines in panel E.

286 The fact that the average offer and the average MAO move around quite a
287 bit over the course of a run limits the predictive power of the model for this com-
288 bination of low intensity of selection and a low mutation rate. Any average that
289 we find in experiments would be close to the average in the simulations at some
290 points in time, but it would be far away from the averages that the simulations
291 produce at many other points in time. Also, at most points in time, there is not
292 much variation; the variation in panel E is generated by the variability across
293 time, not by the variation at any moment in time. The prediction of this model
294 therefore is that we should observe a close to monomorphic population, where
295 the probability with which we would observe a certain average is the result of a
296 draw from the uniform distribution – which is not the same as observing that
297 whole distribution in one population at the same moment in time. The fact
298 that the MAO of everyone in the population is regularly also above the offers
299 of everyone in the population (almost 50% of the time) also implies that if we
300 really believe in weak selection, we should also conclude that if we now find
301 the average MAO to be below the average offer, then this is just a coincidence,
302 and it could also have been the other way around. Also it would make it very
303 unlikely that between different populations they would correlate, and that the
304 first always turns out to be below the second (Henrich and Boyd, 2001; Henrich
305 et al., 2001; 2006).

306 All of these observations imply that this model does not have much pre-
307 dictive power when selection is weak, and the mutation rate is moderate to
308 low. Therefore, we have to be careful not to draw conclusions on the basis of a
309 comparison of a distribution of offers and MAO's that we find at a given mo-
310 ment in time in an experiment, and the average distribution over time from the
311 dynamics in the model.¹

¹In order to have an indication how stable or unstable the distributions are over time, we can calculate the variance in average offers, or MAO's, over time. If \bar{p}^t is the average offer in

312 5.2 Weak selection, high mutation rate, global mutation

313 The right hand side of Fig 3 depicts the same aspects for a run, also with global
314 mutation, and also with a low intensity of selection ($w = 0.001$), but with a
315 high mutation rate ($u = 0.1$). Here, the averages in the population do not
316 move around as much, and the shape of the distribution at any point in time is
317 relatively close to the distribution of the inflow of mutants, which is a uniform
318 distribution over $[0, 1]$. Given the low intensity of selection, this makes sense.
319 With much less variability over time, this produces a much sharper prediction:
320 the distribution should be close to uniform on $[0, 1]$ at all times. This however
321 does not match the empirical evidence either, because the distributions that we
322 find in experiments typically are not that close to uniform. Moreover, as before,
323 the average offer and the average MAO move close to independently, and this
324 does not predict the average offer to be above the average MAO.

325 5.3 Weak selection, local mutation

326 Most overall properties of the dynamics are similar with local mutations. With
327 weak selection and low mutation rates, the populations are typically also close
328 to monomorphic. Over time, they also move around quite a bit, but the muta-
329 tions being local makes the average move around in much smaller steps, and
330 therefore also much slower (see the left hand side of Fig. 4). The distribution
331 over time is not the same as the "distribution that all mutants come from",
332 because with local mutation, there is no such constant mutant distribution. Be-
333 cause the average is a random walk, restricted to $[0, 1]$, the distribution still
334 ends up looking like a uniform distribution with some deviations at and close
335 to the boundaries, where it spends some extra time. For weak selection and
336 high mutation rates, the mutations being local gives random effects much more
337 amplitude, but otherwise, also these simulations are similar to the ones with
338 global mutations.

the population at time t , and $\bar{p} = \frac{1}{T} \sum_{t=1}^T \bar{p}^t$, then

$$\frac{1}{T} \sum_{t=1}^T (\bar{p}^t - \bar{p})^2$$

is the variance. Simulations with a high variance have little predictive power. The variance of proposals calculated from a run with $u = 0.001$ and $w = 0.001$ is around 0.06 for global and 0.12 for local mutations. The variance of MAO's is around 0.07 for global and 0.08 for local mutations. As a benchmark, one can compare these to the variance if one just randomly draws an offer or MAO from a uniform distribution on $[0, 1]$ every period. In that case, the variance is $\int_0^1 (x - 0.5)^2 = \frac{1}{3} \left[(x - 0.5)^3 \right]_0^1 = \frac{1}{12} \approx 0.083$.

