

From Outcome-Based to Language-Based Preferences[†]

VALERIO CAPRARO, JOSEPH Y. HALPERN, AND MATJAŽ PERC*

We review the literature on models that try to explain human behavior in social interactions described by normal-form games with monetary payoffs. We start by covering social and moral preferences. We then focus on the growing body of research showing that people react to the language with which actions are described, especially when it activates moral concerns. We conclude by arguing that behavioral economics is in the midst of a paradigm shift toward language-based preferences, which will require an exploration of new models and experimental setups. (JEL C70, C90, D11, D90, Z13)

1. Introduction

We review the literature on models of human behavior in social interactions that can be described by normal-form games with monetary payoffs. This is certainly a limited set of interactions, since many interactions are neither one-shot nor simultaneous, nor do they have a social element, nor do they involve just monetary payoffs. Although small, this set of interactions is of great interest from both practical and theoretical perspectives. For example, it includes

games such as the prisoner's dilemma and the dictator game, which capture the essence of some of the most fundamental aspects of our social life, such as cooperation and altruism. Economists have long recognized that people do not always behave so as to maximize their monetary payoffs in these games. Finding a good model that explains behavior has driven much of the research agenda over the years.

Earlier work proposed that we have *social preferences*; that is, our utility function depends not just on our payoffs, but also on the payoffs of others. For example, we might prefer to minimize inequity or to maximize the sum of the monetary payoffs, even at a cost to ourselves. However, a utility function based on social preferences is still *outcome based*; that is, it is still a function of the players' payoffs. We review a growing body of experimental literature showing that outcome-based utility functions cannot adequately describe the whole range of

*Capraro: Department of Economics, Middlesex University; Halpern: Computer Science Department, Cornell University; Perc: Faculty of Natural Sciences and Mathematics, University of Maribor; Department of Medical Research, China Medical University Hospital, China Medical University; and Complexity Science Hub Vienna. Halpern's work supported in part by NSF grants IIS-178108 and IIS-1703846 and MURI grant W911NF-19-1-0217.

[†] Go to <https://doi.org/10.1257/jel.20221613> to visit the article page and view author disclosure statement(s).

human behavior in social interactions. The problem is not with the notion of maximizing expected utility. Even if we consider decision rules other than maximizing the expected utility, such as minimizing regret or maximin, we cannot explain many experimental results as long as the utility is outcome based. Nor does it help to take bounded rationality into account.

The experimental evidence suggests that people have *moral preferences*, that is, preferences for doing what they view as the “right thing.” These preferences cannot be expressed solely in terms of monetary payoffs. We review the literature and discuss attempts to construct a utility function that captures peoples’ moral preferences. We then consider more broadly the issue of language. The key takeaway message is that what matters is not just the monetary payoffs associated with actions, but also how these actions are described. A prominent example of this is when the words being used to describe the strategies activate moral concerns.

The review is structured as follows. In section 2, we review the main experimental regularities that were responsible for the paradigm shift from monetary maximization models to outcome-based social preferences and then to non-outcome-based moral preferences. Motivated by these empirical regularities, in the next sections we review the approaches that have been taken to capture human behavior in one-shot interactions and, for each of them, we describe their strengths and weaknesses. We start with social preferences (section 3) and moral preferences (section 4). We also discuss experimental results showing that the words used to describe the available actions impact behavior. In section 5, we review work on games where the utility function depends explicitly on the language used to describe the available actions. We conclude in section 6 with some discussion of potential directions for future

research. Thinking in terms of moral preferences and, more generally, language-based preferences suggests potential connections to and synergies with other disciplines. Work on moral philosophy and moral psychology could help us understand different types and dimensions of moral preferences, and work in computational linguistics on sentiment analysis (Pang, Lee, and Vaithyanathan 2002) could help explain how we obtain utilities on language.

2. *Experimental Regularities*

The goal of this section is to review a series of experimental regularities that were observed in normal-form games played among anonymous players. Although we occasionally mention the literature on other types of games, the main focus of this review is on one-shot, simultaneous-move, anonymous games.

We start by covering standard experiments in which some people have been shown not to act so as to maximize their monetary payoff. Then we move to experiments in which people have been shown not to act according to any outcome-based preference. Finally, we describe experiments suggesting that people’s preferences take into account the words used to describe the available actions.

The fact that some people do not act so as to maximize their monetary payoff was first shown using the *dictator game*. In this game,¹ the *dictator* is given a certain sum of money and has to decide how much, if any, to give to the *recipient*, who starts with nothing. The recipient has no choice and receives only the amount that the dictator decides to give. Since dictators have no monetary incentives to give, a payoff-maximizing dictator would

¹Following the standard approach, we abuse terminology slightly and call this a game, although it does not specify the utility functions, only the outcome associated with each strategy profile.

keep the whole amount. However, experiments have repeatedly shown that people violate this prediction (Kahneman, Knetsch, and Thaler 1986; Forsythe et al. 1994). Moreover, the distribution of giving tends to be bimodal, with peaks at the zero-offer and at the equal share (Engel 2011).

We can summarize the first experimental regularity as follows:

Experimental regularity 1. A significant number of dictators give some money in the dictator game. Moreover, the distribution of donations tend to be bimodal, with peaks at zero and at half the total.

Another classical game in which people often violate the payoff-maximization assumption is the *ultimatum game*. In its original formulation, the ultimatum game is not a normal-form game, but an extensive-form game, where players move sequentially: a *proposer* is given a sum of money and has to decide how much, if any, to offer to the *responder*. The responder can either accept or reject the offer. If the offer is accepted, the proposer and the responder are paid according to the accepted offer; if the offer is rejected, neither the proposer nor the responder receive any payment. A payoff-maximizing responder clearly would accept any amount greater than 0; knowing this, a payoff-maximizing proposer would offer the smallest positive amount available in the choice set. Behavioral experiments showed that people dramatically violate the payoff-maximizing assumption: responders typically reject low offers and proposers often offer an equal split (Güth, Schmittberger, and Schwarze 1982; Camerer 2003). Rejecting low offers is impossible to reconcile with a theory of payoff maximization. Making a nonzero offer is consistent with payoff maximization, if a proposer believes that the responder will reject too low an offer. However, several researchers have noticed that offers are typically larger than

the amount that proposers believe would result in acceptance (Henrich et al. 2001; Lin and Sunder 2002). This led Camerer (2003, p. 56) to conclude that “some of [proposer’s] generosity in ultimatum games is altruistic rather than strategic.” These observations have been replicated in the normal form, simultaneous-move variant of the ultimatum game that is the focus of this article. In this variant, the proposer and the responder simultaneously choose their offer and minimum acceptable offer, respectively, and then are paid only if the proposer’s offer is greater than or equal to the responder’s minimum acceptable offer.²

Experimental regularity 2. In the ultimatum game, a substantial proportion of responders reject nonzero offers and a significant number of proposers offer an equal split.

The fact that some people do not act so as to maximize their monetary payoff was also observed in the context of (one-shot, anonymous) *social dilemmas*. Social dilemmas are situations in which there is a conflict between the individual interest and the interest of the group (Hardin 1968, Ostrom 1990, Olson 2009). The most-studied social dilemmas are the *prisoner’s dilemma* and the *public-goods game*.³

In the prisoner’s dilemma, two players can either cooperate (*C*) or defect (*D*). If both players cooperate, they receive the reward for cooperation, *R*; if they both defect, they receive the punishment payoff, *P*; if one

²Some authors have suggested that the standard sequential-move ultimatum game elicits slightly lower rejection rates than its simultaneous-move variant (Schotter, Weigelt, and Wilson 1994; Blount 1995), but this does not affect the claim that proposers offer more than necessary from a purely monetary point of view.

³Other well-studied social dilemmas are the *Bertrand competition* (Bertrand 1883) and the *traveler’s dilemma* (Basu 1994). In this review, we focus on the prisoner’s dilemma and the public-goods game; the other social dilemmas do not give rise to conceptually different results, at least within our domain of interest.

player defects and the other cooperates, the defector receives the temptation payoff, T , whereas the cooperator receives the sucker's payoff, S . Payoffs are assumed to satisfy the inequalities: $T > R > P > S$. These inequalities guarantee that the only Nash equilibrium is mutual defection, since defecting strictly dominates cooperating. However, mutual defection gives players a payoff smaller than mutual cooperation.

In the public-goods game, each of n players is given an endowment $e > 0$ and has to decide how much, if any, to contribute to a public pool. Let $c_i \in [0, e]$ be player i 's contribution. Player i 's monetary payoff is $e - c_i + \alpha \sum_{j=1}^n c_j$, where $\alpha \in (1/n, 1)$ is the marginal return for cooperation, that is, the proportion of the public good that is redistributed to each player. Since $\alpha \in (1/n, 1)$, players maximize their individual monetary payoff by contributing 0, but if they do that, they receive less than the amount they would receive if they all contribute their whole endowment ($e < \alpha ne$). Although cooperation is not individually optimal in the prisoner's dilemma or the public-goods game, many people cooperate in behavioral experiments using these protocols (Rapoport and Chamah 1965, Ledyard 1995).

Experimental regularity 3. A significant number of people cooperate in the one-shot prisoner's dilemma and the one-shot public-goods game.

In section 3, we show that these regularities can be explained well by outcome-based preferences, that is, preferences that are a function of the monetary payoffs. But we now discuss a set of empirical findings that cannot be explained by outcome-based preferences. We start with truth telling in tasks in which people can increase their monetary payoff by misreporting private information.

Economists have considered several ways of measuring the extent to which people lie, the most popular ones being

the *sender–receiver game* (Gneezy 2005) and the *die-under-cup task* (Fischbacher and Föllmi-Heusi 2013). In its original version, the sender–receiver game is not a normal-form game, but an extensive-form game. There are two possible monetary distributions, called option A and option B; only the *sender* is informed about the payoffs corresponding to these distributions. The sender can then tell the *receiver* either “Option A will earn you more money than option B,” or “Option B will earn you more money than option A.” Finally, the receiver decides which of the two options to implement.⁴ Clearly, only one of the messages that the sender can send is truthful. Gneezy (2005) showed that many senders tell the truth, even when the truthful message does not maximize the sender's monetary payoff. Gneezy also showed that this honest behavior is not driven by purely monetary preferences over monetary outcomes, since people behaved differently when asked to choose between the same monetary options when there was no lying involved (i.e., when they played a dictator game that was monetarily equivalent to the sender–receiver game), suggesting that people find lying intrinsically costly.

In the die-under-cup task, participants roll a dice under a cup (i.e., privately) and are asked to report the outcome. Participants receive a payoff that depends on the reported outcome, not on the actual outcome. The actual outcome is typically not known to the experimenter, however, by comparing the distribution of the reported outcomes to the uniform distribution, the experimenter

⁴The original variant of the sender–receiver game therefore requires the players to choose their action sequentially. Similar results can be obtained with the normal-form, simultaneous-move variant in which the receiver decides whether to believe the sender's message at the same time that the sender sends it, or even when the receiver makes no active choice (Gneezy, Rockenbach, and Serra-Garcia 2013; Biziou-van-Pol et al. 2015).

can tell approximately what fraction of people lied. Several studies showed that people tend to be honest, even if this goes against their monetary interest (Fischbacher and Föllmi-Heusi 2013; Abeler, Nosenzo, and Raymond 2019; Gerlach, Teodorescu, and Hertwig 2019).

Experimental regularity 4. A significant number of people tell the truth in the sender–receiver game and in the die-rolling task, even if it lowers their monetary payoff.

Another line of empirical work that is hard to reconcile with preferences over monetary payoffs is observed in variants of the dictator game. For example, List (2007) explored people’s behavior in two modified dictator games. In the control condition, the dictator was given \$10.00 and the receiver was given \$5.00; the dictator could then give any amount between \$0 and \$5.00 to the recipient. In line with studies using the standard dictator game, List observed that about 70 percent of the dictators give a nonzero amount, with a peak at \$2.50. In the experimental condition, List added a “take” option: dictators were allowed to take \$1.00 from the recipient.⁵ In this case, List found that the peak at giving \$2.50 disappears and that the distribution of choices becomes unimodal, with a peak at giving \$0. However, only 20 percent of the participants choose the take option. This suggests that, for some participants, giving a positive amount dominates giving \$0 in the baseline, but giving \$0 dominates giving the same positive amount in the treatment. This is clearly inconsistent with outcome-based preferences. Bardsley (2008) and Cappelen et al. (2013) obtained a similar result.

⁵List also considered a treatment with multiple take options, up to \$5.00. The results are similar to those with a single take option of \$1.00, so we focus here only on the latter case.

Experimental regularity 5. A significant number of people prefer giving over keeping in the standard dictator game without a take option, but prefer keeping over giving in the dictator game with a take option.

In a similar vein, Lazear, Malmendier, and Weber (2012) showed that some dictators give part of their endowment when they are constrained to play a dictator game, but, given the option of receiving the maximum amount of money they could get by playing the dictator game without actually playing it, they choose to avoid the interaction. Indeed, Dana, Cain, and Dawes (2006) found that some dictators would even pay \$1.00 in order to avoid the interaction. Clearly, these results are inconsistent with outcome-based preferences.

Experimental regularity 6. A significant number of people prefer giving over keeping in the standard dictator game without an exit option, but prefer keeping over giving in the dictator game with an exit option.

How a game is framed is also well-known to affect people’s behavior. For example, contributions in the public-goods game depend on whether the game is presented in terms of positive externalities or negative ones (Andreoni 1995), rates of cooperation in the prisoner’s dilemma depend on whether the game is called “the community game” or “the Wall Street game” (Ross and Ward 1996), and using terms such as “partner” or “opponent” can affect participants’ behavior in the trust game (Burnham, McCabe, and Smith 2000). Burnham, McCabe, and Smith (2000) suggested that the key issue (at least, in these situations) is what players perceive as the norms.

Following this suggestion, there was work exploring the effect of changing the norms on people’s behavior. One line explored dictators’ behavior in variants of the dictator game in which the initial endowment is not given to the dictator, but is instead given to the

recipient, or equally shared between the dictator and the recipient, and the dictator can take some of the recipient's endowment; this is called the "take" frame. Some experiments showed that people tend to be more altruistic in the dictator game in the "take" frame than in the standard dictator game (Swope et al. 2008, Krupka and Weber 2013); moreover, this effect is driven by the perception of what the socially appropriate action is (Krupka and Weber 2013). However, this result was not replicated in other experiments (Dreber et al. 2013, Eckel and Grossman 1996, Halvorsen 2015, Hauge et al. 2016). A related stream of papers pointed out that including morally loaded words in the instructions of the dictator game can impact dictators' giving (Brañas-Garza 2007; Capraro and Vanzo 2019; Capraro et al. 2019; Chang, Chen, and Krupka 2019), and that the behavioral change can be explained by a change in the perception of what dictators think the morally right action is (Capraro and Vanzo 2019). Although there is debate about whether the "take" frame can impact people's behavior in the dictator game, there is general agreement that the wording of the instructions can impact dictators' behavior by activating moral concerns.

The fact that the wording of the instructions can impact behavior has also been observed in other games. For example, Eriksson et al. (2017) found that the language in which the rejection option is presented significantly impacts rejection rates in the ultimatum game. Specifically, receivers are more likely to decline an offer when this option is labeled "rejecting the proposer's offer" than when it is labeled "reducing the proposer's payoff." Moreover, in line with the discussion above regarding the dictator game, Eriksson et al. found this effect to be driven by the perception of what the morally right action is.

A similar result was obtained in the *trade-off game* (Capraro and Rand 2018,

Tappin and Capraro 2018), where a decision-maker has to unilaterally decide between two allocations of money that affect the decision-maker and two other participants. One allocation equalizes the payoffs of the three participants; the other allocation maximizes the sum of the payoffs, but is unequal. Capraro and Rand (2018) and Tappin and Capraro (2018) found that minor changes in the language with which the actions are presented significantly impacts decision-makers' behavior. For example, naming the efficient choice "more generous" and the equitable choice "less generous" makes subjects more likely to choose the efficient choice, while naming the efficient choice "less fair" and the equitable choice "more fair" makes subjects more likely to choose the equitable choice.

Morally loaded words have also been shown to affect behavior in the prisoner's dilemma, where participants have been observed to cooperate at different rates depending on whether the strategies are named "I cooperate/I cheat" or "A/B" (Mieth, Buchner, and Bell 2021). Furthermore, in both the one-shot and iterated prisoner's dilemma, moral suasion, that is, providing participants with cues that make the morality of an action salient, increases cooperation (Capraro et al. 2019, Dal Bó and Dal Bó 2014). This suggests that cooperative behavior is partly driven by a desire to do what is morally right.

Experimental regularity 7. Behavior in several experimental games, including the dictator game, the prisoner's dilemma, the ultimatum game, and the trade-off game, depends on the instructions used to introduce the games, especially when they activate moral concerns.

