Title
Moral flexibility in applying queuing norms can be explained by contractualist principles and game-theoretic considerations

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 46(0)

Authors
White, Joshua P
Bhui, Rahul
Cushman, Fiery
et al.

Publication Date
2024

Peer reviewed

Moral flexibility in applying queuing norms can be explained by contractualist principles and game-theoretic considerations

Joshua P. White[1], Rahul Bhui[1], Fiery Cushman[2], Josh Tenenbaum[3] & Sydney Levine[2,3,4]
jpwhite@mit.edu, rbhui@mit.edu, cushman@fas.harvard.edu, jbt@mit.edu.au, sydneyl@allenai.org

[1]Sloan School of Management, MIT, Cambridge, MA, 02139, USA
[2]Department of Psychology, Harvard University, Cambridge, MA, 02139, USA
[3]Department of Brain and Cognitive Sciences, MIT, Cambridge, MA, 02139, USA
[4]Allen Institute for AI

Abstract

People sometimes display moral flexibility by deciding that a commonly accepted moral norm ought not to apply in particular circumstances. But how? We explore this question in the context of queuing. We show that people's judgements about the moral permissibility of queue-cutting can be explained through cognitive processes related to moral contractualism: universalization, virtual bargaining, and functional thinking. Participants were presented vignettes depicting prospective queue-cutters, and asked whether it was morally permissible to queue-cut in those circumstances. We model these judgements with reference to the existence of a game-theoretic equilibrium supporting queue cutting in a repeated game, and to considerations of whether queue cutting would subvert or enact the function of a queue: if you pay the waiting cost, you should get the reward. These results support the notion that moral flexibility is in part related to contractualist moral principles.

Keywords: Moral permissibility; Moral flexibility; contractualism; virtual bargaining; line-cutting; game-theory.

Introduction

Rules are a large part of human moral judgement. Yet rules cannot be formulated *a priori* to apply to all situations, and people often decide that it is morally acceptable for them to be broken or ignored in certain circumstances.[1] Indeed, this flexibility of moral judgement is a necessary consequence of our complex moral judgement which depends heavily on contextual factors which change over time and place. But how do people make such judgements?

One possibility that has long been a tradition in moral philosophy (see, e.g.,Rawls, 1971; Scanlon, 2000), but is only recently being considered widely in moral cognition research (see, e.g., Levine, Kleiman-Weiner, Chater, Cushman, & Tenenbaum, 2018; Le Pargneux, Chater, & Zeitoun, 2023; Levine, Chater, Tenenbaum, & Cushman, 2023), is that these judgements are underlied by contractualist principles — moral judgements are intimately related to agreement between individuals who are faced with the problem of interdependent choice (Levine et al., 2023).

One particular model-based mechanism of contractualist moral judgement, known as virtual bargaining, proposes that people consult their internal model of bargaining with other relevant parties, asking "what would we all hypothetically agree to had we engaged in a bargaining process" as a lodestar defining their moral sentiments (Levine, Kleiman-Weiner, Chater, Cushman, & Tenenbaum, 2022). Yet, this process — virtually simulating a complex negotiation process — is resource intensive, and thus in many cases it may not be an efficient way at solving the moral problem. In such cases, Levine et al., 2023 have suggested that rather than undergoing this costly process, people may simply use cached action standards learned from previous bargaining processes (real or virtual) to inform their moral judgements.

In the present study, we consider the degree to which contractualist principles, and in particular virtual bargaining, universalization and cached action standards, are related to moral flexibility in the context of when it is morally permissible to cut in line.

One possible cached action standard that may arise in the context of queueing is that people, over repeated experience, infer the 'function' of the line, and make moral judgements in relation to whether cutting would subvert that function. Thus, people may apply the rule that "If you pay the cost of waiting in line, you should get the reward" as a heuristic in their moral judgement. This cognitive process is resource-efficient, simply requiring individuals to ensure that the reward received and the cost paid are similar.