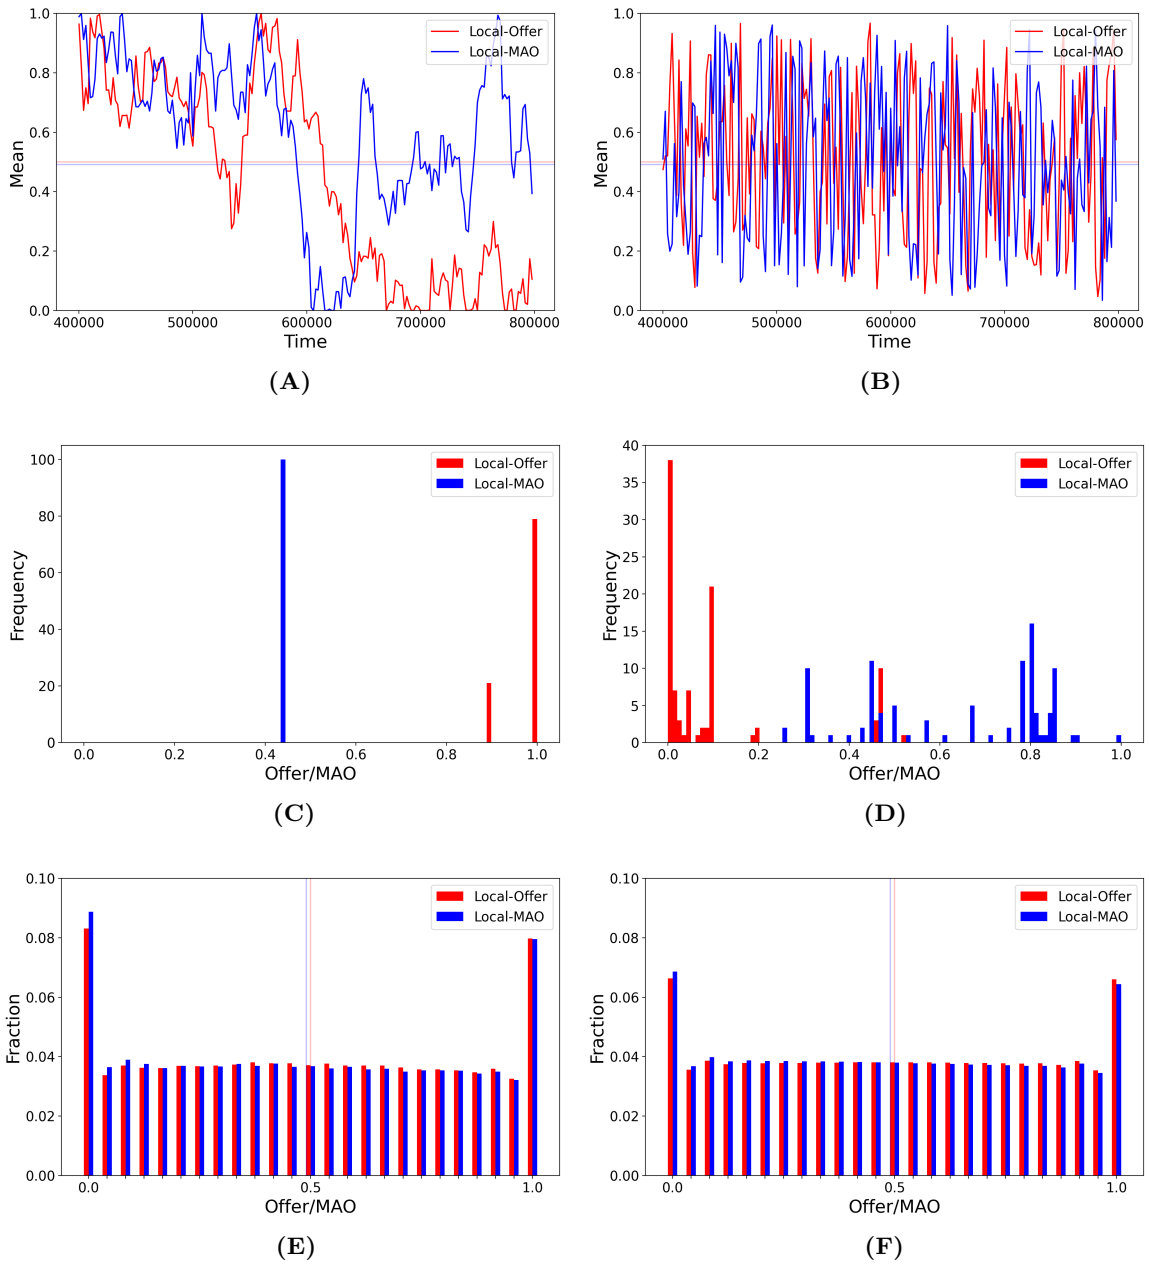


Figure 4: Local mutation, $w = 0.001$, and $u = 0.001$ (left) and $u = 0.1$ (right). The top panels give the average offer and MAO over time for a part of the run. The middle panels give the distribution of strategies at some random moment in the simulation run. The bottom panels give the average distribution over time, where we collected strategies within intervals of length 0.05.

339 6 Discussion

340 We have shown that, in a mutation-selection model of the ultimatum game,
341 changing from global to local mutations reduces the bias in the mutations. The
342 dynamics with local mutations moreover capture the effect of the difference
343 between how costly the “mistakes” of proposers’ under-offering versus respon-
344 ders’ rejecting are. The change from global to local mutations comes with lower
345 average offers and lower average MAO’s. That means that in order to get average
346 offers and average MAO’s up to the levels found in experiments, one would
347 have to resort to unrealistically high mutation rates – where mutations moreover
348 always result in meaningful mutant strategies – or intensities of selection that
349 are so low, that the offers and MAO’s, when averaged over time, become poor
350 indicators of what the model would predict at any moment in time.

351 There are a few ways to deal with the mismatch between what is found in lab
352 experiments in the ultimatum game and what evolves in the mutation-selection
353 model with local mutations. One is to accept that this model would not predict
354 the behaviour in WEIRD populations, but try consider non-WEIRD populations
355 that show lower offers and lower MAO’s. Here it is important to realize that
356 the mutation rates required are still rather high, given that a mutation here