3. *Social Preferences*

In order to explain the seven regularities listed in section 2, one has to go beyond preferences for maximizing monetary payoffs. The first generation of utility functions that

do this appeared in the 1970s. These *social preferences* share the underlying assumption that the utility of an individual depends not only on the individual's monetary payoff, but also on the monetary payoff of the other players involved in the interaction.⁶ In this sense, these social preferences are all particular instances of *outcome-based preferences*, that is, utility functions that depend only on (i) the individuals involved in the interaction and (ii) the monetary payoffs associated with each strategy profile. Social preferences typically explain experimental regularities 1–3 well. However, they have difficulties with experimental regularities 4–7. In this section, we review this line of work. For more comprehensive reviews, we refer the readers to Camerer (2003) and Dhami (2016). Part of this ground was also covered by Sobel (2005).

Economists have long recognized the need to include other-regarding preferences in the utility function. Earlier work focused on economies with one private good and one public good. In these economies, there are n players; player i is endowed with wealth w_i , which she can allocate to the private good or the public good. This is a quite general class of games (e.g., the dictator game, prisoner's dilemma, and the public-goods game

can all be expressed in this form), although it does not cover several other games of interest for this review (e.g., the ultimatum and trade-off games). Let x_i and g_i be i 's allocation to the private good and contribution to the public good, respectively. Economists first assumed that i 's utility u_i depended only and monotonically on x_i and $G = \sum g_j$. According to this model, the government forcing an increase of contributions to public goods (e.g., by increasing taxes) will result in a decrease of private contributions, dollar for dollar. Specifically, if the government takes one dollar from a particular contributor and puts it in the public good while keeping everything else fixed (say, by changing the tax structure appropriately), then that contributor can restore the equilibrium by reducing his contribution by one dollar (Warr 1982, Roberts 1984, Bernheim 1986, Andreoni 1988). The prediction that this would happen was violated in empirical studies (Abrams and Schitz 1978, Abrams and Schmitz 1984). Motivated by these limitations, Andreoni (1990) introduced a theory of *warm-glow giving*, where the utility function captures the intuition that individuals receive positive utility from the very act of giving to the public good. Formally, this corresponds to considering, instead of a utility function of the form $u_i = u_i(x_i, G)$, one of the form $u_i = u_i(x_i, G, g_i)$. Note that the latter utility function is still outcome based, because all of its arguments are functions of monetary outcomes. Warm-glow theory has been applied successfully to several domains. However, when it comes to explaining the experimental regularities listed in section 2, it has two significant limitations. The first is its domain of applicability: the ultimatum and trade-off games cannot be expressed in terms of economies with one private good and one public good in any obvious way. The second is more fundamental: as we show at the end of this section, it cannot explain

⁶There has also been work on explaining human behavior in terms of bounded rationality. The idea behind this approach is that computing a best response may be computationally difficult, so players do so only to the best of their ability. Although useful in many domains, these models do not explain deviations from the payoff-maximizing strategy in situations in which computing this strategy is obvious, such as in the dictator game (experimental regularity 1). While some people may give in the dictator game because they incorrectly computed the payoff-maximizing strategy, did not read the instructions, or played randomly, this does not begin to explain the overall behavior of dictators. In a recent analysis of over 3,500 dictators, all of whom had correctly answered a comprehension question regarding which strategy maximizes their monetary payoff, Brañas-Garza, Capraro, and Rascón-Ramírez (2018) found an average donation of 30.8 percent. Similar observations also apply to the other experimental regularities listed in section 2.

experimental regularities 4–7 because it is outcome based.

More recently, economists have started defining the utility function directly on the monetary payoffs of the players involved in the interaction. These utility functions, by construction, can be applied to any economic interaction. The simplest such utility function is just a linear combination of the individual's payoff and the payoffs of the other players (Ledyard 1995). Formally, let (x_1, \dots, x_n) be a monetary allocation among n players. The utility of player i given this allocation is

$$u_i(x_1, \dots, x_n) = x_i + \alpha_i \sum_{j \neq i} x_j,$$

where α_i is an individual parameter representing i 's level of altruism. Preferring to maximize payoff is the special case where $\alpha_i = 0$. Players with $\alpha_i > 0$ care positively about the payoff of other players. Consequently, this utility function is consistent with altruistic behavior in the dictator game and with cooperative behavior in social dilemmas. Players with $\alpha_i < 0$ are spiteful. These are players who prefer to maximize the differences between their own monetary payoff and the monetary payoff of other players. Thus, this type of utility function is also consistent with the existence of people who reject positive offers in the ultimatum game.

However, it was soon observed that this type of utility function is not consistent with the quantitative details of ultimatum-game experiments. Indeed, from the rate of rejections observed in an experiment, one can easily compute the distribution of the spitefulness parameter. Since proposers are drawn from the same population, one can then use this distribution to compute what offer should be made by proposers. The offers that the proposers should make, according to the α calculated, are substan-

tially larger than those observed in the experiment (Levine 1998).

Starting from this observation, Levine (1998) proposed a generalization of Ledyard's utility function that assumed that agents have information (or beliefs) about the level of altruism of the other players and base their own level of altruism on that of the other players. This allows us to formalize the intuition that players may be more altruistic toward altruistic players than toward selfish or spiteful players. Specifically, Levine proposed the utility function

$$u_i(x_1, \dots, x_n) = x_i + \sum_{j \neq i} \frac{\alpha_i + \lambda \alpha_j}{1 + \lambda} x_j,$$

where $\lambda \in [0, 1]$ is a parameter representing how sensitive players are to the level of altruism of the other players. If $\lambda = 0$, then we obtain Ledyard's model, where i 's level of altruism toward j does not depend on j 's level of altruism toward i ; if $\lambda > 0$, players tend to be more altruistic toward altruistic players than toward selfish and spiteful players. Levine showed that this model fits the empirical data in several settings quite well, including the ultimatum game and prisoner's dilemma.

One year later, Fehr and Schmidt (1999) introduced a utility function based on a somewhat different idea. Instead of caring directly about the absolute payoffs of the other players, Fehr and Schmidt assumed that (some) people care about minimizing payoff differences. Following this intuition, they introduced a family of utility function that can capture inequity aversion:

$$\begin{aligned} u_i(x_1, \dots, x_n) = & x_i - \frac{\alpha_i}{n-1} \\ & \times \sum_{j \neq i} \max\{x_j - x_i, 0\} \\ & - \frac{\beta_i}{n-1} \sum_{j \neq i} \max\{x_i - x_j, 0\}. \end{aligned}$$

Note that the first sum is greater than zero if and only if some player j receives more than player i ($x_j > x_i$). Therefore, α_i can be interpreted as a parameter representing the extent to which player i dislikes disadvantageous inequalities. Similarly, β_i can be interpreted as a parameter representing the extent to which player i dislikes advantageous inequities. Fehr and Schmidt also assumed that $\beta_i \leq \alpha_i$ and $\beta_i \in [0, 1)$. The first assumption means that players dislike having a payoff higher than that of some other player at most as much as they dislike having a payoff lower than that of some other player. To understand the assumption that $\beta_i < 1$, suppose that player i has a payoff larger than the payoff of all the other players. Then i 's utility function reduces to

$$u_i(x_1, \dots, x_n) = (1 - \beta_i)x_i + \frac{\beta_i}{n-1} \sum_{j \neq i} x_j.$$

If $\beta_i \geq 1$, then the component of the utility function corresponding to the monetary payoff of player i is nonpositive, so player i maximizes utility by giving away all his money, an assumption that seems implausible (Fehr and Schmidt 1999). Finally, the assumption $\beta_i \geq 0$ simply means that there are no players who prefer to be better off than other players, even at a cost to themselves.⁷ This way of capturing inequity aversion has been shown to fit empirical data well in a number of contexts, including the standard ultimatum game, variants with multiple proposers and with multiple responders, and the public-goods game.

A different type of utility function capturing inequity aversion was introduced by Bolton and Ockenfels (2000). Like Fehr and Schmidt (1999), Bolton and Ockenfels assume that players' utility functions take into account inequities among players. To define

this function, they assume that the monetary payoffs of all players are nonnegative. Then they define

$$\sigma_i = \begin{cases} \frac{x_i}{c}, & \text{if } c > 0; \\ \frac{1}{n}, & \text{if } c = 0; \end{cases}$$

where $c = \sum_{j=1}^n x_j$, so σ_i represents i 's relative share of the total monetary payoff. Bolton and Ockenfels further assume that i 's utility function (which they refer to as i 's *motivational function*) depends only on i 's monetary payoff x_i and on his relative share σ_i , and satisfies four assumptions. We refer to Bolton and Ockenfels (2000) for the formal details; here, we focus on an intuitive description of two of these assumptions, the ones that characterize inequity aversion (the other two assumptions are made for mathematical convenience). One assumption is that, keeping the relative payoff σ_i constant, i 's utility is increasing in x_i . Thus, for two choices that give the same relative share, player i 's decision is consistent with payoff maximization. The second assumption is that, holding x_i constant, i 's utility is strictly concave in σ_i , with a maximum at the point at which player i 's monetary payoff is equal to the average payoff. Thus, keeping monetary payoff constant, players prefer an equal distribution of monetary payoffs. This utility function was shown to fit empirical data quite well in a number of games, including the dictator game, ultimatum game, and prisoner's dilemma. (We defer a comparison of Bolton and Ockenfels's approach with that of Fehr and Schmidt.)

Shortly after the explosion of inequity-aversion models, several economists observed that some decision-makers appear to act in a way that *increases* inequity, if this increase results in an increase in the total payoff of the participants (Charness and Grosskopf 2001, Kritikos and Bolle 2001, Andreoni and Miller 2002, Charness and Rabin 2002). This observation is hard to reconcile with

⁷Fehr and Schmidt (1999) make this assumption for simplicity, although they acknowledge that they believe that there are subjects with $\beta_i < 0$.

inequity aversion models, and suggests that people not only prefer to minimize inequity, but also prefer to maximize social welfare.

To estimate these preferences, Andreoni and Miller (2002) conducted an experiment, found the utility function in a particular class of utility functions that best fit the experimental results, and showed that this utility function also fits data from other experiments well. In more detail, they conducted an experiment in which participants made decisions in a series of modified dictator games where the cost of giving is in the set $\{0.25, 0.5, 1, 2, 3\}$. For example, when the cost of giving is 0.25, sending one token to the recipient results in the recipient receiving four tokens. Andreoni and Miller found that 22.7 percent of the dictators were perfectly selfish (so their behavior could be rationalized by the utility function $u(x_1, x_2) = x_1$), 14.2 percent of dictators split the monetary payoff equally with the recipient (so their behavior could be rationalized by the *Rawlsian* utility function $u(x_1, x_2) = \min\{x_1, x_2\}$),⁸ and 6.2 percent of the dictators gave to the recipient only when the price of giving was smaller than 1 (and thus their behavior could be rationalized by the *utilitarian* utility function $u(x_1, x_2) = (1/2)x_1 + (1/2)x_2$). To rationalize the behavior of the remaining 57 percent of the dictators, Andreoni and Miller fit their data to a utility function of the form

$$u_1(x_1, x_2) = [\alpha x_1^\rho + (1 - \alpha)x_2^\rho]^{1/\rho}.$$

Here, α represents the extent to which the dictator (player 1) cares about his own monetary payoff relatively to that of the recipient (player 2), while ρ represents the convexity of preferences. Andreoni and

Miller found that subjects can be divided into three classes that they called *weakly selfish* ($\alpha = 0.758, \rho = 0.621$; note that if $\alpha = 1$, we get self-regarding preferences), *weakly Rawlsian* ($\alpha = 0.654, \rho = -0.350$; note that if $0 < \alpha < 1$ and $x_1, x_2 > 0$, then as $\rho \rightarrow -\infty$, we converge to Rawlsian preferences), and *weakly utilitarian* ($\alpha = 0.576, \rho = 0.669$; note that if $\alpha = 0.5$ and $\rho = 1$, we get utilitarian preferences). Moreover, they showed that the model also fits well experimental results on the standard dictator game, the public-goods game, and prisoner's dilemma. Finally, this model can also explain the results mentioned above showing that some decision-makers act so as to increase inequity, if the increase leads to an increase in social welfare.

Charness and Rabin (2002) considered a more general family of utility functions that includes the ones mentioned earlier as special cases. Like Andreoni and Miller, they conducted an experiment to estimate the parameters of the utility function that best fit the data. For simplicity, we describe Charness and Rabin's utility function in the case of two players; we refer to their paper for the general formulation. They considered a utility function for player 2 of the form

$$u_2(x_1, x_2) = (\rho r + \sigma s)x_1 + (1 - \rho r - \sigma s)x_2,$$

where: (1) $r = 1$ if $x_2 > x_1$, and $r = 0$ otherwise; (2) $s = 1$ if $x_2 < x_1$, and $s = 0$ otherwise.⁹ Intuitively, ρ represents how important it is to agent 2 that she gets a higher payoff than agent 1, while σ represents how important it is to agent 2 that agent 1 gets a higher payoff than she does.

⁸Named after John Rawls, a philosopher, who argued, roughly speaking, that in a just society, the social system should be designed so as to maximize the payoff of those worst off.

⁹The general form of the utility function has a third component that takes into account reciprocity in sequential games. Since in this review we focus on normal-form games, we ignore this component here.

Charness and Rabin did not make any a priori assumptions regarding ρ and σ , but they showed that by setting ρ and σ appropriately, one can recover the earlier models:

- Assume that $\sigma \leq \rho \leq 0$. In this case, player 2's utility is increasing in $x_2 - x_1$. So, by definition, this case corresponds to competitive preferences.
- Assume that $\sigma < 0 < 1/2 < \rho < 1$. If $x_1 < x_2$, then player 2's utility is $\rho x_1 + (1 - \rho)x_2$, and thus depends positively on both x_1 and x_2 because $0 < \rho < 1$. Moreover, since $\rho > 1/2$, player 2 prefers increasing player 1's payoff to his own; that is, player 2 prefers to decrease inequity (since $x_2 > x_1$). If $x_2 < x_1$, then player 2's utility is $\sigma x_1 + (1 - \sigma)x_2$, and thus depends negatively on x_1 and positively on x_2 , so again player 2's prefers to decrease inequity.
- If $0 < \sigma \leq \rho \leq 1$, then player 2's utility depends positively on both x_1 and x_2 . Charness and Rabin define these as social welfare preferences. (Note that in the case $\sigma = \rho = 1/2$, one obtains utilitarian preferences $u_2(x_1, x_2) = (x_1 + x_2)/2$; more generally, Charness and Rabin apply the term “social welfare preferences” to those preferences where the individual payoffs are both weighted positively.)

To test which of these cases better fits the experimental data, Charness and Rabin conducted a series of dictator game experiments. They found that the standard assumption of narrow self-interest explains only 68 percent of the data, assuming competitive preferences (i.e., $\sigma \leq \rho \leq 0$) explains even less (60 percent), and assuming a preference for inequity aversion (i.e., $\sigma < 0 < 1/2 < \rho < 1$) is consistent with 75 percent of the data. The best results were obtained by assuming a pref-

erence for maximizing social welfare (i.e., $0 < \sigma \leq \rho \leq 1$), which explains 97 percent of the data.

Engelmann and Strobel (2004) also showed that assuming a preference for maximizing social welfare leads to better predictions than assuming a preference for inequity aversion. They considered a set of decision problems designed to compare the relative ability of several classes of utility functions to make predictions. They considered utility functions corresponding to payoff maximization, social welfare maximization, Rawlsian maximin preferences, the utility function proposed by Fehr and Schmidt, and the utility function proposed by Bolton and Ockenfels. They found that the best fit of the data was provided by a combination of social welfare concerns, maximin preferences, and selfishness. Moreover, Fehr and Schmidt's inequity aversion model outperformed that of Bolton and Ockenfels. However, this increase in performance was entirely driven by the fact that, in many cases, the predictions of Fehr and Schmidt reduced to maximin preferences.

Team-reasoning models (Gilbert 1987; Bacharach 1999; Sugden 2000) and equilibrium notions such as *cooperative equilibrium* (Halpern and Rong 2010; Capraro 2013) also take social welfare maximization seriously. The underlying idea of these approaches is that individuals do not always act so as to maximize their individual monetary payoff, but may also take into account social welfare. For example, in the prisoner's dilemma, social welfare is maximized by mutual cooperation, so team-reasoning models predict that people cooperate. However, since these models typically assume that the utility of a player is a function of the sum of the payoffs of all players, they cannot explain behaviors in zero-sum games, such as the dictator game. Another limitations of these approaches is their inability to explain the behavior of people who choose to mini-

mize both their individual payoffs and social welfare, such as responders who reject offers in the ultimatum game.