A virtual bargaining approach to the issue of line-cutting, on the other hand, would propose that people imagine a negotiation between a prospective line cutter and the other queuers, and ask, "Would they all agree to let this person cut in the line". If the answer is "yes", by the virtual bargaining approach, it would be judged to be morally permissible (even in the absence of an actual negotiation — it is only 'virtual'). In the context of the line, however, this virtual bargaining would involve internally simulating a negotiation process between a number of different parties, who may each have different interests. So how might one have a model of when such a line-cutting agreement may arise in such complicated circumstances? We propose that game theory may help provide answers. By assuming that all individuals are rationally self-interested, and desire outcomes that best serve their preferences, game-theoretic principles can allow us to understand under what circumstances many people would be likely to agree for another to cut in front of them. We later use the existence (and 'strength') of game-theoretic equilib-

[1]Note, we do not in this article draw a sharp distinction between 'moral' and 'non-moral' rules in the social domain. Rather, we treat all social normative decision-making related to interdependent choice as 'moral'. Indeed, a universal and sharp boundary between moral and non-moral norms may not even exist (Levine et al., 2021; Stich, 2018).

In L. K. Samuelson, S. L. Frank, M. Toneva, A. Mackey, & E. Hazeltine (Eds.), *Proceedings of the 46th Annual Conference of the Cognitive Science Society*.

rium to model the virtual bargaining process and thus moral judgements. First, however, we introduce some underlying research about the morality and social dynamics of queuing, before introducing the game-theoretic model (developed by Allon & Hanany, 2012) that allows for a queue-cutting equilibrium in certain circumstances.

Queuing

Although the exact details and norms can vary, queuing is used cross-culturally to allow for the fair allocation and sharing of resources. In the United States for example, a First-in-first-out (FIFO) priority rule has been described as an underlying social norm that helps establish distributive justice, by giving preference to those who wait the longest (Schwartz, 1975; Mann, 1969). Indeed, field experiments have shown that a FIFO queuing norm is deeply embedded in American society with most unjustified cutting attempts being rebuffed (Milgram, Liberty, Toledo, & Wackenhut, 1986). Yet, field experiments have also shown that providing adequate justification improves the chances of a cutting attempt being allowed (Schmitt, Dubé, & Leclerc, 1992). This is real-world evidence of the flexibility of the moral norm of queuing — while it is heavily entrenched in US society, it can be flexibly updated to apply to new situations.

Yet, when queuers have different waiting costs, a FIFO queue is inefficient in comparison to the '$c\mu$ priority rule', wherein those with the highest waiting costs are served first (assuming the rate of service, μ, is equal). This has been known in the operations research literature for some time now, where economic incentives have been devised to allow for people to pay to cut the line and thus to achieve a socially efficient outcome (Hassin & Haviv, 2003; Erlichman & Hassin, 2015). Such findings are also backed up by empirical work, showing that when there is large variance in preferences, markets are preferred to lines for allocating resources (Shaddy & Shah, 2022). Similarly, experiments have shown that people are more willing to allow another customer to cut into line in front of them the more money they are offered (Oberholzer-Gee, 2006). Thus, self-interest and economic considerations are relevant to judgement making regarding lines. That is, social norms and punishment etc., while important (e.g., people punish people who deviate from FIFO, even when it doesn't affect outcomes; Helweg-Larsen & LoMonaco, 2008), are not inviolable. The game theory model we consider next is consistent with both these processes: self-interest is the fundamental process by which the equilibrium arises, and punishment enforces social norms when they are violated.

The Game Theory of cutting in line

Allon and Hanany (2012) have considered the question of when game-theoretic equilibria for line-cutting behaviour can arise in the context of queuing theory. Assuming the dynamics of an M/M/1 queue, they introduce the possibility of two types of queuers, high cost, c_H, and low cost, c_L, that appear in the queue with probability α and $1-\alpha$, respectively. In this scenario, the high cost queuers have a higher cost for waiting than the low cost queuers. For example, imagine person A arrives to the airport late, and has an important flight scheduled to leave in 1 hour compared to another person in the line, person B, who plans to relax in the airport lounge before their flight in 5 hours. In this case, the costs of waiting are clearly higher for person A – if they wait too long they may miss their important flight, whereas a delay is unlikely to be a large issue for person B. In the language of Allon and Hanany (2012), Person A is a high cost queuer, with c_H, and person B is a low cost queuer, with c_L.

