

Why human societies adopt rigid moral rules: the efficiency–robustness trade-off

Julien Lie-Panis^{a,1}, Léo Fitouchi^b, Nicolas Baumard^c, and Jean-Baptiste André^c

This manuscript was compiled on March 23, 2026

Humans are capable of remarkably flexible moral judgment. Yet societies rely on rigid rules—obligations and prohibitions that apply categorically, even when case-by-case reasoning could yield better outcomes. Why would a species capable of flexibility bind itself to rigid rules? We propose that rigid rules arise as social technologies for managing ambiguity about noncooperation. People often have legitimate reasons for failing to cooperate, yet those reasons are typically opaque to observers, allowing opportunists to disguise selfishness as justified hardship. We formalize this idea with an evolutionary game-theoretic model. Two cooperative equilibria emerge: a flexible norm that accommodates legitimate excuses but is vulnerable to exploitation, and a rigid norm that closes this loophole by mandating cooperation even when inefficient. Comparing these equilibria reveals an efficiency–robustness trade-off: flexibility maximizes welfare when trust is secure, whereas rigidity preserves cooperation when trust is fragile. This explains why rigid rules prevail in low-trust settings—interactions with strangers, formal institutions, or tight societies—while flexibility is more common in high-trust contexts.

morality | deontology | cooperation | norms | evolutionary game theory

“Please keep off the grass—unless you have a good reason.” This sign, once posted on the lawn of an Oxford college, represents a flexible approach to collective action. Rather than enforcing a categorical prohibition (“Keep off the grass!”), it invited individuals to exercise judgment, distinguishing between mere convenience (avoiding a short detour) and genuine need (rushing to the hospital). Rigidly staying off the grass would surely keep the lawn immaculate, but at needless cost. By allowing people to weigh the costs and benefits of every detour, the rule achieved its goal efficiently, protecting the grass while permitting legitimate exceptions.

Humans routinely exercise this kind of context-sensitive moral reasoning, adjusting their judgments to the fine-grained parameters of each situation (1–11). When waiting in line, we readily distinguish between someone who cuts out of impatience and someone who cuts because they are diabetic and urgently need sugar (12). This capacity for flexible moral judgment emerges by middle childhood (13–15), suggesting that it is a central feature of human cognition.

Yet few social rules make use of this human flexibility. Across societies—from small-scale foraging bands (16, 17) to large nation-states (18)—cooperation rests on standardized rules that hold regardless of context. In industrialized societies, drivers must obey speed limits and stop at red lights even when roads are empty and conditions perfectly safe. Laws prohibit stealing others’ property regardless of need, although people readily judge theft to avert starvation as more excusable than theft for greed (19–21; cf. *Southwark London Borough Council v Williams* 22). In many subsistence societies, elaborate sharing rules divide a kill by fixed roles—the spotter, the first striker, the harpoon owner—irrespective of who contributed most that day (23–27). Across these settings, rules apply categorically, even when taking context into account could make cooperation more efficient.

Significance

Human moral life is governed by rigid rules. Across societies, people rely on categorical prohibitions and duties—do not lie, do not steal, always keep your promises—even when they recognize compelling reasons to make exceptions. Using a mathematical model, we show that rigid rules can protect cooperation when motives are opaque. Flexible rules allow justified exceptions but also make it easier for cheaters to hide behind plausible excuses. This creates a trade-off between maximizing welfare and protecting cooperation from those who would exploit excuses to justify selfishness. Our account explains why rigid rules are common in low-trust settings, such as with strangers or in rule-bound societies, and why many moral judgments feel like non-negotiable duties rather than cost-benefit calculations.

Author affiliations: ^aMax Planck Research Group Dynamics of Social Behavior, Max Planck Institute for Evolutionary Biology, 24306 Plön, Germany; ^bDepartment of Social and Behavioral Sciences, Institute for Advanced Studies in Toulouse, Toulouse School of Economics, University of Toulouse Capitole, 31080 Toulouse, France; ^cInstitut Jean Nicod, Département d’études cognitives, Ecole normale supérieure, Université PSL, EHESS, CNRS, 75005 Paris, France

J.L.P., L.F., N.B., and J.B.A. designed research; J.L.P. performed research; J.L.P. and J.B.A. analyzed data; and J.L.P., L.F., N.B., and J.B.A. wrote the paper.

The authors declare no competing interest.

¹To whom correspondence should be addressed. E-mail: jliep@protonmail.com.

123 This contrast between flexible minds and rigid rules
124 presents a puzzle: why would a species of flexible
125 cooperators bind itself to categorical rules that demand
126 cooperation even when it is inefficient?

127 We propose that rigid norms persist because they
128 solve a problem intrinsic to reputation-based cooperation:
129 opacity about why people fail to cooperate.

130 To illustrate, imagine two roommates, Alice and Bob,
131 trying to divide the household chores. They could adopt
132 a flexible arrangement, letting each clean when they
133 judge, in good faith, that it's their turn. This promises
134 efficiency: there's no point in requiring someone to scrub
135 the bathroom when they're sick or buried in work—the
136 cost to them outweighs the benefit of a cleaner apartment.

137 But flexibility creates moral wiggle room: ambiguity
138 that an opportunist can exploit (e.g., 28, 29). If Bob
139 walks into a cluttered kitchen and Alice reports being over-
140whelmed with work, he cannot tell whether this reflects
141 genuine hardship or a convenient excuse. That opacity
142 shifts Alice's incentives by giving her cover to exaggerate
143 her workload when she feels lazy. Flexibility thus opens
144 the door to strategic defection, leaving opportunists room
145 to exploit reputational ambiguity.

146 A rigid rule closes this loophole. Suppose chores are
147 assigned on a fixed schedule—Alice cleans on Sundays and
148 Bob on Thursdays. Failing to clean on one's scheduled
149 day is then unambiguously a defection. The rigid rule
150 draws a bright line between cooperation and defection,
151 eliminating the ambiguity that opportunists depend on.

152 We formalize this idea with a model of reputation-based
153 cooperation. In the model, individuals sometimes face
154 genuine hardship—illness, heavy workload, or personal
155 crises—that makes cooperation socially inefficient. Be-
156 cause observers cannot reliably verify these circumstances,
157 actors have room to disguise selfish behavior as justified
158 hardship.

159 We identify two distinct cooperative equilibria. Under
160 a flexible norm, cooperation accommodates context:
161 individuals who appear to face hardship retain others'
162 trust. Under a rigid norm, by contrast, cooperation is
163 demanded categorically—any failure to cooperate leads
164 to loss of trust, regardless of apparent circumstances.
165 Comparing these equilibria reveals a fundamental trade-off
166 between efficiency and robustness. Flexibility is efficient,
167 allowing individuals to opt out when cooperation would
168 be genuinely costly, but it leaves observers vulnerable
169 to strategic defection. Rigidity, though occasionally
170 inefficient, makes cooperation more robust, extending
171 it to environments where trust is fragile and flexibility too
172 risky.

173 In this sense, rigid moral rules act as a social technology:
174 they trade efficiency for robustness, extending the reach
175 of reputation-based cooperation. Based on this model,
176 we predict that flexibility will prevail in high trust
177 environments—such as families, friendships, and tightly
178 knit teams—or when dealing with easily verified hardships
179 like physical injury. Rigid rules, by contrast, should
180 emerge among strangers, in large groups, or in low-trust
181 societies.