Cappelen et al. (2007) introduced a model in which participants strive to balance their individual “fairness ideal” with their self-interest. They consider three fairness ideals. *Strict egalitarianism* contends that people are not responsible for their effort and talent; according to this view, the fairest distribution of resources is the equal distribution. *Libertarianism* argues that people are responsible for their effort and talent; according to this view, resources should be shared so that each person’s share is in proportion to what she produces. *Liberal egalitarianism* is based on the belief that people are responsible for their effort, but not for their talent; according to this view, resources should be distributed so as to minimize differences due to talent, but not those due to effort. To formalize people’s tendency to implement their fairness ideal, Cappelen et al. considered utility functions of the form

$$u_i(x_i, a, q) = \gamma_i x_i - \frac{\beta_i}{2} [x_i - p_i(a, q)]^2,$$

where $a = (a_1, \dots, a_n)$ is the vector of talents of the players, $q = (q_1, \dots, q_n)$ is the vector of efforts made by the players, x_i is the monetary payoff of player i , and $p_i(a, q)$ is the monetary payoff that player i believes to be his fair share. Finally, γ_i and β_i are individual parameters, representing the extent to which player i cares about his monetary payoff and his fairness ideal. In order to estimate the distribution of the fairness ideals, Cappelen et al. conducted an experiment using a variant of the dictator game that includes a production phase. The quantity produced depends on factors within the participants’ control (effort) and factors beyond participants’ control (talent). The experiment showed that 39.7 percent of the subjects can be viewed as strict egalitarians, 43.4 percent

as liberal egalitarians, and 16.8 percent as libertarians. Although this approach is useful in cases where the initial endowments are earned by the players through a task involving effort and/or talent, when the endowments are received as windfall gains, as in most laboratory experiments, it reduces to inequity aversion, so it suffers from the same limitations that other utility functions based on inequity aversion do.

The frameworks discussed thus far are particularly suitable for studying situations in which there is only one decision-maker or in which the choices of different decision-makers are made simultaneously. In many cases, however, there is more than one decision-maker, and they make their choices sequentially. For nonsimultaneous-move games, scholars have long recognized that intentions play an important role. A particular class of intention-based models—reciprocity models—is based on the idea that people typically reciprocate (perceived) good actions with good actions and (perceived) bad actions with bad actions. Economists have considered several models of reciprocity (Rabin 1993, Levine 1998, Charness and Rabin 2002, Dufwenberg and Kirchsteiger 2004, Sobel 2005, Falk and Fischbacher 2006). Intention-based models have been shown to explain deviations from outcome-based predictions in a number of situations where beliefs about the other players’ intentions may play a role (Falk, Fehr, and Fischbacher 2003; McCabe, Rigdon, and Smith 2003; Fehr and Schmidt 2006; Falk, Fehr, and Fischbacher 2008; Dhaene and Bouckaert 2010). Although we acknowledge the existence and the importance of these models, as we mentioned in the introduction, in this review we focus on normal-form games. Some of these games (e.g., dictator game, die-under-cup paradigm, trade-off game) have only one decision-maker, and therefore beliefs about others’ intentions play no role. In this con-

text, beliefs can still play a role, for example beliefs about others' beliefs, since even in single-agent games there may be others watching (or the decision-maker can play as if there are). This leads us to psychological games, which will be mentioned in section 5.

Outcome-based social preferences explain many experimental results well. In particular, looking at the experimental regularities listed in section 2, they easily explain the first three regularities, at least qualitatively. However, they cannot explain any of the remaining regularities. Indeed, the main limitation of outcome-based preferences is that they depend only on the monetary payoffs.

As we suggested in section 2, one way of explaining these regularities is to assume that people have *moral preferences*. Crucially, these moral preferences cannot be defined solely in terms of the economic consequences of the available actions. We review such moral preferences in section 4.¹⁰

Before moving to moral preferences, it is worth noting that another limitation of outcome-based social preferences is that they predict correlations between different behaviors that are not consistent with experimental data. For example, Chapman et al. (2023) reported on a large experiment showing that eight standard measures of prosociality can actually be clustered in three principal components corresponding to altruistic behavior, punishment, and ineq-

uity aversion, which are virtually unrelated. Chapman et al. observed that this is not consistent with any outcome-based social preferences. However, we will see that this result is consistent with moral preferences.

4. Moral Preferences

In the previous sections, we showed that outcome-based preferences and, more generally, outcome-based decision rules, are inconsistent with experimental regularities 4–7. In this section, we review the empirical literature suggesting that all seven experimental regularities discussed in section 2 can be explained by assuming that people have preferences for following a norm¹¹ and discuss attempts to formalize this using an appropriate utility function. We use the term *moral preferences* as an umbrella term to denote this type of utility function.

Experimental regularities 1–4.—The fact that donations in the standard dictator game, offers and rejections in the ultimatum game, cooperation in social dilemmas, and honesty in lying tasks can be explained by moral preferences was independently shown by many authors.

Krupka and Weber (2013) asked dictators to report, for each available choice, how “socially appropriate” they think other dictators think that choice is; dictators were incentivized to guess the modal answer of other dictators. They found that an equal split was rated as the most socially appropriate choice. They also found that framing effects in the dictator game when passing from the “give”

¹⁰Of course, we can consider outcome-based preferences in the context of decision rules other than expected-utility maximization, such as *maximin* expected utility (Wald 1950, Gärdenfors and Sahlin 1982, Gilboa and Schmeidler 1989), *minimax regret* (Niehans 1948, Savage 1951), and *maximin* expected utility (Wald 1950, Gärdenfors and Sahlin 1982, Gilboa and Schmeidler 1989). For example, with maximin, an agent chooses the act that maximizes her worst-case (expected) utility. With minimax regret, an agent chooses the act that minimizes her worst-case regret (the gap between the payoff for an act a in state s and the payoff for the best act in state s). Although useful in many contexts, none of these approaches can explain experimental regularities 4–7, since they are all outcome based.

¹¹We wrote “a norm,” and not “the norm,” because there are different types of norms. For the aim of this review, it is important to distinguish between personal beliefs about what is right and wrong, beliefs about what others approve or disapprove of, and beliefs about what others actually do. We will get to this distinction in more detail later.

frame to the “take” frame can be explained by a change in the perception of what the most appropriate action is.

Kimbrough and Vostroknutov (2016) introduced a “rule-following” task to measure people’s *norm sensitivity*, specifically, how important it was to them to follow rules of behavior. In this task, each participant has to control a stick figure walking across the screen from left to right. Along its walk, the figure encounters five traffic lights, each of which turns red when the figure approaches it. Each participant has to decide how long to wait at the traffic light (which turns green after five seconds), knowing that she will lose a certain amount of money for each second spent on the waiting. The total amount of time spent waiting is taken as an individual measure of norm sensitivity. Kimbrough and Vostroknutov found that this parameter predicts giving in the dictator game and the public-goods game, and correlates with rejection thresholds in the ultimatum game. Bicchieri and Chavez (2010) showed that ultimatum-game responders reject the same offer at different rates, depending on the other available offers; in particular, responders tend to accept offers that they consider to be fairer, compared to the other available offers.

In a similar vein, Capraro and Rand (2018) and Capraro and Vanzo (2019) found that giving in the dictator game depends on what people perceive to be the morally right thing to do. Indirect evidence that moral preferences drive giving in the dictator game was also provided by research showing that including moral reminders in the instructions of the dictator game increases giving (Brañas-Garza 2007, Capraro et al. 2019). In addition, as mentioned in section 2, Eriksson et al. (2017) showed that framing effects among responders in the ultimatum game can be explained by a change in the perception of what is the morally right thing to do, while Dal Bó and Dal Bó (2014) found that

moral reminders increase cooperation in the iterated prisoner’s dilemma. Furthermore, Capraro (2018) found that, in the ultimatum game, 92 percent of the proposers and 72 percent of responders declare that offering half is the morally right thing to do, while Capraro and Rand (2018) found that 81 percent of the subjects declare that cooperating is the morally right thing to do in the one-shot prisoner’s dilemma.

Finally, the fact that honest behavior in economic games in which participants can lie for their benefit is partly driven by moral preferences was suggested by several authors (Gneezy 2005; Erat and Gneezy 2012; Fischbacher and Föllmi-Heusi 2013; Abeler, Nosenzo, and Raymond 2019). Empirical evidence was provided by Cappelen, Sørensen, and Tungodden (2013), who found that telling the truth in the sender–receiver game in the Pareto white-lie condition correlates positively with giving in the dictator game, suggesting that “aversion to lying not only is positively associated with pro-social preferences, but for many a stronger moral motive than the concern for the welfare of others”. Biziou-van-Pol et al. (2015) replicated the correlation between honesty in the Pareto white-lie condition and giving in the dictator game and additionally showed a similar correlation with cooperation in the prisoner’s dilemma; the authors suggested that cooperating, giving, and truth telling might be driven by a common motivation to do the right thing. Finally, Bott et al. (2019) found that moral reminders decrease tax evasion in a field experiment with Norwegian taxpayers.

Experimental regularities 5–6.—List (2007), Bardsley (2008), and Cappelen et al. (2013) showed that people tend to be more altruistic in the standard dictator game than in the dictator game with a take option (experimental regularity 5). The fact that this behavioral change might reflect moral preferences was suggested by Krupka

and Weber (2013). They found that sharing nothing in the standard dictator game is considered to be far less socially appropriate than sharing nothing in the dictator game with a take option. They also showed that social appropriateness can explain why some dictator-game givers prefer to avoid the interaction altogether, given the chance to do so (experimental regularity 6): dictators rate exiting the game to be far less socially inappropriate than keeping the money in the standard dictator game.

Experimental regularity 7.—In section 2, we reviewed the literature showing that behavior in several games, including the dictator game, the prisoner’s dilemma, the ultimatum game, and the trade-off game depends on the language used to present the instructions, especially when it activates moral concerns.

To summarize, all seven regularities can be qualitatively explained by assuming that people have moral preferences. In what follows, we review the models of moral preferences that have been introduced thus far in the literature. The idea that morality has to be incorporated in economic models has been around since the foundational work of Adam Smith and Francis Y. Edgeworth (Smith 2010, Edgeworth 1881); see Sen (1977), Binmore (1994), Tabellini (2008), Bicchieri (2006), and Enke (2019) for more recent accounts. However, work on utility functions that take moral preferences into account is relatively recent. In this review, we focus on utility functions that can be applied to all or most¹² of the economic interactions that are

described in terms of normal-form games with monetary payoffs.¹³

We proceed chronologically. We start by discussing the work of Akerlof and Kranton (2000). Their motivation was to study “gender discrimination in the workplace, the economics of poverty and social exclusion, and the household division of labor.” To do so, they proposed a utility function that takes into account a person’s *identity*. The identity is assumed to carry information about how a person should behave, which, in turn, is assumed to depend on the social categories to which a person belongs. In this setting, Akerlof and Kranton considered a utility function of the form

$$u_i = u_i(a_i, a_{-i}, I_i),$$

where I_i represents i ’s identity. They showed that their model can qualitatively explain group differences such as the ones that motivated their work. This model is certainly consistent with experimental regularities 1–7. Indeed, it suffices to assume that the identity takes into account a tendency to follow the norms. This model is conceptually similar to a previous model proposed by Stigler and Becker (1977), which is based on the idea that preferences should not be defined over marketed goods, but over general commodities that people transform into consumption goods. Although Stigler and Becker (1977) do not aim to explain the experimental regularities that are the focus

their contexts, these models cannot be readily extended to other types of interaction.

¹³Economists have also studied models of morality in other games (see, e.g., Bénabou and Tirole 2011). Recently, economists have also sought to explain political behavior in terms of moral preferences (Bonomi, Gennaioli, and Tabellini 2021; Enke, Polborn, and Wu 2022). Although these models cannot be readily applied to the economic games that are the focus of this review, they show that the idea that moral preferences can help explain people’s behavior is gaining traction across different areas of research.

¹²Economists have also introduced a number of domain-specific models; for example, models to explain cooperation in the prisoner’s dilemma (Bolle and Ockenfels 1990), honesty in lying tasks (Abeler, Nosenzo, and Raymond 2019; Gneezy, Kajackaite, and Sobel 2018), fairness in principal–agent models (Ellingsen and Johannesson 2008), and honesty in principal–agent models (Alger and Renault 2006, 2007). Although useful in

of this review, their model is consistent with them: for example, people may cooperate in order to maintain good relations, or may act altruistically to experience the warm glow of giving, or may act morally to adhere to their self-image. We refer to Sobel (2005) for a more detailed discussion and for the mathematical equivalence between this model and that of Akerlof and Kranton.

A more specific model, but one based on a similar idea, was introduced by Brekke, Kverndokk, and Nyborg (2003). Their initial aim was to explain field experiments showing that paying people to provide a public good can “crowd out” intrinsic motivations. For example, paying people reduces blood donations (Titmuss 2018), their willingness to accept a nuclear waste repository in their neighborhood (Frey and Oberholzer-Gee 1997), and volunteering (Gneezy and Rustichini 2000). To explain these findings, Brekke, Kverndokk, and Nyborg considered economic interactions in which each player i ($i = 1, \dots, n$) has to put in some effort e_i , measured in units of time, to generate a public good g_i . At most T units of time are assumed to be available, so that i has to decide how much time e_i to contribute to the public good and how much time l_i to use for leisure: $l_i = T - e_i$. The total quantity of public good is $G = G_p + \sum_{i=1}^n g_i$, where G_p is the public provision of the public good. The monetary payoff that player i receives from putting in effort e_i is denoted x_i . The key assumption of the model is that player i has a morally ideal effort, denoted e_i^* . Brekke, Kverndokk, and Nyborg postulated that player i maximizes a utility function of the form

$$u_i = u_i(x_i, l_i) + v_i(G) + f_i(e_i, e_i^*),$$

where u_i and v_i are increasing and concave, while the function $f(e_i, e_i^*)$ is assumed to attain its maximum at $e_i = e_i^*$; we can think of f as taking into account the distance between i 's

actual effort e_i and the ideal effort e_i^* in an inversely related way. As an explicit example, Brekke, Kverndokk, and Nyborg considered the function $f_i(e_i, e_i^*) = -a(e_i - e_i^*)^2$, with $a > 0$. Therefore, *ceteris paribus*, players aim to maximize their monetary payoff, their leisure time, and the public good, while aiming at minimizing the distance from their moral ideal. Brekke, Kverndokk, and Nyborg supposed that, before deciding their action, players consider their morally ideal effort. They assumed that all players share a utilitarian moral philosophy, so that the morally ideal effort is found by maximizing $W = \sum_i u_i$ with respect to e_i . Under these assumptions, they showed that their model is consistent with the crowding-out effect. Specifically, they showed that when a fee is introduced for people who do not contribute to the public good, if this fee is at least equal to the cost of buying g_i units of public good in the market and if this fee is smaller than the utility corresponding to the gain of leisure time due to not contributing, then the moral ideal e_i^* is equal to 0; in other words, the fee becomes a moral justification for not contributing. This intuitively happens because individuals leave the responsibility of ensuring the public good to the organization: the public good is provided by the organization, which buys it using the fees. This is convenient for the individuals, as they gain in leisure time. If we replace time with money, then this utility function can capture the empirical regularities observed in the dictator game and social dilemmas. However, the utility function cannot easily be applied to other settings, such as the ultimatum and trade-off games.

A more general utility function was introduced by Bénabou and Tirole (2006). It tries to take into account altruism and is motivated by a theory of social signaling, according to which people's actions are associated with reputational costs and benefits (Nowak and Sigmund 1998; Smith

and Bird 2000; Gintis, Smith, and Bowles 2001). Players choose a participation level for some prosocial activity from an action set $A \subseteq \mathbb{R}$. Choosing a has a cost $c(a)$ and gives a monetary payoff ya ; y can be positive, negative, or zero. Players are assumed to be characterized by a type $v = (v_a, v_y) \in \mathbb{R}^2$, where v_a is, roughly speaking, the impact of the prosocial factors associated with participation level a on the agent's utility, while v_y is, roughly speaking, the impact of money on the agent's utility. Bénabou and Tirole mentioned that v_a , the utility of choosing a , is determined by at least two factors: the material payoff of the other player and the enjoyment derived from the act of giving. Thus, their approach can capture Andreoni's (1990) notion of warm-glow giving.

Bénabou and Tirole then defined the *direct benefit* of action a to be

$$D(a) = (v_a + v_y y)a - c(a)$$

and the *reputational benefit* to be

$$R(a) = x(\gamma_a E[v_a | a, y] - \gamma_y E[v_y | a, y]),$$

where $E[v | a, y]$ represents the observers' posterior expectation that the player is of type v , given that the player chose a when the monetary incentive to choose a is y . The parameters γ_a and γ_y are assumed to be nonnegative. To understand this assumption, note that by definition, players with high v_a are prosocial; players with high v_y are greedy. Therefore, the hypothesis that γ_a and γ_y are nonnegative formalizes the idea that people like to be perceived as prosocial ($\gamma_a \geq 0$) and not greedy ($\gamma_y \geq 0$). The parameter $x > 0$ represents the visibility of an action, that is, the probability that the action is visible to other players. Bénabou and Tirole then defined the utility function

$$u(a) = D(a) + R(a).$$

They studied this utility function in a number of contexts in which actions can be observed and decision options can be defined in terms of contribution. While useful in its domains of applicability, this utility function cannot be applied to games where choices cannot be described in terms of participation levels, such as trade-off games in which players have the role of distributing money, without themselves being affected by the distribution levels.