So given such a setup, when would we expect queue-cutting behaviour to occur in a stable game-theoretic equilibrium? Clearly, in a one-shot game this is not possible. Even though it would maximize social welfare for Person B to let Person A cut the line, it is not in Person B's self-interest to do so. Interestingly, however, the authors show that in an *iterated* game an equilibrium exists that allows for line-cutting for sufficiently patient individuals. The intuition underlying this equilibrium is simple — "I'll let you cut in front of me now when your waiting costs are high, providing that, when my waiting costs are high, somebody else will do the same for me". The authors show that the conditions for such an equilibrium to arise take the form:[2]

$$\frac{\delta}{1-\delta} \geq \frac{c_L}{c_H - c_L} \frac{1}{\alpha(1-\alpha)}$$

Where δ is the discount rate by which the rewards and costs of future games are discounted. This can also be thought of as the degree of 'patience' of individuals between periods of the game (different times in the line).

Analysing this inequality, we can thus see what conditions are most conducive to a $c\mu$ priority rule equilibrium:

- When δ is almost 1. That is, the gains in future games are not discounted heavily compared to a current game.
- When the high cost queuer's waiting costs, c_H, are proportionally much larger than the low cost queuers, c_L. This means that cutting is more likely if there is a large social welfare gain available.
- The prevalence of high cost queuers (prospective line-cutters), α, is close to .5. If this prevalence is too high, those waiting in the line are likely to be overly imposed upon by continuing requests to cut the line. If the prevalence is too low, those in the line are less likely to be in a high cost situation themselves in the future when they can take advantage of the $c\mu$ priority rule.

We now use this framework to consider to what extent contractualist and virtual bargaining processes may underlie

[2]The equilibrium conditions described here require one extra assumption, that 'service rate' is equal for high and low cost queuers. In our context, this means that the prospective cutter is assumed to take as long with the teacher as the other students. For the majority of the vignettes presented in this study, this assumption is satisfied. See Allon and Hanany (2012) for further details of the full equilibrium conditions when not making this assumption.

moral judgements of line-cutting behaviour. For our purposes, we consider a potential line cutter as a high cost queuer, and all others in the queue as low cost queuers.

We report the results of two separate, but related, studies. In both studies, we examine how moral flexibility in upholding queueing norms may be related to game-theoretic equibrilia as well as inferences about how line-cutting actions may subvert a line's function, both informed by the theory of moral contractualism. In study 2, we extend this notion of 'the function of the line', to include situations where line-cutting may actually *enact* the function of a line.

Method

Participants

Data was collected from a US adult sample for two studies. Participants were recruited through Amazon Mechanical Turk and paid for their participation. Sample characteristics are given in Table 1.

Table 1: Sample Characteristics

	Study 1	Study 2
N	69	236
Final N*	48	200
Age		
M	35.94	38.01
SD	10.12	10.16
Range	22-61	21-68

*After removing subjects who failed attention checks

All participants gave informed consent to participate. All studies received ethical approval by the Committee on the Use of Human Subjects in Research at Harvard University.

Materials and Procedure

Study 1 Participants were first shown an example cartoon scenario in which children in a classroom were lining up to receive a snack. They were then shown an example animation of how the distribution of the snack works in the line – that is, in a first-in-first-out (FIFO) fashion where each child receives the snack in the order that they are lined up. Participants then completed four attention/comprehension checks to ensure sure that they were paying attention to the task and that they understood the FIFO nature of the line therein.

With the nature of the line clearly established, the participants were iteratively shown various vignettes containing variants of this scenario with children lining up to receive a snack (see Figure 1 for an example vignette). In each vignette, there is a focal child ('the cutter') who is introduced, along with a rationale for why they may want to cut to the front of the line. Participants are first asked, "Is it OK for him/her to go to the front or not OK?" (0 = *no*, 1 = *yes*; MORAL PERMISSIBILITY), before being asked, in random order, some 'evaulation variables' about the scenario measured on a Likert scale: the WAITING COST of the prospective cutter ("If he/she doesn't cut the line, how bad will it be for him/her?", 1 = *She'll be fine* to 7 = *Very bad*), and the FREQUENCY with which such a situation occurs ("How often do you think someone [situation]?", 1 = *Never* to 7 = *Always*).[3] The order that each vignette is shown to participants is randomized. In addition to these variables, we hard coded a boolean variable for whether, if the prospective cutter cuts the line, they will get a snack or not (GETS SNACK).

Figure 1: Image accompanying the following vignette: *This girl already got her snack, but she wants to get a napkin. She wants to go to the front of the line instead of waiting in the back? Is that OK or not OK?*

Study 2 The procedure for Study 2 was the same as for study 1 with a few notable differences. First, as we planned to model responses at the group-level, not all participants completed all measures. For each vignette, 141 of the participants completed the MORAL PERMISSIBILITY measure, and 40-50 completed each relevant evaluation variable.