182 More broadly, this logic helps explain a central puzzle
183 in moral psychology: why humans often treat certain
184 rules and values as sacred, insisting on duty regardless
185 of consequence. What appears as deontological rigidity

187 may instead reflect an adaptive response to the problem
188 of sustaining cooperation under uncertainty.

189 Model

190 We develop a model of reputation-based cooperation where
191 individuals sometimes face genuine hardship—situations
192 in which cooperating would be too costly—but others
193 cannot tell for sure whether this is the case.

196 **Players and basic structure.** We consider a repeated game with
197 two roles: $n \gg 1$ signalers and an infinite pool of choosers.
198 Signalers make cooperative decisions across rounds and
199 hold reputations that reflect their actions. Choosers, in
200 contrast, participate in only one round and decide whether
201 to trust a signaler based on their reputation.

202 This setup separates two adaptive problems. For
203 choosers, the challenge is trust—predicting whether a
204 partner will reciprocate trust if given the opportunity.
205 For signalers, the challenge is reputation management—
206 acting in ways that secure the trust of future partners.

207 To capture variation in reputation management strate-
208 gies, we assume signalers differ in their time preferences.
209 Each signaler is characterized by a discount factor δ (with
210 values between 0 and 1), which measures how much they
211 value future payoffs relative to immediate ones. More
212 patient signalers (higher δ) are more willing to incur
213 the immediate costs of cooperation to maintain their
214 reputation with choosers, whereas impatient individuals
215 (lower δ) prioritize short-term gains and are more likely
216 to defect.

217 We assume that discount factors are continuously dis-
218 tributed across the population. This produces a spectrum
219 of patience levels, capturing the full range of cooperative
220 tendencies. In particular, it allows for strategic actors
221 who exploit flexible norms but are restrained by rigid
222 ones—the kind of behavior that rigid rules are designed
223 to prevent.

224 As illustrated in Figure 1, the game proceeds in two
225 phases: signalers first participate in a public goods game,
226 before interacting with choosers in trust games.

228 **Initial collective action under potential hardship.** In the initial
229 interaction, signalers decide whether to contribute to a
230 collective action or free-ride on others' contributions. They
231 play a linear public goods game, where each contribution
232 imposes a private cost on the signaler but generates a
233 public benefit for others.

234 The cost of contribution depends on circumstances.
235 Under normal conditions, contribution imposes a relatively
236 low private cost γ . However, with probability p , a signaler
237 experiences hardship—illness, competing obligations, or a
238 personal crisis—that increases their cost of contribution
239 to γ_H .

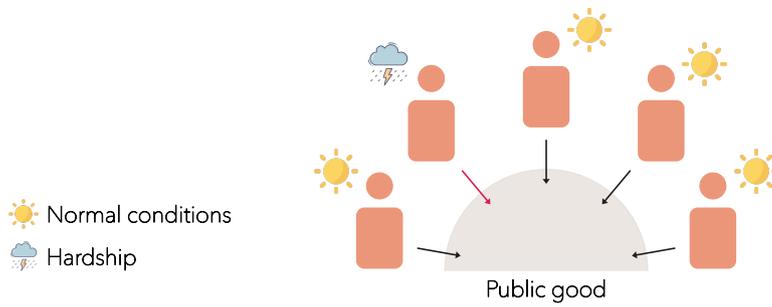
240 Each contribution yields a public benefit β , which is
241 shared equally among the other signalers (excluding the
242 contributor). The net social value of a contribution is
243 therefore $\beta - \gamma$ under normal conditions and $\beta - \gamma_H$ under
244 hardship. We assume:

$$245 0 < \gamma < \beta < \gamma_H,$$

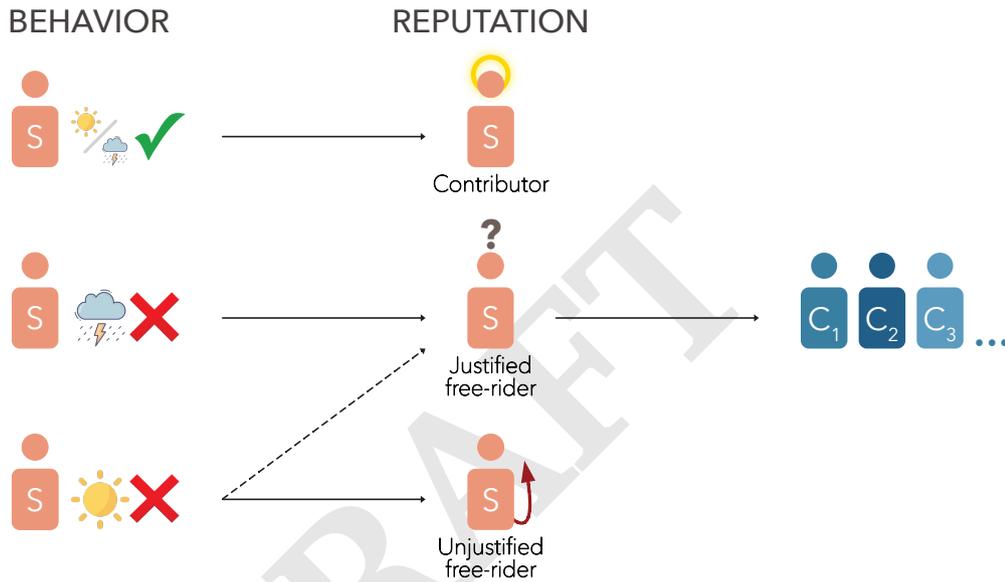
246 making contributions socially efficient under normal
247 conditions ($\beta - \gamma > 0$), but inefficient under hardship
248 ($\beta - \gamma_H < 0$).

251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314

a Initial collective action



b Reputation formation



c Subsequent trust games

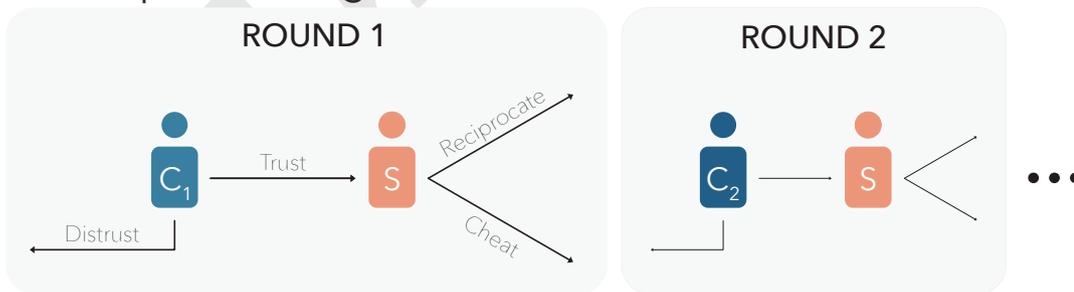


Fig. 1 — Model overview. **a**, Initially, signalers decide whether to contribute to a public good. Contribution is efficient under normal conditions but inefficient under hardship. **b**, Signalers acquire reputations based on their behavior. Those who free-ride under normal conditions are labeled justified free-riders with probability ε , making them indistinguishable from those who free-ride under genuine hardship. The parameter ε captures plausible deniability—how easily opportunists can be mistaken for genuinely hard-pressed individuals. **c**, Subsequently, signalers interact with choosers, who decide whether to trust based on reputation; if trusted, signalers decide whether to reciprocate.

Signalers decide whether to contribute based on their circumstances. Choosers perfectly observe the outcome (contribution or free-riding), but not the underlying circumstances that led to that decision.