357 means that the mutant is a meaningful strategy, and that the game behaviour
358 among non-WEIRD societies is also highly variable (Henrich et al., 2005; 2006).
359 Another way to deal with the mismatch between the empirical evidence and
360 the theoretical predictions is to model the genetics underlying the behaviour
361 differently. It is possible that, with alternative specifications, one would not
362 need excessively high mutation rates. With this in mind, one could also argue
363 that the quantitative predictions from the model are not to be taken as seriously
364 as the original paper does. It seems that the middle ground might have to be
365 somewhere between recognizing (i) that the asymmetry in costliness of mistakes
366 matters, but (ii) that getting close to fair offers of 50% on the basis of cost
367 asymmetries fails short of explaining the evolution of fairness in the ultimatum
368 game.

369 Another conclusion could be that mutations are not the (only) reason why
370 humans reject unequal proposals, and that the answer lies in commitment (Ak-
371 deniz and van Veelen, 2021; Frank, 1988). In the ultimatum game, if the re-
372 sponder can commit to rejecting unfair offers, and if proposer knows this, then
373 it will be in the proposer’s own best interest to accommodate this, and propose
374 a fair split of the money. Committing to rejecting disadvantageous proposals

375 can therefore help the responder by inducing the proposer to make more gen-
376 erous proposals. Akdeniz and van Veelen (2021) review the explanations for
377 the evolution of fairness in the ultimatum game that are commonly studied in
378 the literature, such as population structure and repeated interactions; and ar-
379 gue how commitment can better explain a wide range of empirical phenomena,
380 including a preference for fairness in the ultimatum game.