Second, the set of vignettes was different. In addition to some of the original vignettes, we included vignettes to directly manipulate the extent to which line-cutting may either enact or subvert the lines function, which we have informally formulated as "If you pay the cost of waiting in line, you should get the reward". Thus some vignettes explicitly manipulated the 'reward received' by students who had already waited in line, and others manipulated the 'time waited' (or cost paid) before the cutting request. For example, in the 'Half cookies' vignette, the student has already waited in line and received his snack (a bag of cookies), however half of the cookies were crushed. In the 'Left for toys' vignette, the prospective cutter waits in line until she reaches the middle, after which she leaves to go play with some toys. To use these manipulations within our modeling framework, we include hardcoded variables in study 2 for the REWARD RECEIVED and the (waiting) COST PAID.

In this paper, we refer to the vignettes by short descriptive

[3] Other evaluation variables were measured, but are not used in the analyses in this paper and are thus not further reported.

titles, but the full vignettes are available on OSF here.[4]

Modeling

Framework Our modeling framework, heavily informed by contractualist theories of moral cognition, supposes that the moral flexibility to uphold queuing norms is strongly related to the degree that interested parties would agree to relax or uphold the norm. Following Levine et al., 2023, we consider that there are various cognitive processes, each with differing degrees of cognitive effort, that may be relied upon to make such judgements.

Guided by notions of universalization and virtual bargaining, we consider that moral permissibility will be related to the existence of a $c\mu$ priority norm equilibrium in the particular circumstance of a vignette because, assuming all participants are rationally self-interested, they would all agree to such an outcome. This process engages universalization by imagining what would happen if everybody in such a situation were to cut the line, but also virtual bargaining, if participants simulate the particular game theoretic dynamics / bargaining in their mind.

We also consider moral permissibility to be related to perceptions of whether the line-cutter is subverting or enacting the function of the line. We consider such processes to also be related to contractarian moral judgement, but in a more cached, resource rational way, what Levine et al., 2023 call 'cached action standards'. This is a model-free heuristic, that saves the agent the cognitive cost of simulating a complex negotiation with many parties, by distilling these dynamics into a lines purpose. Here, participants rather ask whether, if the prospective cutter was to cut the queue, the reward received and the cost paid are concomitant. If so, the function of the line would not be subverted by allowing the queue cutting, and it is thus judged not permissible to do so.

Models For study 1, there are two quantities that allow us to model the above dynamics, EQUILIBRIUM STRENGTH, and GETS SNACK. EQUILIBRIUM STRENGTH, E_s, measures the degree to which the scenario is in a game-theoretic equilibrium supporting the $c\mu$ priority rule, given by:

$$E_s = \frac{\delta}{1-\delta} - \frac{c_L}{c_H - c_L}\frac{1}{\alpha(1-\alpha)} \quad (1)$$

GETS SNACK, s, is a binary variable that indexes whether the prospective line cutter would get a snack if they were allowed to cut, allowing them to subvert the function of the line by getting the reward without paying any of the waiting cost.

Combining these two processes together, we model the moral permissibility of line cutting as a stochastic function of EQUILIBRIUM STRENGTH and GETS SNACK as follows:

$$P(\text{Cutting Permissible}) = \frac{1}{1+e^{-[w_1 E_s + w_2 s + \beta]}} \quad (2)$$

The weights, w_1 & w_2, govern the strength of the effect of EQUILIBRIUM STRENGTH and GETS SNACK (respectively) on moral permissibility judgements, and the bias, β, governs whether people err on the side of reiterating the norm (i.e, cutting is *not* permissible) or on allowing exceptions to the norm (i.e., cutting *is* permissible).