A utility function that captures moral motivations and can be applied to all games was introduced by Levitt and List (2007). They argued that the utility of player i when she chooses action a depends on two factors. The first factor is the utility corresponding to the monetary payoff associated with action a . It is assumed to be increasing in the monetary value of a , denoted x_i . The second factor is the moral cost or benefit m_i associated with a . Levitt and List focused on three factors that can affect the moral value of an action. The first is the negative externality that a imposes on other people. They hypothesized that the externality is an increasing function of x_i : the more a player receives from choosing a , the less other participants receive. The second factor is the set n of moral norms and rules that govern behavior in the society in which the decision-maker lives. For example, the very fact that an action is illegal may impose an additional cost for that behavior. The third factor is the extent to which actions are observed. For example, if an illegal or an immoral action is recorded, or performed in front of the experimenter, it is likely that the decision-maker pays a greater moral cost than if the same action is performed when no one is watching. The effect of scrutiny is denoted by s ; greater scrutiny is assumed to increase the moral cost. Levitt and List added this component to take into account the fact that, when the behavior of participants cannot be monitored, people tend to be less prosocial than when it can (Bandiera, Barankay,

and Rasul 2005; List 2006; Benz and Meier 2008). With this notation, Levitt and List proposed that player i maximizes the utility function

$$u_i(a, x_i, n, s) = m_i(a, x_i, n, s) + w_i(a, x_i).$$

This approach can explain experimental regularities 1, 3, 5, and 7. However, the situation described in experimental regularity 2 and some instances of experimental regularities 4 and 6 violate Levitt and List's assumptions. Specifically, the assumption that the negative externalities associated with player i 's action depend negatively on i 's monetary payoff does not hold in the ultimatum game, where rejecting a low offer (the choice that is typically viewed as the moral choice (Bicchieri and Chavez 2010)) decreases both players' monetary payoffs. This is also the case in the sender–receiver game in the Pareto white-lie condition, where telling the truth is viewed as the moral choice, yet minimizes the monetary payoffs of both players. But these are minor limitations; they can be addressed by considering a slightly more general utility function that depends not only on x_i , but also on $\sum_{j \neq i} x_j$ (and dropping the assumption that these two variables are inversely related).

A similar approach was used by López-Pérez (2008), although he focused on extensive-form games. He assumed the existence of a “norm” function Ψ_i that associates, with each information set h for player i in a given game, an action $\Psi_i(h)$ that can be taken at h . Intuitively, $\Psi_i(h)$ is the moral choice at information set h . López-Pérez assumed that player i receives a psychological disutility (in the form of a negative emotion, such as guilt or shame) whenever she violates the norm expressed by Ψ_i . He did not propose a specific functional expression for the utility function; rather, he studied a particular example of a norm, the *e-norm*. This norm is conceptually similar to the

one used by Charness and Rabin (2002), described in section 3. López-Pérez showed that a utility function that gives higher utility to strategies that perform more moral actions qualitatively fits the empirical data in several contexts. Although he did not explicitly show that the experimental regularities presented in section 2 can be explained using his model, it is easy to show that this is the case. Indeed, it suffices to assume that the mapping Ψ_i just associates with the game the morally right action for player i . (We can view a normal-form game as having a single information set, so we can view Ψ_i as just applying to the whole game.) Note, however, that the norm Ψ_i defined in this way is different from the *e-norm* considered by López-Pérez, which is outcome-based.

Andreoni and Bernheim (2009) focused on the dictator game and introduced a utility function that combines elements from theories of social image with inequity aversion. The fact that some people care about how others see them has been recognized by economists for at least two decades (Bernheim 1994, Glazer and Konrad 1996). Combining these ideas with inequity aversion, Andreoni and Bernheim proposed that dictators maximize the utility function

$$u_i(x_i, m, t) = f(x_i, m) + tg\left(x_i - \frac{1}{2}\right),$$

where $x_i \in [0, 1]$ is the monetary payoff of the dictator (the endowment is normalized to 1), $m \geq 0$ is the social image of the dictator, f represents the utility associated with x_i (which is assumed to be increasing in both x_i and m and concave in x_i), $t \geq 0$ is a parameter representing the extent to which the dictator cares about minimizing inequity, and g is a (twice continuously differentiable, strictly concave) function that attains its maximum at 0. This later hypothesis formalizes the intuition that the dictator gets a disutility from not implementing the equal distribution. Andreoni and Bernheim assumed

that there is an audience A that includes the recipient and possibly other people (e.g., the experimenter). The audience observes the donation $1 - x_i$ and then forms beliefs about the dictator's level of fairness t . They assumed that the dictator's social image is some function B of Φ , the cumulative distribution Φ representing A 's beliefs about the dictator's level of fairness. For example, the function B could be the mean of t given Φ . Andreoni and Bernheim showed that with some minimal assumptions about Φ , we can explain donations in the dictator game quite well. In particular, their utility function is consistent with the prevalence of exactly equal splits, that is, the lack of donations slightly above or slightly below 50 percent of the endowment. They also considered a variant of the dictator game in which with some probability p the dictator's choice is not implemented, but instead the recipient receives an amount x_0 close to 0. Intuitively, this should have the effect of creating a peak of voluntary donations at x_0 , since dictators can excuse the outcome as being beyond their control, thus preserving their social image. Andreoni and Bernheim observed this behavior, and showed that this is indeed consistent with their utility function. Unfortunately, it is not clear how to extend this utility function beyond dictator-like games.

A conceptually similar approach was considered by DellaVigna, List, and Malmendier (2012). Instead of formalizing a concern for social image, their utility function takes into account the effect of social pressure (which might affect people's decisions through social image). This approach was motivated by a door-to-door fundraising campaign with three treatments: a *flier* treatment, in which households were informed one day in advance by a flier on their doorknob of the upcoming visit of someone soliciting donations; a *flier with an opt-out checkbox* treatment, where the flier contained a box to be checked in case the household did not want to

be disturbed; and a *baseline*, in which households were not informed about the upcoming visit. Della Vigna, List, and Malmendier found that the flier decreased the frequency of the door being opened, compared to the baseline. Moreover, the flier with an opt-out check box also decreased giving, but the effect was significant only for small donations. To explain these findings, they considered a two-stage game between a *prospect* (potential donor) and a *solicitor*. In the first stage, the prospect may receive a flier and, if so, she notices the flier with probability $r \in (0, 1]$. In the second stage, the solicitor visits the home. The probability of the prospect opening the door is denoted by h : if the prospect did not notice the flier, h is equal to the baseline probability h_0 ; otherwise, the prospect can update this probability at a cost $c(h)$, with $c(h_0) = 0$, $c'(h_0) = 0$, $c'' > 0$. That is, not updating the probability of opening the door has no cost; updating it has a cost that depends monotonically on the adjustment. A donor can donate either in person or through other channels (e.g., by email). Let g be the amount donated in person and g_m the amount donated through other means. A donor's utility is taken to be

$$\begin{aligned}
 u(g, g_m) &= f(w - g - g_m) \\
 &+ \alpha v(g + \delta g_m, g - i) \\
 &+ \sigma(g^\sigma - g) \mathbf{1}\{g < g^\sigma\},
 \end{aligned}$$

where w represents the initial wealth of the donor; $f(w - g - g_m)$ represents the utility of private consumption; α is a parameter representing the extent to which the donor cares about the payoff of the charity, which can be negative;¹⁴ δ is the proportion of the donation made through other channels

¹⁴In fact, if the social pressure (third addend of the utility function) is high enough, someone can end up donating even though she dislikes the charity.

that does not reach the intended recipient (e.g., the cost of an envelope and a stamp); g_{-i} is the (expected) donation made by other donors; σ is a parameter representing the extent to which the donor cares about social pressure; and g^σ is a trade-off donation. If the donor donates less than g^σ when the solicitor is present, then the donor pays a cost depending on $g^\sigma - g$. Della Vigna, List, and Malmendier showed that this utility function captures their experimental results well. However, it is not clear how to extend it to qualitatively different decision contexts.

Kessler and Leider (2012) considered a utility function in which players receive a disutility when they deviate from a norm; this leads to a model similar to that of Brekke, Kverndokk, and Nyborg (2003) that we described above. The main difference is that, instead of considering time, Kessler and Leider (2012) considered money. Players are assumed to have a norm \hat{x} that represents the ideal monetary contribution. The utility of player i is defined as

$$u_i(x_i, x_j, \hat{x}) = \pi_i(x_i, x_j) - \phi_i g(\hat{x} - x_i) \mathbf{1}\{x_i < \hat{x}\},$$

where $\pi_i(x_i, x_j)$ is the monetary payoff of player i when i contributes x_i and j contributes x_j , ϕ_i is a parameter representing i 's norm sensitivity, and g represents the disutility that player i gets from deviating from the norm, so $g(x_i) = 0$ if $x_i \geq \hat{x}$ and $g(x_i)$ increases with $x_i - \hat{x}$ if $x_i < \hat{x}$. Kessler and Leider applied their utility function to an experiment involving four games: an additive two-player public-goods game, a multiplicative public-goods game, a double-dictator game, and a Bertrand game. The set of contributions available depended on the game; they were chosen to ensure a mismatch between the individual monetary payoff-maximizing action and the socially beneficial action. Participants played ten rounds of these

games, with random rematching. Some of these rounds were preceded by a contracting phase in which participants could agree on which contribution to choose. Kessler and Leider observed that the presence of a contracting phase increased contributions. They argued that their utility function fits their data well; in particular, contracting increased contributions by increasing the norm.

In subsequent work, Kimbrough and Vostroknutov (2016) used this approach to explain prosocial behavior in the public-goods game, the dictator game, and the ultimatum game. The key innovation of this work was the estimation of the parameter ϕ_i , which was done using the "rule-following task" discussed earlier in this section. They found that their measure of norm sensitivity significantly correlates with cooperation in the public-goods game, with giving in the dictator game, and with rejection thresholds in the ultimatum game (but not with offers in the ultimatum game, although results trend in the expected direction). In sum, this utility function is very useful in its domain of applicability. Moreover, although it might seem difficult to extend it to situations in which strategies (and norms) cannot be expressed in terms of contributions, it can easily be extended to games where the space of strategy profiles is a metric space (X, d) where d is the metric (i.e., a distance function between strategies). In this setting, we can replace $\hat{x} - x$ in the utility function with $d(\hat{x}, x)$. It is easy to see that this utility function can explain all seven experimental regularities, assuming that \hat{x} coincides with what people view as the moral choice.

Lazear, Malmendier, and Weber (2012) considered a utility function motivated by their empirical finding that some sharers in the dictator game prefer to avoid the dictator-game interaction, given the possibility of doing so. They considered situations where players can choose one of two scenarios. In the scenario with a shar-

ing option, the agent plays the standard dictator game in the role of the dictator: she receives an endowment e , which she has to split between herself (x_1) and the recipient (x_2). In the scenario without a sharing option, the player simply receives $x_1 = e$, while the recipient receives $x_2 = 0$. They proposed a utility that depends on the scenario: $u = u(D, x_1, x_2)$, where $D = 1$ represents the scenario with a sharing opportunity, while $D = 0$ represents the scenario without the sharing opportunity. The fact that the utility depends on the scenario makes it possible to classify people as one of three types. The first type, the “willing sharers,” consists of individuals who prefer the sharing scenario and to share their endowment (at least, to some extent). Formally, this type of player is characterized by the conditions

$$\begin{cases} \max_{x_1 \in [0, e]} u(1, x_1, e - x_1) > u(0, e, 0); \\ \arg \max_{x_1 \in [0, e]} u(1, x_1, e - x_1) < e. \end{cases}$$

The second type, the “non-sharers,” never share. They are defined by the condition

$$\arg \max_{x_1 \in [0, e]} u(1, x_1, e - x_1) = e.$$

The third type, called “reluctant sharers,” are perhaps most interesting. They are determined by the remaining conditions:

$$\begin{cases} \max_{x_1 \in [0, e]} u(1, x_1, e - x_1) < u(0, e, 0); \\ \arg \max_{x_1 \in [0, e]} u(1, x_1, e - x_1) < e. \end{cases}$$

The first condition says that these players prefer to avoid the sharing opportunity when given the possibility of doing so. However, if they are forced to play the dictator game, they share part of their endowment. Lazear, Malmendier, and Weber showed that there is a significant number of people of the third type. Indeed, some subjects even prefer to pay a cost to avoid the sharing opportunity. While this gives a great deal of insight, it is

hard to generalize this type of utility function to other kinds of interaction.

Krupka and Weber (2013) introduced a model in which subjects are torn between following their self-interest and following the “injunctive norm,” that is, what they believe other people would approve or disapprove of (Cialdini, Reno, and Kallgren 1990). Let $A = \{a_1, \dots, a_k\}$ be the set of actions available. Krupka and Weber assumed the existence of a social norm function that associates to each action $a_j \in A$ with a number $N(a_j) \in \mathbb{R}$ representing the degree of social appropriateness of a_j . N is hypothesized to be independent of the individual; it represents the extent to which society views a_j as socially appropriate. $N(a_j)$ can be negative, in which case a_j is viewed as socially inappropriate. The utility of player i is defined as

$$u_i(a_j) = v_i[\pi_i(a_j)] + \gamma_i N(a_j),$$

where $\pi_i(a_j)$ is the monetary payoff of player i associated with action a_j , v_i is the utility associated to monetary payoffs, and γ_i is the extent to which i cares about doing what is socially appropriate. As we mentioned earlier in this section, one of Krupka and Weber’s contributions was to introduce a method of measuring social appropriateness. People were shown the available actions and, for each of them, asked to rate how socially appropriate they were. Participants were also incentivized to guess the mode of the choices made by the other participants. It is not difficult to show that Krupka and Weber’s utility function is consistent with all the experimental regularities, at least if we assume that the most socially appropriate choice coincides with what people believe to be the morally right thing to do. This suggests that one possible limitation of Krupka and Weber’s approach is that it takes into account only the injunctive norm. In general, there are situations where what people believe others would approve of (the injunctive norm) is different from the

choice they believe to be morally right (the personal norm). For example, suppose that a vegan must decide whether to buy a steak for \$5 or a vegan meal for \$10. Knowing that the vast majority of the population eats meat, the vegan might believe that others would view the steak as the most socially appropriate choice. However, since the vegan thinks that eating meat is morally wrong, she would opt for buying the vegan meal. This suggests that there might be situations where social appropriateness might not be a good predictor of human behavior. We return to this issue in section 6.

Alger and Weibull (2013) introduced a notion of *Homo moralis*. They considered a symmetric two-player game where the players have a common strategy set A .¹⁵ Let π_i be i 's payoff function; since the game is symmetric, we have that $\pi_i(a_i, a_j) = \pi_j(a_j, a_i)$. Players may differ in how much they care about morality. Morality is defined by taking inspiration from Immanuel Kant's *categorical imperative*: one should make the choice that is universalizable, that is, the action that would be best if everyone took it (e.g., cooperating in the prisoner's dilemma).¹⁶ More specifically, they defined player i to be a *Homo moralis* if his utility function has the form

$$u_i(a_1, a_2) = (1 - k)\pi_i(a_1, a_2) + k\pi_i(a_1, a_1),$$

where $k \in [0, 1]$ represents the degree of morality. If $k = 0$, then we recover the standard *Homo economicus*; if $k = 1$, we have *Homo kantiansis*, someone who makes the

universalizable choice. Alger and Weibull proved that evolution results in the degree of morality k being equal to the *index of assortativity* of the matching process. We refer to their paper for the exact definition of the index of assortativity. For our purposes, all that is relevant is that it is a nonnegative number that takes into account the potential nonrandomness of the matching process: it is zero with random matching, and greater than zero for nonrandom matching processes. Thus, in the particular case of random matching, Alger and Weibull's theorem shows that evolution results in *homo moralis* with degree of morality $k = 0$, namely, *homo economicus*. If the matching is not random, assortativity can favor the emergence of *homo moralis* with a degree of morality strictly greater than zero.

In subsequent work, Alger and Weibull (2016) showed that *Homo moralis* preferences are evolutionary stable, according to an appropriate definition of evolutionary stability for preferences. This approach is certainly useful for understanding the evolution of morality. However, if the payoff function π is equal to the monetary payoff, then *Homo moralis* preferences are outcome-based, so cannot explain experimental regularities 4–7. On the other hand, if the payoff function is not equal to the monetary payoff (or otherwise outcome based), then how should it be defined? Despite these limitations, it is important to note that there are some framing effects that can be explained by the model of Alger and Weibull (2013). Suppose, for example, that market interactions are formed through a matching process with less assortativity than are business partnership interactions, or coauthorship interactions, then a lower degree of morality should be expected in market interactions.

Kimbrough and Vostroknutov (2020) proposed a utility function similar to the one proposed by Krupka and Weber (2013) that was discussed earlier. Specifically, they pro-

¹⁵Their results also apply to nonsymmetric games where players do not know a priori which role they will play.