As shown in Figure 1b, this partial observability generates reputational ambiguity. Signalers who contribute become contributors, regardless of circumstances. Those who free-ride under hardship become justified free-riders. However, those who free-ride under normal conditions

mistakenly become justified free-riders with probability ε , and are otherwise labeled unjustified free-riders.

The parameter ε captures the degree of plausible deniability: how easily opportunistic free-riders can appear genuinely hard-pressed. We model this as observational noise for simplicity, but an equivalent formulation would let individuals who free-ride under normal conditions falsely claim hardship at no cost, succeeding with probability ε —an opportunity they would all take in equilibrium.

315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378

This assumption preserves the strategic logic of plausible deniability without adding unnecessary complexity.

Subsequent trust games. After the initial stage, each signaler interacts with a sequence of choosers. In every round, a new chooser decides whether to trust the signaler, incurring cost c to provide benefit b (with $0 < c < b$). If trusted, the signaler can reciprocate—incurring cost c to return benefit b to the chooser—or cheat, keeping the benefit without giving back. These trust games are one-shot for choosers but repeated for signalers, creating an incentive to maintain a cooperative reputation over time.

After each round, a signaler's reputation updates to reflect their most recent behavior. The full reputation combines their initial collective action label (contributor, justified free-rider, or unjustified free-rider) with their most recent behavior in trust games (reciprocator or cheater), if any. Choosers rely on this composite reputation when deciding whether to place their trust.

Results

The analysis reveals two cooperative equilibria, Perfect Bayesian Equilibria that sustain reputation-based cooperation (see Methods for a formal definition): a flexible norm that accommodates apparent hardship, and a rigid norm that categorically demands cooperation regardless of context.

For clarity, we assume:

$$\frac{\gamma}{1-\varepsilon} < c < \gamma_H.$$

This assumption ensures that contributing under normal conditions is less demanding than reciprocating trust (left inequality), and that contributing under hardship is more demanding (right inequality). It simplifies the behavioral thresholds below without changing any qualitative result (see Supplementary Information for the general case).

Two equilibria: flexible and rigid norms.

Flexible norm. In this equilibrium, choosers interpret free-riding leniently: if a signaler appears to have faced hardship, they are still deemed trustworthy. In the initial trust game, before a signaler's reciprocation behavior is known, trust is extended to contributors and justified free-riders, but withheld from unjustified free-riders. Later, once reciprocation is observed, choosers trust reciprocators and distrust cheaters.

Signalers never contribute under hardship, since free-riding then carries no reputational penalty. Under normal circumstances, however, they face a trade-off: free-riding avoids the immediate cost of contribution, γ , but risks losing the next chooser's trust—worth b —with probability $1 - \varepsilon$, if the signaler's true circumstances are observed. Signalers then contribute if and only if their discount factor satisfies $\delta \geq \delta_{\text{normal}}^{\text{flex}}$, where:

$$\delta_{\text{normal}}^{\text{flex}} = \frac{\gamma}{(1-\varepsilon)b}. \quad [\text{F1}]$$

In subsequent trust games, reciprocation incurs cost c but secures next chooser's trust, worth b . Signalers

reciprocate if and only if their discount factor satisfies $\delta \geq \delta_{\text{recip.}}$, where:

$$\delta_{\text{recip.}} = \frac{c}{b}.$$

Because $\frac{\gamma}{1-\varepsilon} < c$, this threshold is higher than $\delta_{\text{normal}}^{\text{flex}}$; contributing under normal circumstances requires less patience than reciprocating trust.

From these thresholds, three signaler types emerge. Impatient signalers ($\delta < \delta_{\text{normal}}^{\text{flex}}$) always free-ride and always cheat, regardless of circumstances or reputation. Intermediate signalers ($\delta_{\text{normal}}^{\text{flex}} \leq \delta < \delta_{\text{recip.}}$) contribute under normal circumstances but cheat when trusted. Patient signalers ($\delta \geq \delta_{\text{recip.}}$) contribute under normal circumstances and always reciprocate trust.

The flexible norm is sustainable whenever initial trust decisions are well-founded. The binding constraint involves justified free-riders, who must reciprocate frequently enough to make trusting them worthwhile. The flexible norm exists if and only if:

$$\frac{p\mathbb{P}(\delta \geq \delta_{\text{recip.}})}{p + (1-p)\varepsilon\mathbb{P}(\delta < \delta_{\text{normal}}^{\text{flex}})} \geq \frac{c}{b}. \quad [\text{F2}]$$

This condition ensures that justified free-riders are, on average, sufficiently likely to reciprocate relative to the cost-to-benefit ratio of trust, $\frac{c}{b}$. The numerator gives the probability that a reciprocator acquires the justified free-rider reputation—necessarily by facing hardship—while the denominator gives the overall probability of that reputation: all signalers become justified free-riders when they face hardship, with probability p , and free-riders also acquire this reputation when they face normal conditions and successfully misrepresent them as hardship, with probability $(1-p)\varepsilon$.

Rigid norm. In this equilibrium, choosers adopt a zero-tolerance policy: only contributors are trusted, while any instance of free-riding—regardless of apparent circumstances—results in loss of trust. In the initial trust game, only contributors are trusted, while both justified and unjustified free-riders are distrusted. Later, choosers trust reciprocators and distrust cheaters.

Signalers now face stronger incentives to contribute under normal conditions. They do so if and only if their discount factor satisfies $\delta \geq \delta_{\text{normal}}^{\text{rigid}}$, where:

$$\delta_{\text{normal}}^{\text{rigid}} = \frac{\gamma}{b}. \quad [\text{R1}]$$

This threshold is lower than under the flexible norm ($\frac{\gamma}{b} < \frac{\gamma}{b(1-\varepsilon)}$). Signalers whose discount factor falls between these two thresholds can be understood as opportunists: they contribute when rigidity prevents misrepresentation but free-ride when flexibility allows reputational ambiguity to be exploited—opportunists, in short, exploit moral wiggle room when it is available.

The rigid norm also creates incentives to contribute under hardship. A signaler facing hardship contributes if and only if their discount factor satisfies $\delta \geq \delta_{\text{hardship}}^{\text{rigid}}$, where:

$$\delta_{\text{hardship}}^{\text{rigid}} = \frac{\gamma_H}{\gamma_H + b - c}. \quad [\text{R2}]$$

Signalers reciprocate trust if and only if their discount factor satisfies $\delta \geq \delta_{\text{recip.}} = \frac{c}{b}$, as with the flexible