For study 2, we introduced further elements into the vignettes, by deliberately manipulating the degree to which students had 'paid the cost' and 'got the reward' in a previous line, to see whether they are sensitive, in a graded way, to whether cutting will subvert or enact the function of the line. We thus augmented the model as:

$$P(\text{Cutting Permissible}) = \frac{1}{1+e^{-[w_1 E_s + w_2 s + s w_3 (r-c) + \beta]}} \quad (3)$$

where E_s, s, w_1, w_2 and β are as above. w_3 is the weight parameter for $r-c$, where r is the reward already received, and c is the cost already paid (in terms of time waiting in the line already) normed to 1. For example, the 'Half cookies' vignette in which a child has just received a half full bag of cookies after having waited their full turn in line will be coded $r = 0.5$ and $c = 1$, and thus $r - c = -0.5$

We parameterise these models as follows:

- We get α, by transforming the FREQUENCY Likert variable to a probability with the function $1/(1+e^{-(freq-4)})$.[5]
- We use the WAITING COST Likert variable for C_H, and set C_L to 1. In the game theoretic model, these values are both considered as costs per unit of time of waiting. However, what is most important is not the exact values of C_L and C_H, but the ratio between them. Given the wording of the Likert scale for the WAITING COST variable, this specification is likely to maintain this relationship.
- For models that include E_s, we set β at 0 so that, *ceteris paribus*, $P(\text{Cutting Permissible}) > 0.5$ for situations that are in a $c\mu$ priority scheme equilibrium (i.e, E_s is positive).
- All remaining parameters (w_1, w_2, w_3, δ) were estimated via non-linear least square minimization as implemented by the Levenberg-Marquardt Algorithm in the R package `minpack.lm` (Elzhov, Mullen, Spiess, & Bolker, 2023).

Model comparison We compare the fit of the main model in study 1 (equation 2; 'S1 main model') to simpler models that include either E_s ('GT model') or s ('Function model') but not both, to ascertain whether both factors are necessary. In study 2, we compare our main model (equation 3; 'S2 main model') to the main model from study 1 ('S1 model') as well as a model that omits $w_2 s$; 'No snack').

In both studies, we also compare our main models to a linear regression model that includes the same variables, but importantly, not the functional form underlying E_s. This helps ascertain the degree to which this functional form, which has more cognitively interpretable parameters and relates to game-theoretic dynamics, adds predictive utility rather than

[4] https://osf.io/kh69f/

[5] An error in data collection meant that for 3 vignettes in study 1, FREQUENCY was collected on a Likert scale from 1-8. In these cases, the transformation $1/(1+e^{-(freq-4.5)})$ is used.

a simple linear model. Models were compared with the Bayesian Information Criterion (BIC) to account for model complexity.

Results

Figure 2 compares the results of the various models to human performance in both study 1 (Panel A) and study 2 (Panel B). The main models for both studies show high in-sample correlation ($rs > 0.9$) with mean human judgements of the permissibility of line-cutting, $P(OK)$, across the vignettes. Table 2 displays the number of parameters and BIC for all models (lowest BICs bolded). Inspection of this table reveals that the main model is preferred in BIC in both studies 1 and 2. In all cases, our preferred models outperform the simple regression models, both in terms of in-sample correlation and by BIC.

Table 2: Model Results

	k	BIC
Study 1		
GT model	2	0.9
Function model	2	-12.5
S1 main model	3	**-21.2**
Regression	4	-16.1
Study 2		
S1 main model	3	-1.4
No snack	3	2.5
S2 main model	4	**-16.0**
Regression	5	-10.5

Discussion

This paper modeled the flexibility of moral judgements using principles from the theory of moral contractualism – that moral cognition proceeds by asking not directly about harms and duties, but what would people agree to if they negotiated. Using line-cutting as a case study, and borrowing from game-theoretic queueing theory, we find evidence that these principles can explain the flexibility of moral judgements in this context. Our approach was to use the existence of game-theoretic equilibria as diagnostic of situations where all parties to the line would agree to the cutting-attempt had they all engaged in a negotiation process, because in the long run it would be better for all parties, and no party could reliably be better off (in the long run) by taking another strategy. We tested whether participants' judgements were related to the strength of equilibria, and found evidence that they were. This could be seen as evidence of a virtual bargaining process, wherein participants simulated a negotiation process over time (i.e, an iterated game), deciding whether each individual would be better off under, and would thus agree to, a rule that allows for people to cut into line when they have high waiting costs. The process may also involve universalization — deciding whether the world in which everybody in a particular situation felt at liberty to cut in line results in better outcomes than the world in which nobody does.

We also found strong evidence that people's judgements were very sensitive to violations of the lines' function. Across both studies 1 and 2, people were much less likely to endorse line-cutting in situations when the line was being cut in order to for the child to get a snack. This may be that people have internalized cached action standards which help enforce the function of a line: ensuring that "if you pay the cost, you get the reward".