¹⁶See Laffont (1975) for a macroeconomic application of this principle. Roemer (2010) defined a notion of *Kantian equilibrium* in public-goods type games; these are strategy profiles in which no player would prefer all other players to change their contribution levels by the same multiplicative factor.

posed that people are torn between maximizing their material payoff and (not) doing what they think society would (dis)approve of. This corresponds to the utility function

$$u_i(a) = v_i[\pi_i(a)] + \phi_i \eta(a),$$

where $v_i[\pi_i(a)]$ represents the utility from the monetary payoff $\pi_i(a)$ corresponding to action a , ϕ_i represents the extent to which player i cares about (not) doing what she thinks society would (dis)approve of, and $\eta(a)$ is a measure of the extent to which society approves or disapproves a . The key difference between this and Krupka and Weber's approach is in the definition of η , which corresponds to Krupka and Weber's N . While Krupka and Weber defined N empirically, by asking experimental subjects what they believe others would find socially (in)appropriate, Kimbrough and Vostroknutov defined η in terms of the cumulative dissatisfaction that players experience when a certain strategy profile is realized, rather than other possible strategy profiles. Kimbrough and Vostroknutov's notion of "dissatisfaction" corresponds to what is more typically called *regret*: the difference between what a player could have gotten and what she actually got. Thus, the cumulative dissatisfaction is the sum of the regrets of the players. In their model, Kimbrough and Vostroknutov assumed that the normatively best outcome is the one that minimizes aggregated dissatisfaction. This implies that this model predicts that Pareto-dominant strategy profiles are always more socially appropriate than Pareto-dominated strategy profiles. Therefore, while this model explains well some types of moral behaviors, such as cooperation, it fails to explain moral behaviors that are Pareto dominated, such as honesty when lying is Pareto optimal, and framing effects, in the trade-off game, that push people to choose the equal but Pareto-dominated allocation of money.

Capraro and Perc (2021) introduced a utility function which, instead of considering people's tendency to follow the injunctive norm, considers their tendency to follow their personal norms, that is, their internal standards about what is right or wrong in a given situation (Schwartz 1977). Specifically, Capraro and Perc proposed the utility function

$$u_i(a) = v_i[\pi_i(a)] + \mu_i P_i(a),$$

where $v_i[\pi_i(a)]$ represents the utility from the monetary payoff $\pi_i(a)$ corresponding to action a , μ_i represents the extent to which player i cares about doing what she thinks to be morally right, and $P_i(a)$ represents the extent to which player i thinks that a is morally right. From a practical perspective, the main difference between this utility function and those proposed by Krupka and Weber (2013) and by Kimbrough and Vostroknutov (2016) is that personal norms are individual, that is, $P_i(a)$ depends specifically on player i , while the injunctive norm depends only on the society where an individual lives, so is the same for all individuals in the same society. This utility function is also consistent with all seven regularities described in section 2.

Finally, Bašić and Verrina (2020) proposed a utility function that combines personal and injunctive norms:

$$u_i(a) = v_i[\pi_i(a)] + \gamma_i S(a) + \delta_i P_i(a),$$

where γ and δ represent the extent to which people care about following the injunctive and the personal norms, respectively. This utility function was suggested to the authors by experimental evidence in support of the fact that injunctive norms and personal norms are differentially associated with giving in the dictator game (standard version and a variant with a tax) and with behavior in the ultimatum game (as well as in a third-party punishment game). Also,

this utility function is consistent with all the seven experimental regularities that we are considering in this review.

At the end of section 3, we observed that another limitation of social preferences is that they predict correlations between different prosocial behaviors that are not observed in experimental data. We conclude this section by observing that the lack of these correlations is consistent with at least some moral preferences. Indeed, several models of moral preference (e.g., Akerlof and Kranton 2000, Levitt and List 2007, López-Pérez 2008, Krupka and Weber 2013, Kimbrough and Vostroknutov 2016, Capraro and Perc 2021, Bašić and Verrina 2020) do not assume that morality is unidimensional. If morality is multidimensional and different people may weigh different dimensions differently, then it is possible that, for some people, the right thing to do is to act altruistically, while for others it is to punish antisocial behavior, and for yet others, it is to minimize inequity; this would rationalize the experimental findings of Chapman et al. (2023). We will come back to the multidimensionality of morality in section 6.

5. Language-Based Preferences

In this section, we go beyond moral preferences and consider settings where an agent's utility depends not just on what happens, but on how the agent feels about what happens, which is largely captured by the agent's beliefs, and how what happens is described, which is captured by the agent's language.

Work on a formal model, called a *psychological game*, for taking an agent's beliefs into account in the utility function started with Geanakoplos, Pearce, and Stacchetti (1989), and was later extended by Battigalli and Dufwenberg (2009) to allow for dynamic beliefs, among other things. The overview by Battigalli and Dufwenberg (2022) shows how the fact that the utility function in psycho-

logical games can depend on beliefs allows us to capture, for example, guilt feelings (whether Alice leaves a tip for a taxi driver, Bob, in a foreign country might depend on how guilty Ann would feel if she didn't leave a tip, which in turn depends on what Alice believes Bob is expecting in the way of a tip), reciprocity (how kind player i is to j depends on i 's beliefs regarding whether j will reciprocate), other emotions (disappointment, frustration, anger, regret), image (your belief about how others perceive you), and expectations (how what you actually get compares to your expectation of what you will get; see also Kőszegi and Rabin 2006), among other things.

Having the agent's utility function depend on *language*, how the world is described, provides a yet more general way to express preferences. (It is more general since the language can include a description of the agent's beliefs; see below.) Experimental regularity 7 describes several cases in which people's behavior depends on the words used to describe the available actions. Language-based preferences allow us to explain experimental regularity 7, as well as all the other experimental regularities regarding social interactions listed in section 2, in a straightforward way. From this point of view, we can interpret language-based models as a generalization of moral preferences, where the utility of a sentence is assumed to carry the moral value of the action described by that sentence. Language-based models are strictly more general than moral preferences. In this section, we will show that they can explain other well-known regularities (e.g., ones not involving social interactions) that have been found in behavioral experiments, such as the Allais paradox.

A classic example is standard *framing effects*, where people's decisions depend on whether the alternatives are described in terms of gains or losses (Tversky and Kahneman 1985). It is well known, for

example, that presenting alternative medical treatments in terms of survival rates versus mortality rates can produce a marked difference in how those treatments are evaluated, even by experienced physicians (McNeil et al. 1982). A core insight of *prospect theory* (Kahneman and Tversky 1979) is that subjective value depends not (only) on facts about the world, but on how those facts are viewed (as gains or losses, dominated or undominated options, etc.). And how they are viewed often depends on how they are described in language. For example, Thaler (1980) observed that the credit card lobby preferred that the difference between the price charged to cash and credit card customers be presented as a discount for paying cash rather than as a surcharge for paying by credit card. The two different descriptions amount to taking different reference points.¹⁷

Such language dependence is ubiquitous. We celebrate tenth and one hundredth anniversaries specially, and make a big deal when the Dow Jones Industrial Average crosses a multiple of 1,000, all because we happen to work in a base-10 number system. Prices often end in .99, since people seem to perceive differently the difference between \$19.99 and \$20 and the difference between \$19.98 and \$19.99. We refer to Strulov-Shlain (2023) for some recent work on and an overview of this well-researched topic.

One important side effect of the use of language, to which we return below, is that it emphasizes categories and clustering. For example, as Krueger and Clement (1994) showed, when people were asked to estimate the average high and low temperatures in

Providence, Rhode Island, on various dates, while they were fairly accurate, their estimates were relatively constant for any given month and then jumped when the month changed; the difference in estimates for two equally spaced days was significantly higher if the dates were in different months than if they were in the same month. This clustering arises in likelihood estimation as well. We often assess likelihoods using words like “probable,” “unlikely,” or “negligible,” rather than numeric representations, and when numbers are used, we tend to round them (Manski and Molinari 2010).

The importance of language to economics was already stressed by Rubinstein (1998) in his book *Economics and Language*. In chapter 4, for example, he considers the impact of having an agent’s preferences be definable in a simple propositional language. There have been various formal models that take language into account. For example:

- Lipman (1999) considers “pieces of information” that an agent might receive, and takes the agent’s state space to be characterized by maximal pieces of information. He also applies his approach to framing problems, among other things.
- Although Ahn and Ergin (2010) do not consider language explicitly, they do allow for the possibility that there may be different descriptions of a particular event and use this possibility to capture framing. For them, a “description” is a partition of the state space.
- Blume, Easley and Halpern (2021) take an agent’s object of choice to be programs, where a program is either a primitive program or has the form **if t then a else b** , where a and b are themselves programs, and t is a test. A test is just a propositional formula, so language plays a significant role in the agent’s preferences. Blume, Easley, and

¹⁷Tversky and Kahneman emphasized the distinction between gains and losses in prospect theory, but they clearly understood that other features of a description were also relevant. However, note that prospect theory applied to monetary outcome results is yet another instance of outcome-based preferences, so cannot explain experimental regularities 4–7.

Halpern also show how framing effects can be captured in their approach.

- There are several approaches to decision-making that can be viewed as implicitly based on language. For example, a critical component of Gilboa and Schmeidler's (2001) *case-based* decision theory is the notion of a *similarity function*, which assesses how close a pair of problems are to each other. We can think of problems as descriptions of choice situations in some language.
- Jehiel's (2005) notion of *analogy-based expectation equilibrium* assumes that there is some way of partitioning situations in bundles that, roughly speaking, are treated the same way when it comes to deciding how to move in a game. Again, we can think of as these bundles as ones whose descriptions are similar.
- Finally, Mullainathan (2002) assumes that people use coarse categories (similar in spirit to Jehiel's analogy bundles, the categories partition the space of possibilities) to make predictions.

While none of these approaches directly models the language used, many of the examples they use are language-based.

- While not directly part of the utility function, the role of vagueness and ambiguity in language, how it affects communication, and its economic implications have been studied and modeled (see, e.g., Blume and Board 2014, Halpern and Kets 2015).

We focus here on *language-based games* (Bjorndahl, Halpern, and Pass 2013), where the utility function directly depends on the language. As we shall see, language-based games provide a way of formalizing all the examples above. The following example, which deals with surprise,

gives a sense of how language-based games work.

EXAMPLE 5.1. (Bjorndahl, Halpern, and Pass 2013) Alice and Bob have been dating for a while now, and Bob has decided that the time is right to pop the big question. Though he is not one for fancy proposals, he does want it to be a surprise. In fact, if Alice expects the proposal, Bob would prefer to postpone it entirely until such time as it might be a surprise. Otherwise, if Alice is not expecting it, Bob's preference is to take the opportunity.

We can summarize this scenario by the payoffs for Bob given in table 1.

TABLE 1
THE SURPRISE PROPOSAL

	p	$\neg p$
B_{Ap}	0	1
$\neg B_{Ap}$	1	0

In this table, we denote Bob's two strategies, proposing and not proposing, by p and $\neg p$, respectively, and use B_{Ap} ($\neg B_{Ap}$) to denote that Alice is expecting (not expecting) the proposal. (More precisely, B_{Ap} says that Alice believes that Bob will propose; we are capturing Alice's expectations by her beliefs.) Thus, although Bob is the only one who moves in this game, his utility depends not just on his moves, but on Alice's expectations.

This choice of language already illustrates one of the features of the language-based approach: coarseness. We used quite coarse language to describe Alice's expectation: she either expects the proposal or she doesn't. Since the expectation is modeled using belief, this example can be captured using a psychological game as well. Of course, whether or not Alice expects a proposal may be more

than a binary affair: she may, for example, consider a proposal unlikely, somewhat likely, very likely, or certain. In a psychological game, Alice’s beliefs would be expressed by placing an arbitrary probability $\alpha \in [0, 1]$ on p . But there is good reason to think that an accurate model of her expectations involves only a small number k of distinct “levels” of belief, rather than a continuum. Table 1, for simplicity, assumes that $k = 2$, though this is easily generalized to larger values.

Once we fix a language (which is just a finite or infinite set of formulas), we can take a *situation* to be a maximal consistent set of formulas; that is, a complete description of the world in that language.¹⁸ In the example above, there are four situations: $\{p, B_{Ap}\}$ (Bob proposes and Alice expects the proposal), $\{p, \neg B_{Ap}\}$ (Bob proposes but Alice is not expecting it), $\{\neg p, B_{Ap}\}$ (Bob does not propose although Alice is expecting him to), and $\{\neg p, \neg B_{Ap}\}$ (Bob does not propose and Alice is not expecting a proposal). An agent’s language describes all the features of the game that are relevant to the player. An agent’s utility function associates a utility with each situation, as in table 1 above. Standard game theory is the special case where, given a set Σ_i of strategies (moves) for each player i , the formulas have the form $play_i(\sigma_i)$ for $\sigma_i \in \Sigma_i$. The situations are then strategy profiles.

A normal-form psychological game can be viewed as a special case of a language-based game where (i) the language talks only about agents’ strategies and agents’ possibly higher-order beliefs about these strategies (e.g., Alice’s beliefs about Bob’s beliefs about Alice’s beliefs about the proposal), and (ii) those beliefs are described using probabilities. For example, taking α to denote Alice’s

probability of p , psychological game theory might take Bob’s utility function to be the following:

$$u_B(x, \alpha) = \begin{cases} 1 - \alpha, & \text{if } x = p; \\ \alpha, & \text{if } x = \neg p. \end{cases}$$

The function u_B agrees with table 1 at its extreme points if we identify B_{Ap} with $\alpha = 1$ and $\neg B_{Ap}$ with $\alpha = 0$. Otherwise, for the continuum of other values that α may take between 0 and 1, u_B yields a convex combination of the corresponding extreme points. Thus, in a sense, u_B is a continuous approximation to a scenario that is essentially discrete.

The language implicitly used in psychological games is rich in one sense—it allows a continuum of possible beliefs—but is poor in the sense that it talks only about belief. That said, as we mentioned above, many human emotions can be expressed naturally using beliefs, and thus studied in the context of psychological games. The following example illustrates how.

EXAMPLE 5.2. (Bjorndahl, Halpern, and Pass 2013) Alice and Bob play a classic prisoner’s dilemma game with one twist: neither wishes to live up to low expectations. Specifically, if Bob expects the worst of Alice (i.e., expects her to defect), then Alice, indignant at Bob’s opinion of her, prefers to cooperate. Likewise for Bob. On the other hand, in the absence of such low expectations from their opponent, each will revert to their classical preferences.

The standard prisoner’s dilemma is summarized in table 2:

TABLE 2
THE CLASSICAL PRISONER’S DILEMMA

	c	d
c	(3, 3)	(0, 5)
d	(5, 0)	(1, 1)

¹⁸What counts as a maximal consistent set of formulas depends on the semantics of the language. We omit the (quite standard) formal details here; they can be found in Bjorndahl, Halpern, and Pass (2013).

Let u_A, u_B denote the two players' utility functions according to this table. Let the language consist of the formulas of the form $play_i(\sigma), B_i[play_i(\sigma)]$, and their negations, where $i \in \{A, B\}$ and $\sigma \in \{c, d\}$. Given a situation S , let σ_S denote the unique strategy profile determined by S . We can now define a language-based game that captures the intuitions above by taking Alice's utility function u'_A on situations S to be

$$u'_A(S) = \begin{cases} u_A(\sigma_S) - 6, & \text{if } play_A(d), \\ & B_B play_A(d) \in S; \\ u_A(\sigma_S), & \text{otherwise;} \end{cases}$$

and similarly for u'_B .

More generally, we could take Alice's utility to be $u_A(\sigma_S) - 6\theta$ if $play_A(d), B_B play_A(d) \in S$, where θ is a measure of the extent to which Alice's indignance affects her utility. And yet more generally, if the language lets us talk about the full range of probabilities, Alice's utility can depend on the probability she ascribes to $play_A(d)$. (Although we have described the last variant using language-based games, it can be directly expressed using psychological games.)

Using language lets us go beyond expressing the belief dependence captured by psychological games. For one thing, the coarseness of the language lets us capture some well-known anomalies in the preferences of consumers. For example, we can formalize the explanation hinted at earlier for why prices often end in 0.99. Consider a language that consists of price ranges like "between \$55 and \$55.99" and "between \$60 and \$64.99." With such language, the agent is forced to ascribe the same utility to \$59.98 and \$59.99, while there can be a significant difference between the utilities of \$59.99 and \$60. Intuitively, we think of the agent as using two languages: the (typically quite rich) language used to describe the world and the (perhaps

much coarser) language over which utility is defined. Thus, while the agent understands the difference between a price of \$59.98 and \$59.99 perfectly well, her utility function may be insensitive to that difference.

Using coarse language effectively limits the set of describable outcomes, and thus makes it easier for a computationally bounded agent to determine her own utilities. These concerns suggest that there might be even more coarseness at higher ranges. For example, suppose that the language includes terms like "around \$20,000" and "around \$300." If we assume that both "around \$20,000" and "around \$300" describe intervals (centered at \$20,000 and \$300, respectively), it seems reasonable to assume that the interval described by "around \$20,000" is larger than that described by "around \$300." Moreover, it seems reasonable that \$19,950 should be in the first interval, while \$250 is not in the second. With this choice of language (and the further assumptions), we can capture consumers who might drive an extra 5 kilometers to save \$50 on a \$300 purchase but would not be willing to drive an extra 5 kilometers to save \$50 on a \$20,000 purchase (this point was already made by Thaler 1980): a consumer gets the same utility if they pay \$20,000 or \$19,950 (since in both cases, they are paying "around \$20,000"), but does not get the same utility paying \$250 rather than \$300.