Domains of sustainable cooperation

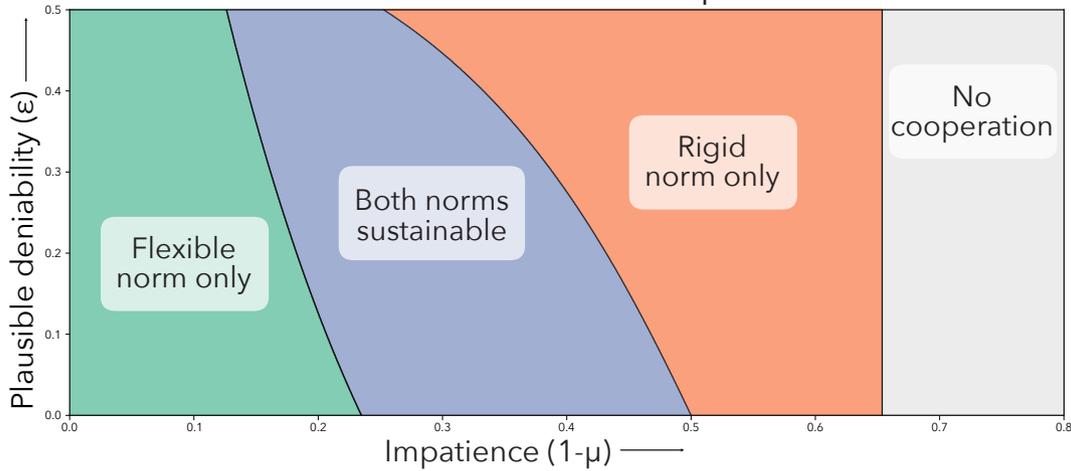


Fig. 2 — Domains of sustainable cooperation. Each point represents a social environment defined by impatience $1 - \mu$ (x-axis) and plausible deniability ε (y-axis). Colors indicate where the flexible (F2) and rigid (R3-R4) norms are sustainable. Green: only the flexible norm is sustainable; orange: only the rigid norm; blue: both; white: cooperation fails. Discount factors follow a normal distribution with mean μ and standard deviation $\sigma = 0.3$, truncated to the interval $(0, 1)$. Parameters: $c = 1, b = 2, \gamma = 0.5, \beta = 0.75, \gamma_H = 2, p = 0.2$. We vary impatience $1 - \mu$ between 0 and 0.75 (i.e., μ between 1 and 0.25) and plausible deniability ε between 0 and $1 - \frac{\gamma}{c} = 0.5$, ensuring that $\gamma/(1 - \varepsilon) < c$ holds throughout.

norm. From these thresholds, four types of signalers emerge. Very impatient signalers ($\delta < \delta_{\text{normal}}^{\text{rigid}}$) always free-ride and cheat; impatient signalers ($\delta_{\text{normal}}^{\text{rigid}} \leq \delta < \delta_{\text{recip.}}$) contribute under normal circumstances but cheat when trusted; patient signalers ($\delta_{\text{recip.}} \leq \delta < \delta_{\text{hardship}}^{\text{rigid}}$) contribute under normal circumstances and reciprocate; and very patient signalers ($\delta \geq \delta_{\text{hardship}}^{\text{rigid}}$) contribute even under hardship and reciprocate.

The rigid norm is sustainable whenever initial trust decisions are well-founded. There are two binding conditions, involving contributors and justified free-riders respectively. The rigid norm exists if and only if:

$$\frac{p \mathbb{P}(\delta_{\text{recip.}} \leq \delta < \delta_{\text{hardship}}^{\text{rigid}})}{p \mathbb{P}(\delta < \delta_{\text{hardship}}^{\text{rigid}}) + (1 - p) \varepsilon \mathbb{P}(\delta < \delta_{\text{normal}}^{\text{rigid}})} < \frac{c}{b}, \quad [\text{R3}]$$

$$\frac{(1 - p) \mathbb{P}(\delta \geq \delta_{\text{recip.}}) + p \mathbb{P}(\delta \geq \delta_{\text{hardship}}^{\text{rigid}})}{(1 - p) \mathbb{P}(\delta \geq \delta_{\text{normal}}^{\text{rigid}}) + p \mathbb{P}(\delta \geq \delta_{\text{hardship}}^{\text{rigid}})} \geq \frac{c}{b}. \quad [\text{R4}]$$

Condition (R3) ensures that justified free-riders are sufficiently unlikely to reciprocate, making it sensible to distrust them, while condition (R4) ensures that contributors are sufficiently likely to reciprocate, making trust worthwhile. In both cases, the numerators give the probability that a reciprocator holds the relevant reputation, and the denominators give its overall probability in the population.

Comparing flexible and rigid norms. These two equilibria represent distinct approaches to moral judgment under uncertainty. The flexible norm extends trust to those who appear to have faced hardship; the rigid norm withholds it. As captured by conditions (F2) and (R3), whether such trust is well-founded depends on how likely these justified free-riders are to reciprocate trust.

That likelihood, in turn, depends on two features of the environment: the distribution of discount factors, which shapes the population's cooperativeness, and plausible

deniability, which measures how easily selfish behavior can be disguised as hardship.

To illustrate, we assume signalers' discount factors follow a normal distribution with mean μ and standard deviation σ , truncated to the interval $(0, 1)$. Figure 2 shows the resulting domains of sustainable cooperation as functions of impatience ($1 - \mu$) and plausible deniability (ε). When the population is patient and deniability is low, only the flexible norm is sustainable: trust can safely extend to justified free-riders, who are likely to reciprocate and to have faced genuine hardship. As impatience and deniability increase, the flexible norm collapses, and cooperation can persist only under the rigid norm. When impatience becomes extreme, even contributors can no longer be trusted, and cooperation breaks down altogether.

Rigid norms do more than offer an alternative to flexible ones—they enable cooperation in contexts where flexibility would fail. Figure 3 shows how rigid norms achieve this by comparing performance where both norms are possible. The rigid norm is more robust to strategic defection but less efficient: it prevents opportunistic exploitation of ambiguity (Figure 3a), but incentivizes cooperation under hardship, leading to an overall payoff loss (Figure 3b).

Thus, rigid norms function as social technologies that trade efficiency for robustness against strategic defection, extending the domain of cooperation by eliminating exploitable ambiguity.

Discussion

Rigid rules as social technologies. Why do societies adopt rigid moral rules when humans are perfectly capable of flexible, context-sensitive moral reasoning? Our model addresses this question by examining how cooperation evolves when people's circumstances are private and reputations ambiguous.

Individuals in our model sometimes face genuine hardship—situations in which cooperating would be socially inefficient. Yet observers cannot perfectly assess

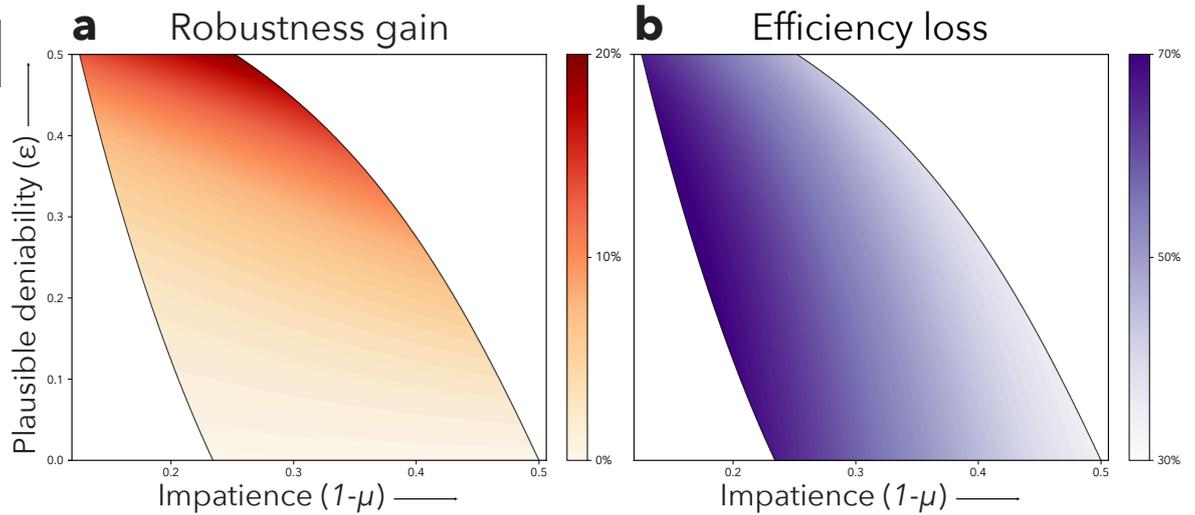
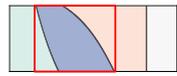