In study 2, we manipulated the reward and cost in a graded way to see whether, in certain circumstances, cutting may be allowed to help 'enact' the purpose of the line (say, by allowing a kid who lined up but got a defective snack to get another snack without having to line up again). Our results showed that people were responsive to this graded manipulation.

Limitations and Future Directions

The data presented herein, while helping to further our understanding of the different processes involved in the flexibility of the application of moral rules, leave unanswered many of the deepest questions in the cognitive science of morality.

One of these, is the study of individual difference in moral cognition. When looking into the data, it can be found that there is large individual difference in the pattern of responses among participants. While our present model assumes that participants all undertake the same computation/moral strategy and treats any inter-individual differences as stochastic, this is unlikely. The data rather shows clear differences between participants in this regard. Indeed, this is similar to other studies in moral cognition which also show large differences between people in the strategies or computations used when making moral judgements (see e.g., Andrejević, White, Feuerriegel, Laham, & Bode, 2022; Levine, Kleiman-Weiner, Schulz, Tenenbaum, & Cushman, 2020). Further study looking to model this data at an individual level may be fruitful, to try to understand if there are person-centric or situation-centric factors which can explain the strategy variance, or indeed if people combine different moral strategies in some way when making their judgements. One interesting direction would be to ascertain whether different information environments can bias people to engage in different moral strategies.

Another limitation of this work is the application of queuing theory to the present scenario. In particular, some assumptions of an M/M/1 queue do not perfectly apply to the scenario as shown to participants. Most pertinently, the vignettes presented a wider array of circumstances and factors than those that are implemented in the M/M/1-based game-theoretic model of Allon & Hanany, 2012; such as pro-social action, the possibility of being 'at fault', and other social benefits or costs that could arise from lining up in a classroom scenario (such as wanting to stand near friends, or away from bullies). While we attempted in our modelling to account for some of these other factors, by adding further parameters to our models (e.g., GETS SNACK), we believe further work could do so in a more principled and cognitively meaningful and unified way. Future work could ensure greater correspondence between the vignettes and the assumptions underlying