This can be viewed as an application of *Weber's law*, which asserts that the minimum difference between two stimuli necessary for a subject to discriminate between them is proportional to the magnitude of the stimuli; thus, larger stimuli require larger differences between them to be perceived. Although traditionally applied to physical stimuli, Weber's law has also been shown to be applicable in the realm of numerical perception: larger numbers are

subjectively harder to discriminate from one another (Moyer and Landauer 1967, Restle 1978).

As we observed earlier, we can understand the partitions that arise in Jehiel’s notion of a coarsening of the language. This is even more explicit in Mullainathan’s notion of categories. The observation of Manski and Molinari (2010) that people often represent likelihoods using words suggests that coarseness can arise in the representation of likelihood. To see the potential impact of this on decision theoretic concerns, consider the following analysis of Allais’s paradox (Allais 1953).

EXAMPLE 5.3. Consider the two pairs of gambles described in table 3.

Gamble 1a		Gamble 1b	
1	\$1 million	0.89	\$1 million
		0.1	\$5 million
		0.01	\$0
Gamble 2a		Gamble 2b	
0.89	\$0	0.9	\$0
0.11	\$1 million	0.1	\$5 million

The first pair is a choice between (1a) \$1 million for sure, versus (1b) a 0.89 chance of \$1 million, a 0.1 chance of \$5 million, and a 0.01 chance of nothing. The second is a choice between (2a) a 0.89 chance of nothing and a 0.11 chance of \$1 million, versus (2b) a 0.9 chance of nothing and a .1 chance of \$5 million. The “paradox” arises from the fact that most people choose (1a) over (1b), and most people choose (2b) over (2a) (Allais 1953), but these preferences are not simul-

taneously compatible with expected-utility maximization.

Suppose that we apply the observations of Manski and Molinari (2010) to this setting. Specifically, suppose that probability judgments such as “there is a 0.11 chance of getting \$1 million” are represented in language with only finitely many levels of likelihood. In particular, suppose that the language has only the descriptions “no chance,” “slight chance,” “unlikely,” and their respective opposites, “certain,” “near certain,” and “likely,” interpreted as in table 4.

Once we represent likelihoods using words in language rather than numbers, we have to decide how to determine (expected) utility. For definiteness, suppose that the utility of a gamble as described in this language is determined using the interval midpoint representative given in the third column of table 4. Thus, a “slight chance” is effectively treated as a 0.025 probability, a “likely” event as a 0.9 probability, and so on.

Range	Description	Representative
1	certain	1
[0.95, 1)	near certain	0.975
[0.85, 0.95)	likely	0.9
(0.05, 0.15]	unlikely	0.1
(0, 0.05]	slight chance	0.025
0	no chance	0

Revisiting the gambles associated with the Allais paradox, suppose that we replace the actual probability given in table 3 by the word that represents it (i.e., replace 1 by “certain,” 0.89 by “likely,” and so on). This is how we assume that an agent might represent what

he hears. Then, when doing an expected utility calculation, the word is replaced by the probability representing that word, giving us table 5.

TABLE 5
THE ALLAIS PARADOX, COARSELY APPROXIMATED

Gamble 1a		Gamble 1b	
1	\$1 million	0.9	\$1 million
		0.1	\$5 million
		0.025	\$0
Gamble 2a		Gamble 2b	
0.9	\$0	0.9	\$0
0.1	\$1 million	0.1	\$5 million

Using these numbers, we can calculate the revised utility of (1b) to be $0.9 \cdot u_A(\$1 \text{ million}) + 0.1 \cdot u_A(\$5 \text{ million}) + 0.025 \cdot u_A(\$0)$, and this quantity may well be less than $u_A(\$1 \text{ million})$, depending on the utility function u_A . For example, if $u_A(\$1 \text{ million}) = 1$, $u_A(\$5 \text{ million}) = 3$, and $u_A(\$0) = -10$, then the utility of gamble (1b) evaluates to 0.95. In this case, Alice prefers (2b) to (2a) but also prefers (1a) to (1b). Thus, this choice of language rationalizes the observed preferences of many decision-makers. (Rubinstein 1998 offered a closely related analysis.)

It is worth noting that this approach to evaluating gambles will lead to discontinuities; the utility of a gamble that gets, say, \$1,000,000 with probability x and \$5,000,000 with probability $1 - x$ does not converge to the utility of a gamble that gets \$1,000,000 with probability 1 as x approaches 1. Indeed, we would get discontinuities at the boundaries of every range. We expect almost everyone to treat certainty specially, and so have a special category for the range $[1, 1]$; what people take as the range for other descriptions will vary. Andreoni and Sprenger

(2009) present an approach to the Allais paradox that is based on the discontinuity of the utility of gambles at 1 and present experimental evidence for such a discontinuity. We can view the language-based approach as providing a potential explanation for this discontinuity.

Going back to example 5.2, note that cooperating is rational for Alice if she thinks that Bob is sure that she will defect, since cooperating in this case would yield a minimum utility of 0, whereas defecting would result in a utility of -1 . On the other hand, if Alice thinks that Bob is *not* sure that she will defect, then since her utility in this case is determined classically, it is rational for her to defect, as usual. Bjorndahl, Halpern, and Pass (2013) define a natural generalization of Nash equilibrium in language-based games and show that, in general—and, in particular in this game—it does not exist, even if mixed strategies are allowed. The problem is the discontinuity in payoffs. Intuitively, a Nash equilibrium is a state of play where players are happy with their choice of strategies *given correct beliefs about what their opponents will choose*. But there is a fundamental tension between a state of play where everyone has correct beliefs and one where some player successfully surprises another.

Bjorndahl, Halpern, and Pass (2013) also define a natural generalization of the solution concept of *rationalizability* (Bernheim 1984, Pearce 1984), and show that all language-based games where the language satisfies a natural constraint have rationalizable strategies. But the question of finding appropriate solution concepts for language-based games remains open. Moreover, the analysis of Bjorndahl, Halpern, and Pass (2013) was carried out only for normal-form games. Geanakoplos, Pearce, and Stacchetti (1989) and Battigalli and Dufwenberg (2009) consider extensive-form psychological games. Extending language-based games to the

extensive-form setting will require dealing with issues like the impact of the language changing over time.

We conclude this section by observing that, interpreting the choice of language as a framing of the game, language-based games can be seen as a special case of framing. There have already been attempts to provide general models of the effects of framing. For example, Tversky and Simonson (1993) considered situations in which an agent's choices may depend on the background set B (i.e., the set of all available choices) and the choice set S (i.e., the set of offered choices). Tversky and Simonson introduced a choice function $V_B(x, C) = v(x) + \beta f_B(x) + \theta g(x, S)$ consisting of three components: $v(x)$ is the context-free value of x , independent of B , $f_B(x)$ captures the effect of the background, and $g(x, S)$ captures the effect of the choice set. Salant and Rubinstein (2008) and Ellingsen et al. (2012) assumed that there is a set \mathcal{F} of frames and that the utility function depends on the specific frame $F \in \mathcal{F}$.

These models of framing effects can easily explain all seven regularities by choosing suitable frames and utility functions. For example, Ellingsen et al. applied their model to the prisoner's dilemma and were able to explain changes in the rate of cooperation depending on the name of the game ("community game" versus "stock market game"), under the assumption that the frame affects the beliefs about the opponent's level of altruism. Although these models can be applied to explain framing effects specifically generated by language, they do not model the effect of language directly. As experimental regularity 7 shows, many framing effects are in fact ultimately due to language. Language-based games provide a way of capturing these language effects directly. Moreover, they allow us to ask questions that are not asked in the standard framing literature, for example, why people's behavior changes when the price of gas goes from \$3.99 to \$4.00, but not when

it goes from \$3.98 to \$3.99. (This would not typically be called a framing effect; but we can reinterpret it as a framing effect by assuming that there is a frame $F \in \mathcal{F}$ such that "over \$4" and "under \$4" are different categories in F .)

6. Future Research and Outlook

The key takeaway message of this article is that the monetary payoffs associated with actions are not sufficient to fully explain people's behavior. What matters are not just the monetary payoffs, but also the language used to present the available actions. It follows that economic models of behavior should also take language into account. We believe that the shift from outcome-based preferences to language-based preferences will have a profound impact on economics. We conclude our review with a discussion of some lines of research that we believe will play a prominent role in this shift. These lines of research are quite interdisciplinary, involving psychology, sociology, philosophy, and computer science.

In the previous sections, we have highlighted experimental results suggesting that, at least in some cases, people seem to have moral preferences. However, we were deliberately vague about where these moral preferences come from. Do they arise from personal beliefs about what is right and wrong, beliefs about what others approve or disapprove of, or beliefs about what others actually do?

Moral psychologists and moral philosophers have long argued that there are several types of norms, which sometimes conflict with one other. An important distinction is between *personal norms* and *social norms* (Schwartz 1977). Personal norms refer to internal standards about what is considered to be right and wrong; they are not externally motivated by the approval or disapproval of others. It might happen that others either

approve or disapprove of them, but this is not what drives the personal norms. Social norms, on the other hand, refer to “rules and standards that are understood by members of a group, and that guide and/or constrain behavior without the force of laws” (Cialdini and Trost 1998). Two important types of social norms are *injunctive norms* and *descriptive norms*, defining, respectively, what people think others would approve or disapprove of and what people actually do (Cialdini, Reno, and Kallgren 1990). A unified theory of norms has been proposed more recently by Cristina Bicchieri (2006). According to her theory, there are three main classes of norms: *personal normative beliefs*, which are personal beliefs about what should happen in a given situation; *empirical expectations*, which are personal beliefs about what one expects others to do; and *normative expectations*, which are personal beliefs about what others think one should do.

Although the different types of norms often align, as we discussed earlier, they may conflict. When descriptive norms conflict with injunctive norms, people tend to follow the descriptive norm, as shown by a famous field experiment in which people are observed to litter more in a littered environment than in a clean environment (Cialdini, Reno, and Kallgren 1990). Similarly, when empirical and normative expectations are in conflict, people tend to follow the empirical expectations. One potential explanation for this is that people are rarely punished when everyone is engaging in the same behavior (Bicchieri and Xiao 2009). Little is known about what happens when personal norms are in conflict with descriptive or injunctive norms, or when personal normative beliefs are in contrast with empirical and normative expectations. The example that we mentioned in section 4 of a vegan who does not eat food containing animal-derived ingredients while realizing that it is an injunctive norm to do so suggests that, at least in some

cases, personal judgments about what is right and what is wrong represent the dominant motivation for behavior.

In the context of the games considered in this review, some experimental work points toward a significant role of personal norms, at least in one-shot and anonymous interactions. For example, Capraro and Rand (2018) created a laboratory setting in which the personal norm was pitted against the descriptive norm in the trade-off game, and found that participants tended to follow the personal norm. More recently, Catola et al. (2021) showed that personal norms predict cooperative behavior in the public-goods game better than social norms do. The role of personal norms was also highlighted by Bašić and Verrina (2020). They found that both personal and social norms shape behavior in the dictator and ultimatum games, which led them to propose a utility function that takes into account both types of norms, as reviewed in section 4. In any case, we believe that an important direction for future empirical research is an exploration of how people resolve norm conflicts. This may be a key step in allowing us to create a new generation of utility functions that take into account the relative effect of different types of norms.

Another major direction for future work is the exploration of how the heterogeneity in individual personal norms affects economic choices. People differ in their judgments about what they consider right and wrong. For example, some people embrace utilitarian ethics, which dictates that the action selected should maximize total welfare and minimize total harm (Mill 2016, Bentham 1996). Others embrace *deontological ethics*, according to which the rightness or wrongness of an action is entirely determined by whether the action respects certain moral norms and duties, regardless of its consequences (Kant 2002). It has been suggested that people’s personal norms can be

decomposed into fundamental dimensions, although there is some debate about the number and the characterization of these dimensions. According to moral foundations theory, there are six dimensions: care/harm, fairness/cheating, loyalty/betrayal, authority/subversion, sanctity/degradation, and liberty/oppression (Haidt and Joseph 2004, Graham et al. 2013, Iyer et al. 2012, Haidt 2012). According to morality-as-cooperation theory, there are seven dimensions: helping kin, helping your group, reciprocating, being brave, deferring to superiors, dividing disputed resources, and respecting prior possession (Curry 2016; Curry, Mullins, and Whitehouse 2019; Curry, Chesters, and Van Lissa 2019). Each individual assigns different weights to these dimensions. These weights have been shown to play a key role in determining a range of important characteristics, including political orientation (Graham, Haidt, and Nosek 2009). There has been very little work exploring the link between moral dimensions and prosocial behavior. Some preliminary evidence does suggest that different forms of prosocial behavior may be associated with different moral foundations, not necessarily correlated among themselves. For example, while both dictator game giving and ultimatum game rejections appear to be associated with moral preferences, giving appears to be primarily driven by the fairness dimension of morality (Schier, Ockenfels, and Hofmann 2016), whereas ultimatum game rejections seem to be primarily driven by the in-group dimension (Capraro and Rodriguez-Lara 2021). It is worth noting that the fairness and in-group dimensions are not correlated between themselves (Haidt and Joseph 2004, Graham et al. 2013, Iyer et al. 2012, Haidt 2012). Therefore, these preliminary results speak in favor of the multidimensionality of prosociality and may provide a rationalization for the result by Chapman et al. (2023) that the giving cluster of social prefer-

ences is not correlated with the punishment cluster. As said, this line of research has just started.¹⁹ Understanding the link between moral dimensions and prosocial behaviors is a necessary step for building models that can explain human behavior with greater precision.

Since our moral preferences are clearly affected by our social interactions, and it is well known that the structure of an individual's social network affects preferences and outcomes in general (Easley and Kleinberg 2010), we believe that another important line of research is how moral preferences are shaped by social connections. For example, it is known that cooperation is strongly affected by the structure of the social network. Hubs in such networks can act as strong cooperative centers and exert a positive influence on the peripheral nodes (Santos, Pacheco, and Lenaerts 2006). The ability to break ties with unfair or exploitative partners and make new ones with those of better reputation also favorably affects cooperation (Perc and Szolnoki 2010). More recent research has shown that other forms of moral behavior, such as truth telling and honesty, are also strongly affected by the properties of social networks (Capraro, Perc, and Vilone 2020). And a case has been made for further explorations being much needed along these lines (Capraro and Perc 2018), for example, by studying how network structure affects different types of moral behavior, including equity, efficiency, and trustworthiness.

Another line of research involves language-based preferences. To the extent that people do have language-based preferences, it would be useful to be able to predict

¹⁹There is some work exploring the link between moral foundations and cooperation in the prisoner's dilemma (Clark et al. 2017). Moreover, a working paper by Bonneau (2021) explores the role of different moral foundations in explaining altruistic behavior in different forms of the dictator game, including the dictator game with a take option and the dictator game with an exit option.

how people will behave in an economic decision problem described by a language. A relatively new area of research in computational linguistics may be relevant in this regard. *Sentiment analysis* (e.g., Pang, Lee, and Vaithyanathan 2002; Pang and Lee 2004; Esuli and Sebastiani 2007) aims to determine the attitude of a speaker or a writer to a topic from the information contained in a document. For example, we may want to determine the feelings of a reviewer about a movie from his review. Among other things, sentiment analysis attempts to associate to a description (in a given context) its *polarity*, that is, a number in the interval $[-1, 1]$ expressing a positive, negative, or indifferent sentiment. One could perhaps use sentiment analysis to define a utility function by taking the polarity of the description of strategies into account. The idea is that people are reluctant to perform actions that evoke negative sentiments, like stealing, but are eager to perform actions that evoke positive sentiments; the utility function could take this into account (in addition to taking into account the monetary consequences of actions). Being able to associate utility to words would also allow us to measure the explanatory power of language-based models. Experimental regularity 7 shows that language-based models explain behavior in some games better than models based solely on monetary outcomes. However, to the best of our knowledge, there has been no econometric study measuring the exploratory power of language-based models.

We have only scratched the surface of potential research directions. We believe that economics is taking brave steps into uncharted territory, here. New ideas will be needed, and bridges to other fields will need to be built. The prospects are most exciting!