Fig. 3 — The efficiency–robustness trade-off. Each point represents a social environment defined by impatience $1 - \mu$ (x -axis) and plausible deniability ε (y -axis). We restrict attention to the domain where both norms are sustainable (blue region in Figure 2) and compare their performance. **a**, Robustness gain from adopting the rigid norm. Warm shades (yellow to red) show the fraction of opportunists who, under normal conditions, would free-ride under the flexible norm but contribute when cooperation is categorically required. **b**, Efficiency loss from adopting the rigid norm. Purple shades show the normalized payoff difference from the initial collective action (flexible – rigid). Parameters and distribution of discount factors as in Figure 2.

these circumstances: with some probability, an individual who defects under normal conditions is mistaken for someone acting under hardship. This probability captures the degree of plausible deniability—how easily selfish behavior can appear as genuine constraint. Individuals also differ in their time preferences: more future-oriented individuals are more willing to incur the immediate costs of cooperation to maintain their reputation, whereas present-oriented individuals prioritize short-term gains.

The analysis reveals two cooperative equilibria, representing two approaches to moral judgment under uncertainty: a flexible norm that accommodates apparent hardship, and a rigid norm that categorically demands cooperation regardless of context. These equilibria succeed under different conditions. The flexible norm is sustainable when trust is secure; when plausible deniability is low and individuals are future-oriented, those who appear hard-pressed are genuinely trustworthy, and contextual discretion does not invite exploitation. In contrast, when individuals are present-oriented and hardship is easy to fake, flexible trust becomes too risky, but cooperation can be sustained under the rigid norm.

The rigid norm thus extends cooperation to settings where flexible trust would fail. Comparing it to the flexible norm reveals an efficiency–robustness trade-off: by treating all defections alike, rigidity makes trust decisions robust to strategic defection, but at the cost of sometimes enforcing inefficient cooperation. Rigid norms appear as social technologies that trade efficiency for robustness.

Empirical predictions. Our model predicts that rigid rules—those that demand cooperation even in apparent hardship—should arise when flexibility is most risky; that is, when the potential for strategic defection is highest.

First, rigidity should emerge around circumstances that are easy to fake. Consider pregnancy: women in late stages, where their condition is more visible, often receive accommodation, such as seats on public transport or flexibility at work. Such accommodations are much rarer in early pregnancy, when needs are comparable but

less apparent. More generally, hard-to-verify claims like stress or fatigue should receive less accommodation than observable injuries.

Second, moral rules should be applied more rigidly to people perceived as impulsive or untrustworthy. This may explain why young children are taught categorical moral rules with little room for contextual nuance—they are not yet seen as reliable enough to exercise moral discretion (see also 13). A similar logic shapes attitudes toward welfare: when recipients are perceived as lazy rather than unlucky, people tend to favor strict eligibility rules (30).

Third, rigid rules should be more frequent in distant relationships. We predict that individuals should cooperate more flexibly with family and friends, but adopt more rigid rules with colleagues or strangers. This sheds light on relational psychology (31). Recent work suggests that formality is a fundamental dimension of relational psychology, distinguishing public relationships that adhere to rules and regulations from private relationships with looser, more casual styles (32). Consistent with our prediction, relationships become more formal in modernized societies, where interactions increasingly involve strangers rather than intimate social circles.

Fourth, distrust should foster rigidity. This helps explain why some societies are tight—with many strong norms and a low tolerance of deviant behavior—and others are loose (18). Tightness is associated with ecological and social threats, such as disease, natural disasters, resource scarcity, population density, and intergroup conflict (33–35). Existing accounts attribute this pattern to coordination demands: groups under threat require stronger norms to sustain collective action (formalized in 36). Our model complements this view by identifying a cognitive mechanism that makes rigidity appealing. In threatening environments, immediate needs loom larger (37), making investment in one’s cooperative reputation less worthwhile. Recognizing this, people expect others to be less reliably cooperative and see flexibility as too risky, turning instead to more rigid rules—the loss in efficiency is outweighed by the gain in robustness. As

763 shown by Nettle and Saxe (38), a similar mechanism
764 underlies support for authoritarian leaders and punitive
765 regulation, explaining why these political preferences also
766 rise under threat (39–42).

768 **The logic of moral rigidity.** Moral judgments are often strik-
769 ingly deontological: people treat certain rules, values,
770 and rights as sacred or inalienable, refusing to weigh
771 them against costs, benefits, or pragmatic compromises
772 (1, 43–46). One way to understand this rigidity is through
773 the efficiency–robustness trade-off. As we have shown,
774 categorical rules close loopholes, making cooperation
775 easier to sustain. The same logic extends more broadly:
776 rigid principles keep expectations clear, make departures
777 easier to detect or sanction, and prevent minor allowances
778 from hardening into precedents that legitimize further
779 slippage, gradually weakening the cooperative order (47).
780 People may therefore endorse some rules categorically not
781 because they view them as intrinsically right, but because
782 they sense that strict adherence is the most reliable way
783 to sustain cooperation in the long run.

784 Our account complements existing explanations for
785 moral rigidity. Rules can act as cognitive shortcuts
786 (48), helping individuals approximate—and coordinate
787 on—mutually beneficial bargains (49). Bright-line norms
788 can also facilitate coordination among third-party en-
789 forciers (50) and bystanders (51). And principled, trade-off-
790 insensitive behavior can serve as a signal of commitment,
791 helping individuals build trust (52, 53).

794 **Implications for evolutionary game theory.** Our model differs
795 from the standard framework for reputation-based coop-
796 eration in three ways.

797 First, we introduce heterogeneity in cooperative types.
798 In standard indirect reciprocity models, individuals are
799 identical and interactions are repeated (54–56). Norms—
800 understood as rules for assigning good or bad reputation—
801 are then evaluated by their capacity to generate and
802 stabilize cooperation (57–61). We instead assume that
803 signalers vary in patience, integrating the logic of signaling
804 with that of repeated cooperation (building on 62; see also
805 63, 64). Patient signalers reciprocate trust to attract the
806 trust of future partners, whereas impatient ones cheat. Dif-
807 ferent norms then emerge depending on which reputations
808 predict reciprocation and which predict cheating. This
809 avoids the scoring dilemma: cooperation in the trust game
810 is sustainable because first-order reputations (reciprocator
811 or cheater) perfectly predict future behavior.

812 Second, we add a public goods game, following the
813 structure of Panchanathan and Boyd (65). The signaling
814 framework clarifies when contributions can be stabilized:
815 signalers must be sufficiently patient for choosers to infer
816 trustworthiness from contribution.

817 Third, we introduce opaque exculpatory circumstances,
818 a feature rarely modeled (although see 66). This creates
819 opportunities for strategic misrepresentation, leading to
820 reputational ambiguity. Flexible assessments emerge when
821 choosers can trust despite this ambiguity; rigidity becomes
822 necessary when trust is too risky. Adopting a signaling per-
823 spective thus reframes norm evolution: stricter assessment
824 rules—akin to Stern Judging in indirect reciprocity models
825 (61, 67–69)—arise as adaptive responses to observers’
826 inferential challenges.