381 References

- 382 Ashkan Akdeniz and Matthijs van Veelen. The evolution of morality and the
383 role of commitment. *Evolutionary Human Sciences*, pages 1–53, 2021.
- 384 Gary E Bolton and Axel Ockenfels. ERC: A theory of equity, reciprocity, and
385 competition. *American Economic Review*, 90(1):166–193, 2000.
- 386 Stephane Debove, Nicolas Baumard, and Jean-Baptiste André. Models of the
387 evolution of fairness in the ultimatum game: a review and classification. *Evol-*
388 *ution and Human Behavior*, 37(3):245–254, 2016.
- 389 Ernst Fehr and Klaus M Schmidt. A theory of fairness, competition, and co-
390 operation. *Quarterly Journal of Economics*, 114(3):817–868, 1999.
- 391 Robert H Frank. *Passions Within Reason: The strategic role of the emotions*.
392 WW Norton & Co, 1988.
- 393 John Gale, Kenneth G Binmore, and Larry Samuelson. Learning to be im-
394 perfect: The ultimatum game. *Games and Economic Behavior*, 8(1):56–90,
395 1995.
- 396 Werner Güth, Rolf Schmittberger, and Bernd Schwarze. An experimental ana-
397 lysis of ultimatum bargaining. *Journal of Economic Behavior & Organization*,
398 3(4):367–388, 1982.
- 399 Joseph Henrich and Robert Boyd. Why people punish defectors: Weak con-
400 formist transmission can stabilize costly enforcement of norms in cooperative
401 dilemmas. *Journal of Theoretical Biology*, 208(1):79–89, 2001.
- 402 Joseph Henrich, Robert Boyd, Samuel Bowles, Colin Camerer, Ernst Fehr, Her-
403 bert Gintis, and Richard McElreath. In search of homo economicus: behavi-
404 oral experiments in 15 small-scale societies. *American Economic Review*, 91
405 (2):73–78, 2001.
- 406 Joseph Henrich, Robert Boyd, Samuel Bowles, Colin Camerer, Ernst Fehr, Her-
407 bert Gintis, Richard McElreath, Michael Alvard, Abigail Barr, Jean Ens-
408 minger, et al. “economic man” in cross-cultural perspective: Behavioral ex-
409 periments in 15 small-scale societies. *Behavioral and Brain Sciences*, 28(6):
410 795–815, 2005.

- 411 Joseph Henrich, Richard McElreath, Abigail Barr, Jean Ensminger, Clark Bar-
412 rett, Alexander Bolyanatz, Juan Camilo Cardenas, Michael Gurven, Edwins
413 Gwako, Natalie Henrich, et al. Costly punishment across human societies.
414 *Science*, 312(5781):1767–1770, 2006.
- 415 Martin A Nowak, Karen M Page, and Karl Sigmund. Fairness versus reason in
416 the ultimatum game. *Science*, 289(5485):1773–1775, 2000.
- 417 Hessel Oosterbeek, Randolph Sloof, and Gijs Van De Kuilen. Cultural dif-
418 ferences in ultimatum game experiments: Evidence from a meta-analysis.
419 *Experimental Economics*, 7(2):171–188, 2004.
- 420 David G Rand, Corina E Tarnita, Hisashi Ohtsuki, and Martin A Nowak. Evol-
421 ution of fairness in the one-shot anonymous ultimatum game. *Proceedings of*
422 *the National Academy of Sciences*, 110(7):2581–2586, 2013.