REFERENCES

Abeler, Johannes, Daniele Nosenzo, and Collin Raymond. 2019. "Preferences for Truth-Telling." *Econo-*

- metrica* 87 (4): 1115–53.
- Abrams, Burton A., and Mark D. Schitz. 1978. "The 'Crowding-Out' Effect of Governmental Transfers on Private Charitable Contributions." *Public Choice* 33 (1): 29–39.
- Abrams, Burton A., and Mark D. Schmitz. 1984. "The Crowding-Out Effect of Governmental Transfers on Private Charitable Contributions: Cross-Section Evidence." *National Tax Journal* 37 (4): 563–68.
- Ahn, David S., and Haluk Ergin. 2010. "Framing Contingencies." *Econometrica* 78 (2): 655–95.
- Akerlof, George A., and Rachel E. Kranton. 2000. "Economics and Identity." *Quarterly Journal of Economics* 115 (3): 715–53.
- Alger, Ingela, and Régis Renault. 2006. "Screening Ethics When Honest Agents Care about Fairness." *International Economic Review* 47 (1): 59–85.
- Alger, Ingela, and Régis Renault. 2007. "Screening Ethics When Honest Agents Keep Their Word." *Economic Theory* 30 (2): 291–311.
- Alger, Ingela, and Jörgen W. Weibull. 2013. "Homo Moralis—Preference Evolution under Incomplete Information and Assortative Matching." *Econometrica* 81 (6): 2269–2302.
- Alger, Ingela, and Jörgen W. Weibull. 2016. "Evolution and Kantian Morality." *Games and Economic Behavior* 98: 56–67.
- Allais, M. 1953. "Le comportement de l'Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l'École Américaine." *Econometrica* 21 (4): 503–46.
- Andreoni, James. 1988. "Privately Provided Public Goods in a Large Economy: The Limits of Altruism." *Journal of Public Economics* 35 (1): 57–73.
- Andreoni, James. 1990. "Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving." *Economic Journal* 100 (401): 464–77.
- Andreoni, James. 1995. "Cooperation in Public-Goods Experiments: Kindness or Confusion?" *American Economic Review* 85 (4): 891–904.
- Andreoni, James, and B. Douglas Bernheim. 2009. "Social Image and the 50–50 Norm: A Theoretical and Experimental Analysis of Audience Effects." *Econometrica* 77 (5): 1607–36.
- Andreoni, James, and John Miller. 2002. "Giving According to GARP: An Experimental Test of the Consistency of Preferences for Altruism." *Econometrica* 70 (2): 737–53.
- Andreoni, James, and Charles Sprenger. 2009. "Certain and Uncertain Utility: The Allais Paradox and Five Decision Theory Phenomena." <https://econweb.ucsd.edu/~jandreoni/WorkingPapers/AndreoniSprengerFivePhenomena.pdf>.
- Bacharach, Michael. 1999. "Interactive Team Reasoning: A Contribution to the Theory of Co-operation."

- Research in Economics* 53 (2): 117–47.
- Bandiera, Oriana, Iwan Barankay, and Imran Rasul. 2005. “Social Preferences and the Response to Incentives: Evidence from Personnel Data.” *Quarterly Journal of Economics* 120 (3): 917–62.
- Bardsley, Nicholas. 2008. “Dictator Game Giving: Altruism or Artefact?” *Experimental Economics* 11 (2): 122–33.
- Bašić, Zvonimir, and Eugenio Verrina. 2020. “Personal Norms—and Not Only Social Norms—Shape Economic Behavior.” MPI Collective Goods Discussion Paper 2020/25.
- Basu, Kaushik. 1994. “The Traveler’s Dilemma: Paradoxes of Rationality in Game Theory.” *American Economic Review* 84 (2): 391–95.
- Battigalli, Pierpaolo, and Martin Dufwenberg. 2009. “Dynamic Psychological Games.” *Journal of Economic Theory* 144 (1): 1–35.
- Battigalli, Pierpaolo, and Martin Dufwenberg. 2022. “Belief-Dependent Motivations and Psychological Game Theory.” *Journal of Economic Literature* 60 (3): 833–82.
- Bénabou, Roland, and Jean Tirole. 2006. “Incentives and Prosocial Behavior.” *American Economic Review* 96 (5): 1652–78.
- Bénabou, Roland, and Jean Tirole. 2011. “Identity, Morals, and Taboos: Beliefs as Assets.” *Quarterly Journal of Economics* 126 (2): 805–55.
- Bentham, Jeremy. 1996. *The Collected Works of Jeremy Bentham: An Introduction to the Principles of Morals and Legislation*. Oxford: Clarendon Press.
- Benz, Matthias, and Stephan Meier. 2008. “Do People Behave in Experiments as in the Field? Evidence from Donations.” *Experimental Economics* 11 (3): 268–81.
- Bernheim, B. Douglas. 1984. “Rationalizable Strategic Behavior.” *Econometrica* 52 (4): 1007–28.
- Bernheim, B. Douglas. 1986. “On the Voluntary and Involuntary Provision of Public Goods.” *American Economic Review* 76 (4): 789–93.
- Bernheim, B. Douglas. 1994. “A Theory of Conformity.” *Journal of Political Economy* 102 (5): 841–77.
- Bertrand, Joseph. 1883. “Review of *Theorie Mathématique de la Richesse Sociale* and *Recherches sur les Principes Mathématiques de la Theorie des Richesses*.” *Journal des Savants*: 499–508.
- Bicchieri, Cristina. 2006. *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge, UK: Cambridge University Press.
- Bicchieri, Cristina, and Alex Chavez. 2010. “Behaving as Expected: Public Information and Fairness Norms.” *Journal of Behavioral Decision Making* 23 (2): 161–78.
- Bicchieri, Cristina, and Erte Xiao. 2009. “Do the Right Thing: But Only If Others Do So.” *Journal of Behavioral Decision Making* 22 (2): 191–208.
- Binmore, Ken. 1994. *Game Theory and the Social Contract: Playing Fair*. Cambridge, MA: MIT Press.
- Biziou-van-Pol, Laura, Jana Haenen, Arianna Novaro, Andrés Occhipinti Liberman, and Valerio Capraro. 2015. “Does Telling White Lies Signal Pro-social Preferences?” *Judgment and Decision Making* 10: 538–48.
- Bjorndahl, Adam, Joseph Y. Halpern, and Rafael Pass. 2013. “Language-Based Games.” In *Proceedings of the 14th Conference on Theoretical Aspects of Rationality and Knowledge*, 39–48.
- Blount, Sally. 1995. “When Social Outcomes Aren’t Fair: The Effect of Causal Attributions on Preferences.” *Organizational Behavior and Human Decision Processes* 63 (2): 131–44.
- Blume, Andreas, and Oliver Board. 2014. “Intentional Vagueness.” *Erkenntnis* 79: 855–99.
- Blume, Lawrence E., David Easley, and Joseph Y. Halpern. 2021. “Constructive Decision Theory.” *Journal of Economic Theory* 196: 105305.
- Bolle, Friedel, and Peter Ockenfels. 1990. “Prisoners’ Dilemma as a Game with Incomplete Information.” *Journal of Economic Psychology* 11 (1): 69–84.
- Bolton, Gary E., and Axel Ockenfels. 2000. “ERC: A Theory of Equity, Reciprocity, and Competition.” *American Economic Review* 90 (1): 166–93.
- Bonneau, Maxime. 2021. “Can Different Moral Foundations Help Explain Different Forms of Altruistic Behavior?” Unpublished.
- Bonomi, Giampaolo, Nicola Gennaioli, and Guido Tabellini. 2021. “Identity, Beliefs, and Political Conflict.” *Quarterly Journal of Economics* 136 (4): 2371–2411.
- Bott, Kristina M., Alexander W. Cappelen, Erik Ø. Sørensen, and Bertil Tungodden. 2019. “You’ve Got Mail: A Randomised Field Experiment on Tax Evasion.” *Management Science* 66 (7): 2801–3294.
- Brañas-Garza, Pablo. 2007. “Promoting Helping Behavior with Framing in Dictator Games.” *Journal of Economic Psychology* 28 (4): 477–86.
- Brañas-Garza, Pablo, Valerio Capraro, and Ericka Rascón-Ramírez. 2018. “Gender Differences in Altruism on Mechanical Turk: Expectations and Actual Behaviour.” *Economics Letters* 170: 19–23.
- Brekke, Kjell Arne, Snorre Kverndokk, and Karine Nyborg. 2003. “An Economic Model of Moral Motivation.” *Journal of Public Economics* 87 (9–10): 1967–83.
- Burnham, Terence, Kevin McCabe, and Vernon L. Smith. 2000. “Friend-or-Foe Intentionality Priming in an Extensive Form Trust Game.” *Journal of Economic Behavior and Organization* 43 (1): 57–73.
- Camerer, Colin F. 2003. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton, NJ:

- Princeton University Press.
- Cappelen, Alexander W., Astri Drange Hole, Erik Ø. Sørensen, and Bertil Tungodden. 2007. "The Pluralism of Fairness Ideals: An Experimental Approach." *American Economic Review* 97 (3): 818–27.
- Cappelen, Alexander W., Ulrik H. Nielsen, Erik Ø. Sørensen, Bertil Tungodden, and Jean-Robert Tyran. 2013. "Give and Take in Dictator Games." *Economics Letters* 118 (2): 280–83.
- Cappelen, Alexander W., Erik Ø. Sørensen, and Bertil Tungodden. 2013. "When Do We Lie?" *Journal of Economic Behavior and Organization* 93: 258–65.
- Capraro, Valerio. 2013. "A Model of Human Cooperation in Social Dilemmas." *PLoS ONE* 8 (8): e72427.
- Capraro, Valerio. 2018. "Social versus Moral Preferences in the Ultimatum Game: A Theoretical Model and an Experiment." <http://dx.doi.org/10.2139/ssrn.3155257>.
- Capraro, Valerio, Glorianna Jagfeld, Rana Klein, Mathijs Mul, and Iris van de Pol. 2019. "Increasing Altruistic and Cooperative Behaviour with Simple Moral Nudges." *Scientific Reports* 9 (1): 11880.
- Capraro, Valerio, and Matjaž Perc. 2018. "Grand Challenges in Social Physics: In Pursuit of Moral Behavior." *Frontiers in Physics* 6: 107.
- Capraro, Valerio, and Matjaž Perc. 2021. "Mathematical Foundations of Moral Preferences." *Journal of the Royal Society Interface* 18 (175): 20200880.
- Capraro, Valerio, Matjaž Perc, and Daniele Vilone. 2020. "Lying on Networks: The Role of Structure and Topology in Promoting Honesty." *Physical Review E* 101: 032305.
- Capraro, Valerio, and David G. Rand. 2018. "Do the Right Thing: Experimental Evidence That Preferences for Moral Behavior, Rather than Equity or Efficiency per SE, Drive Human Prosociality." *Judgment and Decision Making* 13 (1): 99–111.
- Capraro, Valerio, and Ismael Rodriguez-Lara. 2021. "Moral Preferences in Bargaining Games." <https://ssrn.com/abstract=3933603>.
- Capraro, Valerio, and Andrea Vanzo. 2019. "The Power of Moral Words: Loaded Language Generates Framing Effects in the Extreme Dictator Game." *Judgment and Decision Making* 14 (3): 309–17.
- Catola, Marco, Simone D'Alessandro, Pietro Guarnieri, and Veronica Pizziol. 2021. "Personal Norms in the Online Public Good Game." *Economics Letters* 207: 110024.
- Chang, Daphne, Roy Chen, and Erin Krupka. 2019. "Rhetoric Matters: A Social Norms Explanation for the Anomaly of Framing." *Games and Economic Behavior* 116: 158–78.
- Chapman, Jonathan, Mark Dean, Pietro Ortoleva, Erik Snowberg, and Colin Camerer. 2023. "Econographics." *Journal of Political Economy Microeconomics* 1 (1): 115–61.
- Charness, Gary, and Brit Grosskopf. 2001. "Relative Payoffs and Happiness: An Experimental Study." *Journal of Economic Behavior and Organization* 45 (3): 301–28.
- Charness, Gary, and Matthew Rabin. 2002. "Understanding Social Preferences with Simple Tests." *Quarterly Journal of Economics* 117 (3): 817–69.
- Cialdini, Robert B., Raymond R. Reno, and Carl A. Kallgren. 1990. "A Focus Theory of Normative Conduct: Recycling the Concept of Norms to Reduce Littering in Public Places." *Journal of Personality and Social Psychology* 58 (6): 1015–26.
- Cialdini, Robert B., and Melanie R. Trost. 1998. "Social Influence: Social Norms, Conformity and Compliance." In *The Handbook of Social Psychology*, edited by Daniel T. Gilbert, Susan T. Fiske and Gardner Lindzey, 151–92. New York: McGraw-Hill.
- Clark, C. Brendan, Jeffrey A. Swails, Heidi M. Pontinen, Shannon E. Bowerman, Kenneth A. Kriz, and Peter S. Hendricks. 2017. "A Behavioral Economic Assessment of Individualizing versus Binding Moral Foundations." *Personality and Individual Differences* 112: 49–54.
- Curry, Oliver Scott. 2016. "Morality as Cooperation: A Problem-Centred Approach." In *The Evolution of Morality*, edited by Todd K. Shackelford and Ranald D. Hansen, 27–51. Cham, Switzerland: Springer.
- Curry, Oliver Scott, Matthew Jones Chesters, and Caspar J. Van Lissa. 2019. "Mapping Morality with a Compass: Testing the Theory of 'Morality-as-Cooperation' with a New Questionnaire." *Journal of Research in Personality* 78: 106–24.
- Curry, Oliver Scott, Daniel Austin Mullins, and Harvey Whitehouse. 2019. "Is It Good to Cooperate? Testing the Theory of Morality-as-Cooperation in 60 Societies." *Current Anthropology* 60 (1): 47–69.
- Dal Bó, Ernesto, and Pedro Dal Bó. 2014. "'Do the Right Thing': The Effects of Moral Suasion on Cooperation." *Journal of Public Economics* 117: 28–38.
- Dana, Jason, Daylian M. Cain, and Robyn M. Dawes. 2006. "What You Don't Know Won't Hurt Me: Costly (But Quiet) Exit in Dictator Games." *Organizational Behavior and Human Decision Processes* 100 (2): 193–201.
- DellaVigna, Stefano, John A. List, and Ulrike Malmendier. 2012. "Testing for Altruism and Social Pressure in Charitable Giving." *Quarterly Journal of Economics* 127 (1): 1–56.
- Dhaene, Geert, and Jan Bouckaert. 2010. "Sequential Reciprocity in Two-Player, Two-Stage Games: An Experimental Analysis." *Games and Economic Behavior* 70 (2): 289–303.
- Dhami, Sanjit. 2016. *The Foundations of Behavioral Economic Analysis*. Oxford, UK: Oxford University