827 **Limits and future directions.** Our model isolates one adap-
828 tive problem—judging trustworthiness under opaque
829 circumstances—and shows how it can produce both flexi-
830 ble and rigid assessments. But moral norms are not just
831 solutions to individual inference problems. They involve
832 widely shared expectations about how people ought to
833 behave, the felt obligation to conform, and culturally
834 inherited conventions that come to seem external or
835 binding (70). Real-world cooperation also requires forms
836 of coordination that our analysis sets aside—coordination
837 on the standards we use to assess others, on when and
838 how to enforce them, and on the shared courses of action
839 that make collective endeavors possible. Integrating our
840 approach with models that have examined such dynamics
841 (71–74) is a natural direction for future work.

842 Conclusion

843 Flexible minds, rigid rules: our model shows this contrast
844 arises from a simple informational constraint—people
845 often cannot tell whether noncooperation reflects genuine
846 hardship or convenient excuses. This ambiguity gives
847 opportunists room to exploit flexible rules. Rigid rules
848 close that loophole, making cooperation more robust
849 to strategic defection, even at the cost of sometimes
850 enforcing inefficient behavior. This efficiency–robustness
851 trade-off explains when societies tighten norms and why
852 consequentialist moral reasoning can nonetheless give rise
853 to judgments that seem deontological.

854 Materials and Methods

855 Model description.

856 **Players.** We consider a repeated game with $n \gg 1$ long-run signalers
857 and an infinite pool of short-run choosers. Before play begins, signalers are
858 assigned a discount factor δ according to a continuous distribution with full
859 support on $\Delta = (0, 1)$. They privately observe their discount factor; the
860 distribution is common knowledge.

861 **Interactions.** The game proceeds in two stages. Initially (round 0),
862 signalers decide whether to contribute to a public good. Each contribution
863 generates collective benefit β , divided equally among the $n - 1$ other
864 signalers. The cost of contribution depends on circumstances: γ_H under
865 hardship (probability $p \in (0, 1)$) and γ under normal conditions (probability
866 $1 - p$), where $0 < \gamma < \beta < \gamma_H$. Signalers privately observe their cost
867 before deciding.

868 Subsequently (rounds $t \geq 1$), each signaler plays a repeated trust
869 game with a sequence of choosers. In every round, the signaler is paired
870 with a new chooser, who decides whether to trust them, incurring cost c
871 to provide benefit b (where $0 < c < b$). If trusted, the signaler chooses
872 whether to reciprocate, incurring cost c to return benefit b .

873 **Reputations.** Signalers acquire a public goods game reputation ω_{pgg} and
874 a trust game reputation ω_{tg} . Their full reputation is $\omega = (\omega_{\text{pgg}}, \omega_{\text{tg}}) \in \Omega$.

875 There are three possible values for ω_{pgg} . Signalers who contribute
876 become contributors, those who free-ride under hardship become justified
877 free-riders, and those who free-ride under normal conditions become
878 justified free-riders with probability $\varepsilon \in (0, 1)$ and unjustified free-riders
879 with probability $1 - \varepsilon$.

880 There are three possible values for ω_{tg} . Signalers who have not yet been
881 trusted have unobserved trust game behavior; those who last reciprocated
882 are reciprocators; and those who last cheated are cheaters.

883 For simplicity of exposition, when a signaler’s trust game behavior
884 is unobserved ($\omega_{\text{tg}} = \text{unobserved}$), we refer to them by their public
885 goods game reputation alone—for example, contributors are signalers with
886 reputation (contributor, unobserved). Once a signaler has been trusted,
887 we refer to them by their trust game reputation—for example, reciprocators
888 are signalers with any reputation (\cdot , reciprocator).
889
890

891 **Strategies and beliefs.** A signaler public goods game strategy,
892 represented by a function $\sigma_{\text{pgg}} : \Delta \times \{\text{normal, hardship}\} \rightarrow$
893 $\{\text{contribute, free-ride}\}$, specifies whether to contribute depending on
894 discount factor and circumstances. A signaler trust game strategy, represented
895 by a function $\sigma_{\text{tg}} : \Delta \times \Omega \rightarrow \{\text{reciprocate, cheat}\}$, specifies whether to
896 reciprocate trust depending on discount factor and reputation.

896 Choosers hold shared beliefs $\phi(\cdot | \omega)$ over Δ for each reputation $\omega \in \Omega$,
897 encoding how they assign likelihood to different signaler discount factors,
898 conditional on observed reputation. A chooser strategy, represented by a
899 function $\sigma_{\text{ch}} : \Omega \rightarrow \{\text{trust, distrust}\}$, specifies whether to trust a signaler
900 based on their reputation.

900 A strategy profile is $\sigma = (\sigma_{\text{pgg}}, \sigma_{\text{tg}}, \sigma_{\text{ch}})$.

901 Equilibrium analysis.

903 **Equilibrium concept.** We study the model's Perfect Bayesian Equilibria
904 (PBEs) in pure strategies. A PBE is a pair (σ, ϕ) where players have no
905 profitable deviations at any history (on or off the equilibrium path), and
906 beliefs are updated according to Bayes' rule whenever possible (75).

907 **Cooperative equilibria.** We focus on cooperative PBEs, in which
908 choosers' trust decisions incentivize cooperation at every history. Choosers
909 must trust contributors and distrust unjustified free-riders (incentivizing
910 contributions to the public good), and trust reciprocators and distrust
911 cheaters (incentivizing reciprocation in every subgame). We obtain two
912 cooperative PBEs: the flexible and the rigid norm.

912 In the Supplementary Information, we characterize the model's other
913 PBEs. These fall into two categories: PBEs that are outcome-equivalent to
914 the flexible or rigid norm, differing only in off-path play; and PBEs in which
915 signalers never contribute to the public good.