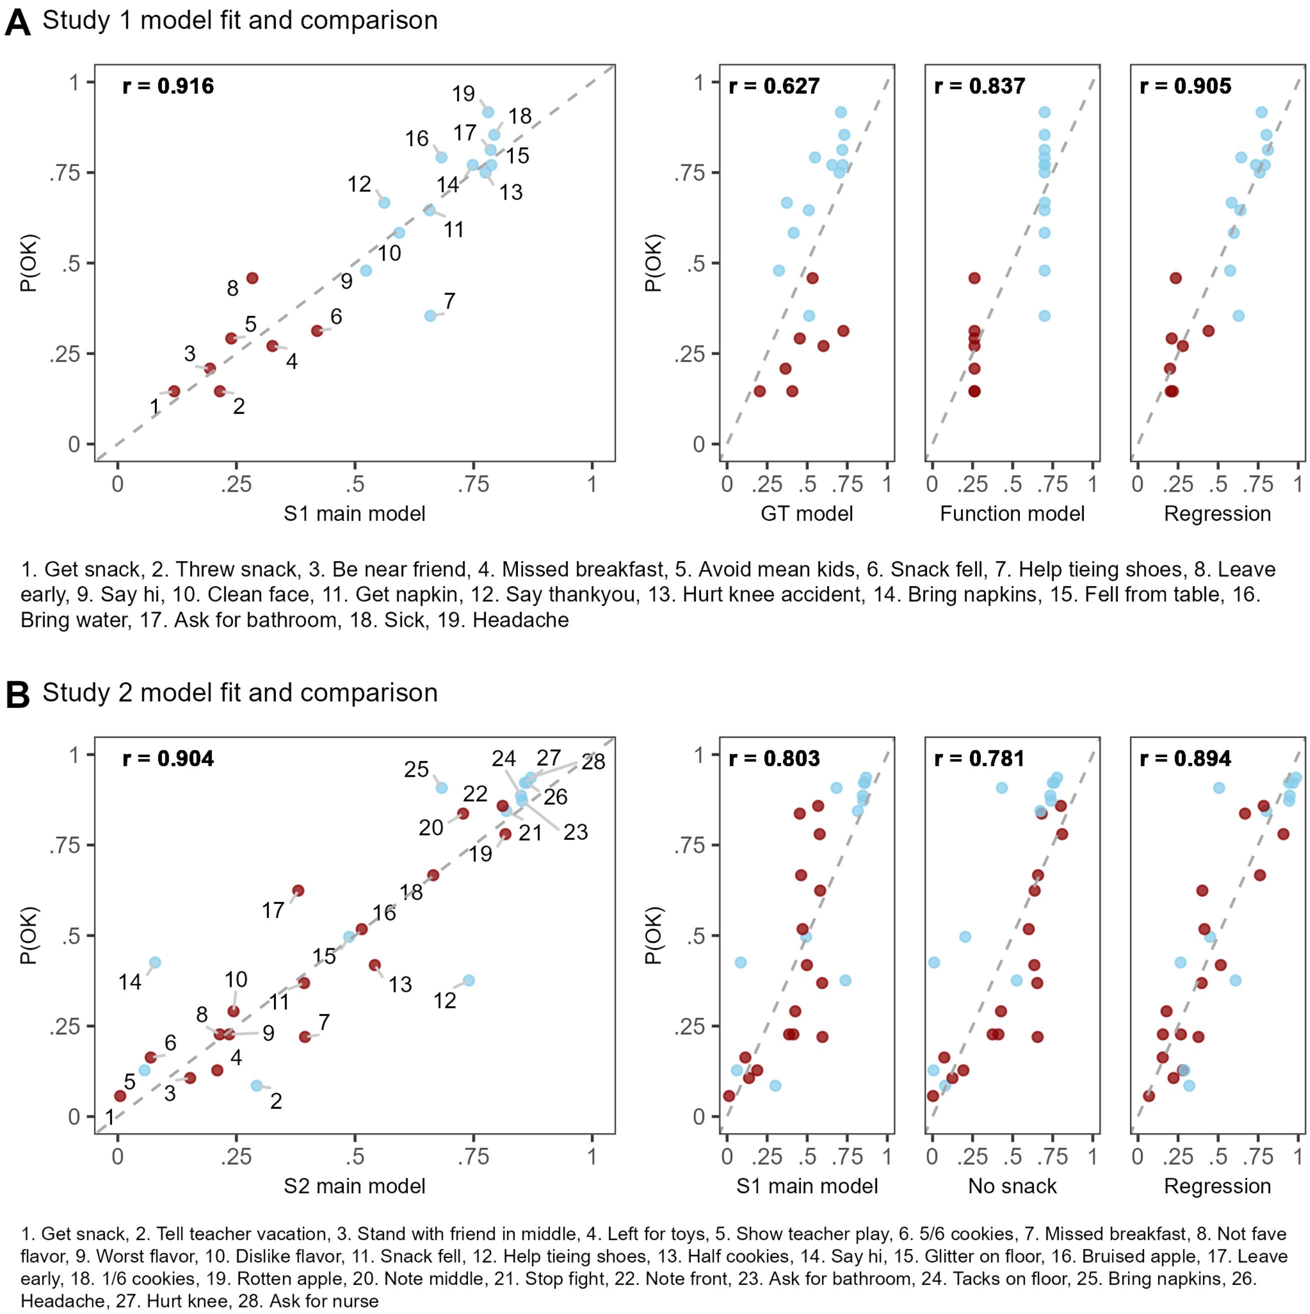

Figure 2: Comparing results of the main models (left), and alternative models (right) for study 1 (Panel A) and study 2 (Panel B). Colour indicates the value of GETS SNACK, red = yes, blue = no.

the game-theoretic models, whether that be by explicitly designing vignettes to accord with such assumptions, or developing new game-theoretic accounts of the vignettes presented herein.

Conclusion

Overall, this paper provides some early steps to combining game-theoretic insights with the toolkit of computational cognitive science to help understand the flexibility of moral cognition from a contractualist perspective, providing evidence that people engage in contractualist reasoning when flexibly considering exceptions to common moral norms such as queuing.

Acknowledgments

This work was supported by ARO MURI: Foundations of Decision Making with Behavioral and Computational Constraints (grant number W911NF-19-1-0217) and the Office of Naval Research (Grant no. N00014-21-1-2170).