- Press.
- Dreber, Anna, Tore Ellingsen, Magnus Johannesson, and David G. Rand. 2013. "Do People Care about Social Context? Framing Effects in Dictator Games." *Experimental Economics* 16 (3): 349–71.
- Dufwenberg, Martin, and Georg Kirchsteiger. 2004. "A Theory of Sequential Reciprocity." *Games and Economic Behavior* 47 (2): 268–98.
- Easley, David, and Jon Kleinberg. 2010. *Networks, Crowds, and Markets: Reasoning about a Highly Connected World*. Cambridge, UK: Cambridge University Press.
- Eckel, Catherine C., and Philip J. Grossman. 1996. "Altruism in Anonymous Dictator Games." *Games and Economic Behavior* 16 (2): 181–91.
- Edgeworth, Francis Ysidro. 1881. *Mathematical Psychics: An Essay on the Application of Mathematics to the Moral Sciences*. London: C. Kegan Paul and Co.
- Ellingsen, Tore, and Magnus Johannesson. 2008. "Pride and Prejudice: The Human Side of Incentive Theory." *American Economic Review* 98 (3): 990–1008.
- Ellingsen, Tore, Magnus Johannesson, Johanna Mollerstrom, and Sara Munkhammar. 2012. "Social Framing Effects: Preferences or Beliefs?" *Games and Economic Behavior* 76 (1): 117–30.
- Engel, Christoph. 2011. "Dictator Games: A Meta Study." *Experimental Economics* 14 (4): 583–610.
- Engelmann, Dirk, and Martin Strobel. 2004. "Inequality Aversion, Efficiency, and Maximin Preferences in Simple Distribution Experiments." *American Economic Review* 94 (4): 857–69.
- Enke, Benjamin. 2019. "Kinship, Cooperation, and the Evolution of Moral Systems." *Quarterly Journal of Economics* 134 (2): 953–1019.
- Enke, Benjamin, Mattias Polborn, and Alex Wu. 2022. "Values as Luxury Goods and Political Polarization." NBER Working Paper 30001.
- Erat, Sanjiv, and Uri Gneezy. 2012. "White Lies." *Management Science* 58 (4): 723–33.
- Eriksson, Kimmo, Pontus Strimling, Per A. Andersson, and Torun Lindholm. 2017. "Costly Punishment in the Ultimatum Game Evokes Moral Concern, in Particular When Framed as Payoff Reduction." *Journal of Experimental Social Psychology* 69: 59–64.
- Esuli, Andrea, and Fabrizio Sebastiani. 2007. "SentiWordNet: A High-Coverage Lexical Resource for Opinion Mining." Instituto di Scienza e Tecnologie dell'Informazione Technical Report 2007-TR-02.
- Falk, Armin, Ernst Fehr, and Urs Fischbacher. 2003. "On the Nature of Fair Behavior." *Economic Inquiry* 41 (1): 20–26.
- Falk, Armin, Ernst Fehr, and Urs Fischbacher. 2008. "Testing Theories of Fairness—Intentions Matter." *Games and Economic Behavior* 62 (1): 287–303.
- Falk, Armin, and Urs Fischbacher. 2006. "A Theory of Reciprocity." *Games and Economic Behavior* 54 (2): 293–315.
- Fehr, Ernst, and Klaus M. Schmidt. 1999. "A Theory of Fairness, Competition, and Cooperation." *Quarterly Journal of Economics* 114 (3): 817–68.
- Fehr, Ernst, and Klaus M. Schmidt. 2006. "The Economics of Fairness, Reciprocity and Altruism—Experimental Evidence and New Theories." In *Handbook of the Economics of Giving, Altruism and Reciprocity*, Vol. 1, edited by Serge-Christophe Kolm and Jean Mercier Ythier, 615–91.
- Fischbacher, Urs, and Franziska Föllmi-Heusi. 2013. "Lies in Disguise? An Experimental Study on Cheating." *Journal of the European Economic Association* 11 (3): 525–47.
- Forsythe, Robert, Joel L. Horowitz, Nathan E. Savin, and Martin Sefton. 1994. "Fairness in Simple Bargaining Experiments." *Games and Economic Behavior* 6 (3): 347–69.
- Frey, Bruno S., and Felix Oberholzer-Gee. 1997. "The Cost of Price Incentives: An Empirical Analysis of Motivation Crowding-Out." *American Economic Review* 87 (4): 746–55.
- Gärdenfors, Peter, and Nils-Eric Sahlin. 1982. "Unreliable Probabilities, Risk Taking, and Decision Making." *Synthese* 53 (3): 361–86.
- Geanakoplos, John, David Pearce, and Ennio Stacchetti. 1989. "Psychological Games and Sequential Rationality." *Games and Economic Behavior* 1 (1): 60–80.
- Gerlach, Philipp, Kinneret Teodorescu, and Ralph Hertwig. 2019. "The Truth about Lies: A Meta-Analysis on Dishonest Behavior." *Psychological Bulletin* 145 (1): 1–44.
- Gilbert, Margaret. 1987. "Modelling Collective Belief." *Synthese* 73 (1): 185–204.
- Gilboa, Itzhak, and David Schmeidler. 1989. "Maxmin Expected Utility with a Non-unique Prior." *Journal of Mathematical Economics* 18 (2): 141–53.
- Gilboa, Itzhak, and David Schmeidler. 2001. *A Theory of Case-Based Decisions*. Cambridge, UK: Cambridge University Press.
- Gintis, Herbert, Eric Alden Smith, and Samuel Bowles. 2001. "Costly Signaling and Cooperation." *Journal of Theoretical Biology* 213 (1): 103–19.
- Glazer, Amihai, and Kai A. Konrad. 1996. "A Signaling Explanation for Charity." *American Economic Review* 86 (4): 1019–28.
- Gneezy, Uri. 2005. "Deception: The Role of Consequences." *American Economic Review* 95 (1): 384–94.
- Gneezy, Uri, Agne Kajackaite, and Joel Sobel. 2018. "Lying Aversion and the Size of the Lie." *American Economic Review* 108 (2): 419–53.
- Gneezy, Uri, Bettina Rockenbach, and Marta Ser-

- ra-Garcia. 2013. "Measuring Lying Aversion." *Journal of Economic Behavior and Organization* 93: 293–300.
- Gneezy, Uri, and Aldo Rustichini. 2000. "Pay Enough or Don't Pay at All." *Quarterly Journal of Economics* 115 (3): 791–810.
- Graham, Jesse et al. 2013. "Moral Foundations Theory: The Pragmatic Validity of Moral Pluralism." In *Advances in Experimental Social Psychology*, Vol. 47, edited by Patricia Devine and Ashby Plant, 55–130. San Diego: Academic Press.
- Graham, Jesse, Jonathan Haidt, and Brian A. Nosek. 2009. "Liberals and Conservatives Rely on Different Sets of Moral Foundations." *Journal of Personality and Social Psychology* 96 (5): 1029–46.
- Güth, Werner, Rolf Schmittberger, and Bernd Schwarze. 1982. "An Experimental Analysis of Ultimatum Bargaining." *Journal of Economic Behavior and Organization* 3 (4): 367–88.
- Haidt, Jonathan. 2012. *The Righteous Mind: Why Good People Are Divided by Politics and Religion*. New York: Pantheon Books.
- Haidt, Jonathan, and Craig Joseph. 2004. "Intuitive Ethics: How Innately Prepared Intuitions Generate Culturally Variable Virtues." *Daedalus* 133 (4): 55–66.
- Halpern, Joseph Y., and Willemien Kets. 2015. "Ambiguous Language and Common Priors." *Games and Economic Behavior* 90: 171–80.
- Halpern, Joseph Y., and Nan Rong. 2010. "Cooperative Equilibrium." In *AAMAS '10: Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*, 1465–66. Richland, SC: International Foundation for autonomous Agents and Multiagent Systems.
- Halvorsen, Trond U. 2015. "Are Dictators Loss Averse?" *Rationality and Society* 27 (4): 469–91.
- Hardin, Garrett. 1968. "The Tragedy of the Commons." *Science* 162 (3859): 1243–48.
- Hauge, Karen Evelyn, Kjell Arne Brekke, Lars-Olof Johansson, Olof Johansson-Stenman, and Henrik Svedsäter. 2016. "Keeping Others in Our Mind or in Our Heart? Distribution Games under Cognitive Load." *Experimental Economics* 19 (3): 562–76.
- Henrich, Joseph et al. 2001. "In Search of Homo Economicus: Behavioral Experiments in 15 Small-Scale Societies." *American Economic Review* 91 (2): 73–78.
- Iyer, Ravi, Spassena Koleva, Jesse Graham, Peter Ditto, and Jonathan Haidt. 2012. "Understanding Libertarian Morality: The Psychological Dispositions of Self-Identified Libertarians." *PLoS ONE* 7 (8): e42366.
- Jehiel, Philippe. 2005. "Analogy-Based Expectation Equilibrium." *Journal of Economic Theory* 123 (2): 81–104.
- Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler. 1986. "Fairness and the Assumptions of Economics." *Journal of Business* 59 (4): S285–300.
- Kahneman, Daniel, and Amos Tversky. 1979. "Prospect Theory, an Analysis of Decision under Risk." *Econometrica* 47 (2): 263–92.
- Kant, Immanuel. 2002. *Groundwork for the Metaphysics of Morals*. New Haven, CT: Yale University Press.
- Kessler, Judd B., and Stephen Leider. 2012. "Norms and Contracting." *Management Science* 58 (1): 62–77.
- Kimbrough, Erik O., and Alexander Vostroknutov. 2016. "Norms Make Preferences Social." *Journal of the European Economic Association* 14 (3): 608–38.
- Kimbrough, Erik O., and Alexander Vostroknutov. 2020. "A Theory of Injunctive Norms." <http://dx.doi.org/10.2139/ssrn.3566589>.
- Kőszegi, Botond, and Matthew Rabin. 2006. "A Model of Reference-Dependent Preferences." *Quarterly Journal of Economics* 121 (4): 1133–65.
- Kritikos, Alexander, and Friedel Bolle. 2001. "Distributional Concerns: Equity-or Efficiency-Oriented?" *Economics Letters* 73 (3): 333–38.
- Krueger, Joachim, and Russell W. Clement. 1994. "Memory-Based Judgments about Multiple Categories: A Revision and Extension of Tajfel's Accentuation Theory." *Journal of Personality and Social Psychology* 67 (1): 35–47.
- Krupka, Erin L., and Roberto A. Weber. 2013. "Identifying Social Norms Using Coordination Games: Why Does Dictator Game Sharing Vary?" *Journal of the European Economic Association* 11 (3): 495–524.
- Laffont, Jean-Jacques. 1975. "Macroeconomic Constraints, Economic Efficiency and Ethics: An Introduction to Kantian Economics." *Economica* 42 (168): 430–37.
- Lazear, Edward P., Ulrike Malmendier, and Roberto A. Weber. 2012. "Sorting in Experiments with Application to Social Preferences." *American Economic Journal: Applied Economics* 4 (1): 136–63.
- Ledyard, John O. 1995. "Public Goods: A Survey of Experimental Research." In *The Handbook of Experimental Economics*, edited by John H. Kagel and Alvin E. Roth, 111–94. Princeton, NJ: Princeton University Press.
- Levine, David K. 1998. "Modeling Altruism and Spitefulness in Experiments." *Review of Economic Dynamics* 1 (3): 593–622.
- Levitt, Steven D., and John A. List. 2007. "What Do Laboratory Experiments Measuring Social Preferences Reveal about the Real World?" *Journal of Economic Perspectives* 21 (2): 153–74.
- Lin, Haijin, and Shyam Sunder. 2002. "Using Experimental Data to Model Bargaining Behavior in Ultimatum Games." In *Experimental Business Research*, edited by Rami Zwick and Amnon Rapoport, 373–97.

- Boston: Springer.
- Lipman, Barton L. 1999. "Decision Theory without Logical Omniscience: Toward an Axiomatic Framework for Bounded Rationality." *Review of Economic Studies* 66 (2): 339–61.
- List, John A. 2006. "The Behavioralist Meets the Market: Measuring Social Preferences and Reputation Effects in Actual Transactions." *Journal of Political Economy* 114 (1): 1–37.
- List, John A. 2007. "On the Interpretation of Giving in Dictator Games." *Journal of Political Economy* 115 (3): 482–93.
- López-Pérez, Raúl. 2008. "Aversion to Norm-Breaking: A Model." *Games and Economic Behavior* 64 (1): 237–67.
- Manski, Charles F., and Francesca Molinari. 2010. "Rounding Probabilistic Expectations in Surveys." *Journal of Business and Economic Statistics* 28 (2): 219–231.
- McCabe, Kevin A., Mary L. Rigdon, and Vernon L. Smith. 2003. "Positive Reciprocity and Intentions in Trust Games." *Journal of Economic Behavior and Organization* 52 (2): 267–75.
- McNeil, Barbara J., Stephen J. Pauker, Harold C. Sox Jr., and Amos Tversky. 1982. "On the Elicitation of Preferences for Alternative Therapies." *New England Journal of Medicine* 306 (21): 1259–62.
- Miehl, Laura, Axel Buchner, and Raoul Bell. 2021. "Moral Labels Increase Cooperation and Costly Punishment in a Prisoner's Dilemma Game with Punishment Option." *Scientific Reports* 11 (1): 10221.
- Mill, John Stuart. 2016. "Utilitarianism." In *Seven Masterpieces of Philosophy*, edited by Steven M. Cahn, 329–75. London and New York: Routledge.
- Moyer, Robert S., and Thomas K. Landauer. 1967. "Time Required for Judgements of Numerical Inequality." *Nature* 215 (5109): 1519–20.
- Mullainathan, Sendhil. 2002. "Thinking through Categories." <https://oz.stern.nyu.edu/seminar/sp03/0313background.pdf>.
- Niehans, Jürg. 1948. "Zur Preisbildung bei ungewissen Erwartungen." *Swiss Journal of Economics and Statistics* 84 (5): 433–56.
- Nowak, Martin A., and Karl Sigmund. 1998. "Evolution of Indirect Reciprocity by Image Scoring." *Nature* 393 (6685): 573–77.
- Olson, Mancur. 2009. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Cambridge, MA: Harvard University Press.
- Ostrom, Elinor. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge, UK: Cambridge University Press.
- Pang, Bo, and Lillian Lee. 2004. "A Sentimental Education: Sentiment Analysis Using Subjectivity Summarization Based on Minimum Cuts." In *Proceedings for the 42nd Annual Meeting of the Association for Computational Linguistics (ACL-04)*, 271–78. Kerrville, TX: Association for Computational Linguistics.
- Pang, Bo, Lillian Lee, and Shivakumar Vaithyanathan. 2002. "Thumbs Up? Sentiment Classification Using Machine Learning Techniques." In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP 2022)*, 79–86. Kerrville, TX: Association for Computational Linguistics.
- Pearce, David G. 1984. "Rationalizable Strategic Behavior and the Problem of Perfection." *Econometrica* 52 (4): 1029–50.
- Perc, Matjaž, and Attila Szolnoki. 2010. "Coevolutionary Games—A Mini Review." *BioSystems* 99 (2): 109–25.
- Rabin, Matthew. 1993. "Incorporating Fairness into Game Theory and Economics." *American Economic Review* 83 (5): 1281–1302.
- Rapoport, Anatol, and Albert M. Chammah. 1965. *Prisoner's Dilemma: A Study in Conflict and Cooperation*. Ann Arbor: University of Michigan Press.
- Restle, Frank. 1978. "Speed of Adding and Comparing Numbers." *Journal of Experimental Psychology* 83 (2): 274–78.
- Roberts, Russell D. 1984. "A Positive Model of Private Charity and Public Transfers." *Journal of Political Economy* 92 (1): 136–48.
- Roemer, John E. 2010. "Kantian Equilibrium." *Scandinavian Journal of Economics* 112 (1): 1–24.
- Ross, Lee, and Andrew Ward. 1996. "Naive Realism in Everyday Life: Implications for Social Conflict and Misunderstanding." In *Values and Knowledge*, edited by Edward S. Reed, Elliot Turiel and Terrance Brown, 103–35. Mahwah, NJ: Lawrence Erlbaum Associates, Inc.
- Rubinstein, Ariel. 1998. *Modeling Bounded Rationality*. Cambridge, MA: MIT Press.
- Salant, Yuval, and Ariel Rubinstein. 2008. "(A, f): Choice with Frames." *Review of Economic Studies* 75 (4): 1287–96.
- Santos, F. C., J. M. Pacheco, and Tom Lenaerts. 2006. "Evolutionary Dynamics of Social Dilemmas in Structured Heterogeneous Populations." *Proceedings of the National Academy of Sciences* 103 (9): 3490–94.
- Savage, L. J. 1951. "The Theory of Statistical Decision." *Journal of the American Statistical Association* 46 (253): 55–67.
- Schier, Uta K., Axel Ockenfels, and Wilhelm Hofmann. 2016. "Moral Values and Increasing Stakes in a Dictator Game." *Journal of Economic Psychology* 56: 107–15.
- Schotter, Andrew, Keith Weigelt, and Charles Wilson. 1994. "A Laboratory Investigation of Multiperson Rationality and Presentation Effects." *Games and*

- Economic Behavior* 6 (3): 445–68.
- Schwartz, Shalom H. 1977. “Normative Influences on Altruism.” *Advances in Experimental Social Psychology* 10: 221–79.
- Sen, Amartya K. 1977. “Rational Fools: A Critique of the Behavioral Foundations of Economic Theory.” *Philosophy and Public Affairs* 6 (4): 317–44.
- Smith, Adam. 2010. *The Theory of Moral Sentiments*. London: Penguin Classics, (Orig. pub. 1759).
- Smith, Eric Alden, and Rebecca L. Bliege Bird. 2000. “Turtle Hunting and Tombstone Opening: Public Generosity as Costly Signaling.” *Evolution and Human Behavior* 21 (4): 245–61.
- Sobel, Joel. 2005. “Interdependent Preferences and Reciprocity.” *Journal of Economic Literature* 43 (2): 392–436.
- Stigler, George J., and Gary S. Becker. 1977. “De Gustibus Non Est Disputandum.” *American Economic Review* 67 (2): 76–90.
- Strulov-Shlain, Avner. 2023. “More than a Penny’s Worth: Left-Digit Bias and Firm Pricing.” *Review of Economic Studies* 90 (5): 2612–45.
- Sugden, Robert. 2000. “Team Preferences.” *Economics and Philosophy* 16 (2): 175–204.
- Swope, Kurtis, John Cadigan, Pamela Schmitt, and Robert Shupp. 2008. “Social Position and Distributive Justice: Experimental Evidence.” *Southern Economic Journal* 74 (3): 811–18.
- Tabellini, Guido. 2008. “Institutions and Culture.” *Journal of the European Economic Association* 6 (2–3): 255–94.
- Tappin, Ben M., and Valerio Capraro. 2018. “Doing Good vs. Avoiding Bad in Prosocial Choice: A Refined Test and Extension of the Morality Preference Hypothesis.” *Journal of Experimental Social Psychology* 79: 64–70.
- Thaler, Richard. 1980. “Towards a Positive Theory of Consumer Choice.” *Journal of Economic Behavior and Organization* 1 (1): 39–60.
- Titmuss, Richard M. 2018. *The Gift Relationship (Reissue): From Human Blood to Social Policy*. Bristol, UK: Policy Press.
- Tversky, Amos, and Daniel Kahneman. 1985. “The Framing of Decisions and the Psychology of Choice.” In *Behavioral Decision Making*, edited by George Wright, 25–41. Boston: Springer.
- Tversky, Amos, and Itamar Simonson. 1993. “Context-Dependent Preferences.” *Management Science* 39 (10): 1179–89.
- Wald, Abraham. 1950. *Statistical Decision Functions*. New York: Wiley.
- Warr, Peter G. 1982. “Pareto Optimal Redistribution and Private Charity.” *Journal of Public Economics* 19 (1): 131–38.