916
917
918 1. F Cushman, L Young, M Hauser, The role of conscious reasoning and intuition in moral
919 judgment: testing three principles of harm. *Psychol. Sci.* **17**, 1082–1089 (2006).
920 2. JD Greene, RB Somerville, LE Nystrom, JM Darley, JD Cohen, An fMRI investigation of
921 emotional engagement in moral judgment. *Science* **293**, 2105–2108 (2001).
922 3. M Kleiman-Weiner, T Gerstenberg, S Levine, JB Tenenbaum, Inference of intention and
923 permissibility in moral decision making in *Proceedings of the 37th Annual Conference of the*
924 *Cognitive Science Society*. (Cognitive Science Society, Austin, TX), pp. 1123–1128 (2015).
925 4. M Kleiman-Weiner, R Saxe, JB Tenenbaum, Learning a commonsense moral theory.
926 *Cognition* **167**, 107–123 (2017).
927 5. A Le Pargneux, F Cushman, Moral judgment is sensitive to bargaining power. *J. Exp. Psychol.*
928 *Gen.* **154**, 263–278 (2025).
929 6. A Le Pargneux, N Chater, H Zeitoun, Contractualist tendencies and reasoning in moral
930 judgment and decision making. *Cognition* **249**, 105838 (2024).
931 7. S Levine, M Kleiman-Weiner, L Schulz, J Tenenbaum, F Cushman, The logic of
932 universalization guides moral judgment. *Proc. Natl. Acad. Sci. United States Am.* **117**,
933 26158–26169 (2020).
934 8. L Young, R Saxe, When ignorance is no excuse: Different roles for intent across moral
935 domains. *Cognition* **120**, 202–214 (2011).
936 9. S Levine, M Kleiman-Weiner, N Chater, F Cushman, JB Tenenbaum, When rules are
937 over-ruled: Virtual bargaining as a contractualist method of moral judgment. *Cognition* **250**,
938 105790 (2024).
939 10. J Kwon, T Zhi-Xuan, J Tenenbaum, S Levine, When it is not out of line to get out of line: The
940 role of universalization and outcome-based reasoning in rule-breaking judgments (2023).
941 11. A Le Pargneux, X Roberts-Gaal, L Woodley, JD Greene, FA Cushman, Hard bargains and
942 even splits: Fairness judgments track bargaining power across diverse cultures. *OSF Prepr.*
943 (2025).
944 12. E Awad, et al., When is it acceptable to break the rules? Knowledge representation of moral
945 judgements based on empirical data. *Auton. Agents Multi-Agent Syst.* **38**, 35 (2024).
946 13. M Jambon, JG Smetana, Moral complexity in middle childhood: children's evaluations of
947 necessary harm. *Dev. Psychol.* **50**, 22–33 (2014).
948 14. J Marshall, K Wynn, P Bloom, Do children and adults take social relationship into account
949 when evaluating people's actions? *Child Dev.* **91**, e1082–e1100 (2020) .eprint:
950 <https://onlinelibrary.wiley.com/doi/pdf/10.1111/cdev.13390>.
951 15. J Marshall, et al., When not helping is nice: Children's changing evaluations of helping during
952 COVID-19. *Dev. Psychol.* **59**, 953–962 (2023).
953 16. C Boehm, *Moral Origins: The Evolution of Virtue, Altruism, and Shame*. (Soft Skull Press, New
954 York), (2012).
955 17. L Glowacki, S Lew-Levy, How small-scale societies achieve large-scale cooperation. *Curr.*
956 *Opin. Psychol.* **44**, 44–48 (2022).
957 18. MJ Gelfand, JR Harrington, JC Jackson, The strength of social norms across human groups.
958 *Perspectives on Psychol. Sci.* **12**, 800–809 (2017).
959 19. L Kohlberg, *Essays on Moral Development, Volume I: The Philosophy of Moral Development*.
960 (Harper & Row, San Francisco), (1981).
961 20. JR Rest, Moral judgment research and the cognitive-developmental approach to moral
962 education. *The Pers. Guid. J.* **58**, 602–605 (1980).
963 21. JR Snarey, J Reimer, L Kohlberg, Development of social-moral reasoning among kibbutz
964 adolescents: A longitudinal cross-cultural study. *Dev. Psychol.* **21**, 3–17 (1985).
965 22. Southwark London Borough Council v Williams (1971).
966 23. J Altman, *Hunter-Gatherers Today: An Aboriginal Economy in North Australia*, Social
967 Anthropology Series. (Australian Institute of Aboriginal Studies, Canberra) No. 7, (1987).

955 Numerical solution.

956 **Evaluation of domains and performance.** We numerically evaluate
957 the domain conditions for the flexible (F2) and rigid (R3-R4) norms on
958 a parameter grid. These conditions are stated in terms of probabilities
959 over discount factors; to evaluate them, we fix the distribution of discount
960 factors to a normal with mean μ and standard deviation $\sigma = 0.3$, truncated
961 to $(0, 1)$ (so μ refers to the pre-truncation mean). We vary impatience
962 $1 - \mu$ and plausible deniability ε , holding other parameters fixed as in the
963 figure captions. Figure 2 displays the domains of sustainable cooperation
964 by identifying parameter combinations for which these conditions hold;
965 Figure 3 compares performance where both norms are sustainable.

966 **Data, Materials, and Software Availability.** No new empirical data were
967 generated for this study. All analytical derivations are provided in the
968 Supplementary Information. The code used to generate Figures 2 and 3 is
969 available in the public OSF repository for this paper (<https://osf.io/2jxry>).

970 **ACKNOWLEDGMENTS.** We thank researchers in the Plotkin Research
971 Group and the Akçay Lab at the University of Pennsylvania for valuable
972 discussions, and Taylor Kessinger in particular for a helpful comment that
973 improved the model. This study was supported by the EUR FrontCog grant
974 ANR-17-EURE-0017 and ANR-10-IDEX-0001-02 to PSL. Funding from
975 the French Agence Nationale de la Recherche (under the Investissement
976 d'Avenir programme, ANR-17-EURE-0010) is gratefully acknowledged.
977 Open Access funding provided by the Max Planck Society.

978
979
980 24. LJ Pospisil, *Kapauku Papuans and Their Law*, Yale University Publications in Anthropology.
981 (Department of Anthropology, Yale University, New Haven) No. 54, (1958).
982 25. MS Alvard, DA Nolin, Rousseau's whale hunt? coordination among big-game hunters. *Curr.*
983 *Anthropol.* **43**, 533–559 (2002).
984 26. D Damas, Central Eskimo systems of food sharing. *Ethnology* **11**, 220–240 (1972).
985 27. G Wenzel, Inuit subsistence and hunter support in nunavut in *Nunavut: Inuit Regain Control of*
986 *Their Lands and Lives*, eds. J Dahl, J Hicks, P Jull. (International Work Group for Indigenous
987 Affairs, Copenhagen), pp. 180–194 (2000).
988 28. J Dana, RA Weber, JX Kuang, Exploiting moral wiggle room: Experiments demonstrating an
989 illusory preference for fairness. *Econ. Theory* **33**, 67–80 (2007).
990 29. Z Grossman, JJ Van Der Weele, Self-image and willful ignorance in social decisions. *J. Eur.*
991 *Econ. Assoc.* **15**, 173–217 (2017).
992 30. L Aaroe, MB Petersen, Crowding out culture: Scandinavians and americans agree on social
993 welfare in the face of deservingness cues. *J. Polit.* **76**, 684–697 (2014).
994 31. AP Fiske, The four elementary forms of sociality: Framework for a unified theory of social
995 relations. *Psychol. Rev.* **99**, 689–723 (1992).
996 32. X Cheng, et al., The conceptual structure of human relationships across modern and historical
997 cultures. *Nat. Hum. Behav.* pp. 1–14 (2025).
998 33. MJ Gelfand, et al., Differences between tight and loose cultures: A 33-nation study. *Science*
999 **332**, 1100–1104 (2011).
1000 34. JR Harrington, MJ Gelfand, Tightness–looseness across the 50 united states. *Proc. Natl.*
1001 *Acad. Sci.* **111**, 7990–7995 (2014).
1002 35. JC Jackson, MJ Gelfand, CR Ember, A global analysis of cultural tightness and its relationship
1003 with ecological threat, social complexity, and social structure. *Proc. Royal Soc. B: Biol. Sci.*
1004 **287**, 20191227 (2020).
1005 36. P Roos, M Gelfand, D Nau, J Lun, Societal threat and cultural variation in the strength of social
1006 norms: An evolutionary basis. *Organ. Behav. Hum. Decis. Process.* **129**, 14–23 (2015).
1007 37. H Mell, N Baumard, JB André, Time is money: Waiting costs explain why selection favors
1008 steeper time discounting in deprived environments. *Evol. Hum. Behav.* **42**, 379–387 (2021).
1009 38. D Nettle, R Saxe, 'if men were angels, no government would be necessary': The intuitive
1010 theory of social motivation and preference for authoritarian leaders. *Collabra: Psychol.* **7**,
1011 28105 (2021).
1012 39. J Duckitt, K Fisher, The impact of social threat on worldview and ideological attitudes. *Polit.*
1013 *Psychol.* **24**, 199–222 (2003).
1014 40. L Laustsen, MB Petersen, Does a competent leader make a good friend? Conflict, ideology
1015 and the psychologies of friendship and followership. *Evol. Hum. Behav.* **36**, 286–293 (2015).
1016 41. L Safra, et al., Childhood harshness predicts long-lasting leader preferences. *Evol. Hum.*
1017 *Behav.* **38**, 645–651 (2017).
1018 42. KEG Williams, AM Votruba, SL Neuberg, MJ Saks, Capital and punishment: Resource scarcity
1019 increases endorsement of the death penalty. *Evol. Hum. Behav.* **40**, 65–73 (2019).
1020 43. PE Tetlock, Thinking the unthinkable: Sacred values and taboo cognitions. *Trends Cogn. Sci.*
1021 **7**, 320–324 (2003).
1022 44. JD Greene, JM Paxton, Patterns of neural activity associated with honest and dishonest moral
1023 decisions. *Proc. Natl. Acad. Sci.* **106**, 12506–12511 (2009).
1024 45. J Haidt, SH Koller, MG Dias, Affect, culture, and morality, or is it wrong to eat your dog? *J.*
1025 *Pers. Soc. Psychol.* **65**, 613–628 (1993).
1026 46. J Baron, I Ritov, Protected values and omission bias as deontological judgments in *Psychology*
1027 *of Learning and Motivation*. (Academic Press) Vol. 50, pp. 133–167 (2009).
1028 47. JB André, L Fitouchi, S Debove, N Baumard, An evolutionary contractualist theory of morality
1029 (2022) Preprint.