Appendix A: Co-occurring versus independent mutations

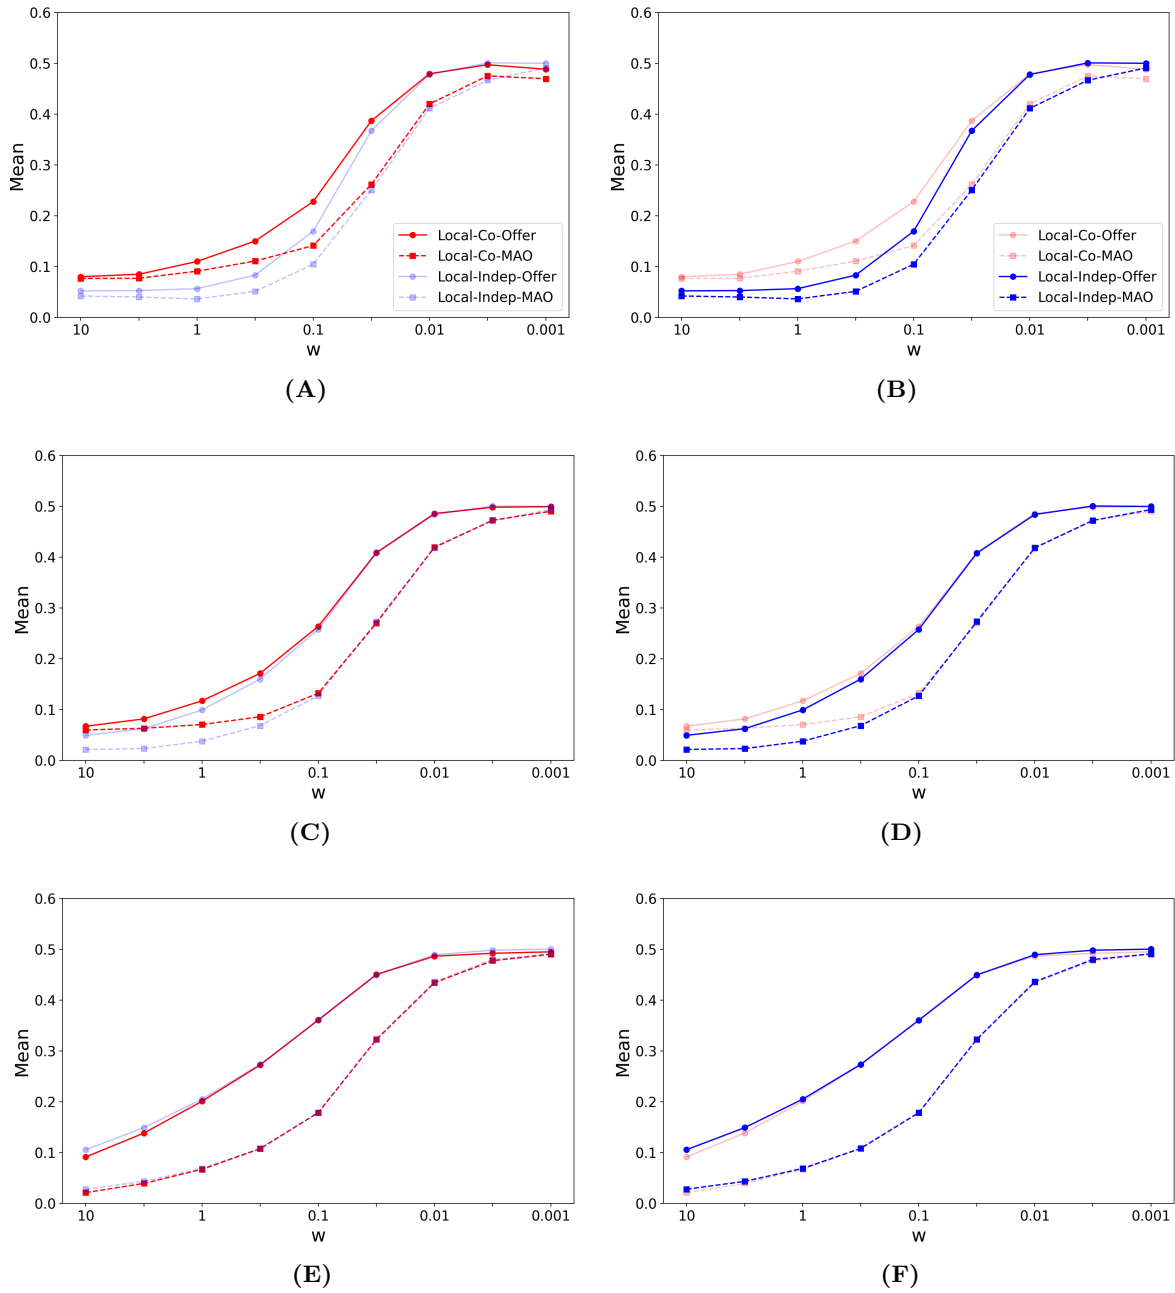


Figure A.1: Average offers and MAO's at $u = 0.001$, $u = 0.01$, and $u = 0.1$. In red the average offers and MAO's with local, co-occurring mutations, and in blue the same, but for local, independent mutations.