References

Allon, G., & Hanany, E. (2012, March). Cutting in Line: Social Norms in Queues. *Management Science*, *58*(3), 493–506. doi: 10.1287/mnsc.1110.1438

Andrejević, M., White, J. P., Feuerriegel, D., Laham, S., & Bode, S. (2022, June). Response time modelling reveals evidence for multiple, distinct sources of moral decision caution. *Cognition*, *223*, 105026. doi: 10.1016/j.cognition.2022.105026

Elzhov, T. V., Mullen, K. M., Spiess, A.-N., & Bolker, B. (2023, September). *Minpack.lm: R Interface to the Levenberg-Marquardt Nonlinear Least-Squares Algorithm Found in MINPACK, Plus Support for Bounds.*

Erlichman, J., & Hassin, R. (2015, August). Strategic Overtaking in a Monopolistic M/M/1 Queue. *IEEE Transactions on Automatic Control*, *60*(8), 2189–2194. doi: 10.1109/TAC.2015.2419151

Hassin, R., & Haviv, M. (2003). *To queue or not to queue: Equilibrium behavior in queueing systems* (Vol. 59). Springer Science & Business Media.

Helweg-Larsen, M., & LoMonaco, B. L. (2008). Queuing Among U2 Fans: Reactions to Social Norm Violations1. *Journal of Applied Social Psychology*, *38*(9), 2378–2393. doi: 10.1111/j.1559-1816.2008.00396.x

Le Pargneux, A., Chater, N., & Zeitoun, H. (2023, December). Contractualist inclinations and reasoning in moral judgment and decision making.
doi: 10.31234/osf.io/p4cyx

Levine, S., Chater, N., Tenenbaum, J., & Cushman, F. (2023, May). *Resource-rational contractualism: A triple theory of moral cognition.* PsyArXiv. doi: 10.31234/osf.io/p48t7

Levine, S., Kleiman-Weiner, M., Chater, N., Cushman, F., & Tenenbaum, J. (2018). The Cognitive Mechanisms of Contractualist Moral Decision-Making. *Cognitive Science*.

Levine, S., Kleiman-Weiner, M., Chater, N., Cushman, F., & Tenenbaum, J. (2022, June). *When rules are over-ruled: Virtual bargaining as a contractualist method of moral judgment.* PsyArXiv. doi: 10.31234/osf.io/k5pu8

Levine, S., Kleiman-Weiner, M., Schulz, L., Tenenbaum, J., & Cushman, F. (2020, October). The logic of universalization guides moral judgment. *Proceedings of the National Academy of Sciences*, *117*(42), 26158–26169. doi: 10.1073/pnas.2014505117

Levine, S., Rottman, J., Davis, T., O'Neill, E., Stich, S., & Machery, E. (2021). Religious Affiliation and Conceptions of the Moral Domain. *Social Cognition*, *39*(1), 139–165. doi: 10.1521/soco.2021.39.1.139

Mann, L. (1969). Queue Culture: The Waiting Line as a Social System. *American Journal of Sociology*, *75*(3), 340–354.

Milgram, S., Liberty, H. J., Toledo, R., & Wackenhut, J. (1986, October). Response to intrusion into waiting lines. *Journal of Personality and Social Psychology*, *51*(4), 683–689. doi: 10.1037/0022-3514.51.4.683

Oberholzer-Gee, F. (2006). A Market for Time Fairness and Efficiency in Waiting Lines. *Kyklos*, *59*(3), 427–440. doi: 10.1111/j.1467-6435.2006.00340.x

Rawls, J. (1971). Atheory of justice. *Cambridge (Mass.).*

Scanlon, T. M. (2000). *What we owe to each other*. Harvard University Press.

Schmitt, B. H., Dubé, L., & Leclerc, F. (1992). Intrusions into waiting lines: Does the queue constitute a social system? *Journal of Personality and Social Psychology*, *63*(5), 806–815. doi: 10.1037/0022-3514.63.5.806

Schwartz, B. (1975). *Queuing and Waiting*. Chicago: University of Chicago Press.

Shaddy, F., & Shah, A. K. (2022, May). When to Use Markets, Lines, and Lotteries: How Beliefs About Preferences Shape Beliefs About Allocation. *Journal of Marketing*, *86*(3), 140–156. doi: 10.1177/00222429211012107

Stich, S. (2018). The Quest for the Boundaries of Morality *. In *The Routledge Handbook of Moral Epistemology.* Routledge.