1019	48. JD Greene, <i>Moral tribes: Emotion, reason, and the gap between us and them.</i> (Penguin), (2014).	1083
1020	49. S Levine, N Chater, JB Tenenbaum, FA Cushman, Resource-rational contractualism: A triple theory of moral cognition. <i>Behav. Brain Sci.</i> (2024) Online ahead of print (received Oct 28, 2024).	1084
1021		1085
1022		1086
1023	50. E Yoeli, NA Dalkiran, B Burum, MA Nowak, M Hoffman, The emergence of categorical norms. <i>Preprint</i> (2022) Preprint available on ResearchSquare.	1087
1024	51. P DeScioli, R Kurzban, A solution to the mysteries of morality. <i>Psychol. Bull.</i> 139 , 477–496 (2013).	1088
1025		1089
1026	52. M Hoffman, E Yoeli, MA Nowak, Cooperate without looking: Why we care what people think and not just what they do. <i>Proc. Natl. Acad. Sci.</i> 112 , 1727–1732 (2015).	1090
1027	53. M Singh, M Hoffman, Commitment and impersonation: A reputation-based theory of principled behavior (2021) Preprint.	1091
1028	54. MA Nowak, K Sigmund, Evolution of indirect reciprocity by image scoring. <i>Nature</i> 393 , 573–577 (1998).	1092
1029		1093
1030	55. K Panchanathan, R Boyd, A tale of two defectors: the importance of standing for evolution of indirect reciprocity. <i>J. Theor. Biol.</i> 224 , 115–126 (2003).	1094
1031	56. O Leimar, P Hammerstein, Evolution of cooperation through indirect reciprocity. <i>Proc. Royal Soc. London. Ser. B: Biol. Sci.</i> 268 , 745–753 (2001).	1095
1032		1096
1033	57. H Ohtsuki, Y Iwasa, The leading eight: social norms that can maintain cooperation by indirect reciprocity. <i>J. Theor. Biol.</i> 239 , 435–444 (2006).	1097
1034	58. H Ohtsuki, Y Iwasa, How should we define goodness?—reputation dynamics in indirect reciprocity. <i>J. Theor. Biol.</i> 231 , 107–120 (2004).	1098
1035		1099
1036	59. Y Murase, C Hilbe, Computational evolution of social norms in well-mixed and group-structured populations. <i>Proc. Natl. Acad. Sci. United States Am.</i> 121 , e2406885121 (2024).	1100
1037	60. C Hilbe, L Schmid, J Tkadlec, MA Nowak, Indirect reciprocity with private, noisy, and incomplete information. <i>Proc. Natl. Acad. Sci.</i> 115 , 12241–12246 (2018).	1101
1038		1102
1039	61. FP Santos, FC Santos, JM Pacheco, Social norm complexity and past reputations in the evolution of cooperation. <i>Nature</i> 555 , 242–245 (2018).	1103
1040	62. J Lie-Panis, JB André, Cooperation as a signal of time preferences. <i>Proc. Royal Soc. B: Biol. Sci.</i> 289 , 20212266 (2022).	1104
1041		1105
1042	63. O Leimar, Reciprocity and communication of partner quality. <i>Proc. Royal Soc. London. Ser. B: Biol. Sci.</i> 264 , 1209–1215 (1997).	1106
1043	64. JB André, The evolution of reciprocity: Social types or social incentives? <i>The Am. Nat.</i> 175 , 197–210 (2010).	1107
1044		1108
1045	65. K Panchanathan, R Boyd, Indirect reciprocity can stabilize cooperation without the second-order free rider problem. <i>Nature</i> 432 , 499–502 (2004).	1109
1046	66. E Yoeli, M Hoffman, <i>Hidden Games: The Surprising Power of Game Theory to Explain Irrational Human Behavior.</i> (Basic Books, New York), (2022).	1110
1047		1111
1048	67. TA Kessinger, CE Tarnita, JB Plotkin, Evolution of norms for judging social behavior. <i>Proc. Natl. Acad. Sci.</i> 120 , e2219480120 (2023).	1112
1049	68. JM Pacheco, FC Santos, FACC Chalub, Stern-judging: A simple, successful norm which promotes cooperation under indirect reciprocity. <i>PLoS Comput. Biol.</i> 2 , e178 (2006).	1113
1050	69. M Kawakatsu, TA Kessinger, JB Plotkin, A mechanistic model of gossip, reputations, and cooperation. <i>Proc. Natl. Acad. Sci.</i> 121 , e2400689121 (2024).	1114
1051		1115
1052	70. D Nettle, C Heintz, Studying culture from an evolutionary psychological perspective: approaches, achievements, and prospects (PsyArXiv) (2025).	1116
1053	71. Y Murase, C Hilbe, Indirect reciprocity under opinion synchronization. <i>Proc. Natl. Acad. Sci.</i> 121 , e2418364121 (2024).	1117
1054		1118
1055	72. TA Kessinger, JB Plotkin, Institutions of public judgment established by social contract and taxation. <i>Proc. Natl. Acad. Sci.</i> 122 , e2506537122 (2025).	1119
1056	73. HP Young, The evolution of social norms. <i>Annu. Rev. Econ.</i> 7 , 359–387 (2015).	1120
1057	74. R Boyd, PJ Richerson, Punishment allows the evolution of cooperation (or anything else) in sizable groups. <i>Ethol. Sociobiol.</i> 13 , 171–195 (1992).	1121
1058		1122
1059	75. D Fudenberg, J Tirole, <i>Game Theory.</i> (MIT Press, Cambridge, MA), (1991).	1123
1060		1124
1061		1125
1062		1126
1063		1127
1064		1128
1065		1129
1066		1130
1067		1131
1068		1132
1069		1133
1070		1134
1071		1135
1072		1136
1073		1137
1074		1138
1075		1139
1076		1140
1077		1141
1078		1142
1079		1143
1080		1144
1081		1145
1082		1146