Quantifying accuracy and bias in motive introspection

Ryan W. Carlson[1], Yoonseo Zoh[2], Adam Morris[2], & M. J. Crockett[2]

[1]Department of Psychology, Yale University

[2]Department of Psychology, Princeton University

Author Note

Correspondence concerning this article should be addressed to Ryan Carlson, Department of

Psychology, Yale University, 2 Hillhouse Avenue, New Haven, CT 06511, United States or M.J.

Crockett, Department of Psychology, Princeton University, Peretsman Scully Hall, Princeton, NJ

08544. Email: ryan.carlson@yale.edu or mj.crockett@princeton.edu

Abstract

Over a century of theorizing in psychology argues that humans are unable or unwilling to accurately report their motives. Yet such views are often assumed, and less often tested. Here, we present a method for systematically comparing people's subjective reports of their motives to the motives revealed through their actual choices. Using this approach in five experiments (N=2191), we find robust evidence that people can accurately report their motives. People's subjective reports are consistently highly correlated with the preferences revealed through their prosocial and consumer choices. People show high introspective accuracy both before and after choice, their level of accuracy exceeds that of independent 'observer' participants, and they remain accurate even when we experimentally shift their social preferences. Together, this work engages with a longstanding debate spanning psychology, cognitive science, and behavioral economics, and offers a novel method for charting the boundaries of peoples' capacity to know themselves.

Keywords: introspection, motives, subjectivity, self-report, inference

Introduction

Human thought and behavior are shaped by a wide array of motives—ranging from hunger and thirst, to wanting to help others and acquire symbols of social status (1–4). Understanding human motives is vital for building interventions to improve people's health and social relationships. Thus, a central question facing social and behavioral scientists is how to reliably probe motives—the psychological states comprising people's wants, desires, and goals.

Since motives are unobservable, one way to measure them is to ask people to introspect on their motives—to look inward at the forces currently guiding them—and directly report them. Alternatively, scientists can measure motives indirectly through behavior or brain activity. Introspection has intuitive appeal, as it often seems easy to access and report one's inner states (5–7). Moreover, some researchers embrace self-report as a method, including people's subjective reports on their conscious thoughts and feelings. For instance, psychologists have recently used subjective reports to systematically study mind-wandering (8), task immersion (9), and subjective confidence (10). Additionally, in many areas of psychology, such as affective science (11–14) and psychopathology (15), subjective reports are viewed as indispensable.

However, scientists tend to prefer indirect approaches when it comes to measuring motives—a preference rooted in over a century of theorizing that humans cannot be trusted to report their motives (16–27). In social psychology, the idea that humans are either unable or unwilling to report on motives echoes through modern research (1, 27–39). Cognitive scientists and developmental psychologists similarly express doubt about people's ability to introspect on their motives and desires (40–44). Moreover, in neighboring fields such as behavioral economics, sociology, and behavioral science, subjective reports on preferences tend to be viewed as "cheap talk" (45, 46) or "flimsy improvisations" (47)—not to be taken at face value. As such, it is unsurprising that recent decades have seen a surge of interest in advancing indirect measures of human motives—such as neuroimaging (19, 48–50), implicit methods (32, 51, 52), and formal models of behavior (53, 54).

Why are subjective reports on motives the target of so much skepticism? Two key classes of concerns are often levied against people's ability to introspect on motives.

The first class of concerns relate to subjective reports being *inaccurate*. Such concerns stem from proposals that humans have little to no introspective access to motives. Some suggest motives largely reside out of reach in the unconscious (18, 27). Others deny that mental states such as motives exist altogether (47). Under such views, when people report their motives, they are not actually introspecting on them directly. Rather, they are responding much the same way an outside observer would—by *inferring* their motives from other sources of data (23, 27, 42, 55, 56). To this end, people may draw on their shared cultural theories (i.e., which motives a person typically has in the situation; 23, 57), self-theories (i.e., which motives uniquely drive them as an individual; 58, 59), or self-observation (i.e., which motives best explain their own past behavior in such situations; 55, 60–62). Sometimes these inferences will yield accurate reports about peoples' motives (including one's own); but such theories are also often mistaken (63). Thus, this self-inference account gives rise to two assumptions about introspective accuracy that persist in the social and behavioral sciences. The first is the *low accuracy* assumption, which asserts that people's subjective reports on motives are generally *inaccurate*—showing only a low correlation with their actual behavior. The second is the *self-other symmetry* assumption—which asserts that since people's introspective accuracy rests largely on inferences from observable information or causal theories, however high or low their accuracy may be in a given domain, it will tend to be roughly symmetrical with, and thus no more accurate than, the inference of an informed observer.

The second class of concerns state that subjective reports of motives are not just inaccurate but tend to be distorted in specific ways. That is, subjective reports are *biased*. Some motives may threaten a person's self-image—researchers often note—and those motives, however dominant in guiding behavior, will be faint or entirely absent in subjective reports. In particular, researchers suggest that even if humans were aware of and could accurately report

their true motives, they would often still choose to obscure or conceal them (64). This concern traces to work finding that people tend to reach conclusions that cast themselves in a favorable light (65–69), as well as evidence that people's subjective reports are shaped by social desirability (70) and self-presentational (71) pressures. Concerns of bias are especially potent in moral settings, where a person's motives speak to their broader moral character (72–74). In such settings, people may be especially likely to conceal motives that violate their personal moral standards—allowing them to preserve their moral self-image (75–78). Accordingly, indirect measures—such as economic games and neuroimaging—are common for measuring prosocial motivation (19, 48, 50, 79–82), for instance.

In modern research, concerns of inaccuracy and bias tend to be stated in qualitative, rather than quantitative, terms: researchers often simply claim motive introspection is untrustworthy as it is inaccurate and/or biased, nodding to classic research (e.g., 23). In turn, these claims have led many researchers to prioritize indirect measures of motives. Here, we pursue a different direction. We believe the classic empirical evidence against accurate motive introspection—rather than serving as grounds for methodologically side-stepping introspective approaches—presents a vital opportunity for more precise and systematic investigations into people's introspective capacity. It is both theoretically and practically important for scientists to understand the conditions which produce inaccuracy and bias in motive introspection, the degree of such inaccuracy and bias, and why it emerges. For instance, many classic demonstrations of introspective inaccuracy probe introspection through peoples' verbally-reported 'reasons' for their past behavior (23, 83), which are difficult to directly compare to behavior, and might promote reliance on non-introspective information, such as cultural scripts or norms. Towards these ends, we developed an approach to quantify introspective accuracy and bias and leveraged this approach to test (a) the extent of accuracy and bias in people's subjective reports on their motives, and (b) whether their accuracy reflects reliance on self-inference versus introspection.

Our approach treats motives as dynamic states that actively guide cognition and behavior (2, 4, 84). Thus, motives are related to but distinct from people's post-hoc explanations of past behavior (i.e., reasons) or their behavioral tendencies across time (i.e., traits and dispositions; 84). Since introspection involves attending to one's current internal states, and motives can change in a given situation, introspection should only reveal motives that are actively guiding one's thoughts and behavior, and not those which have been absent for hours or days (86).

Guided by this view, here we used computational modeling to quantify introspective accuracy and bias. We focus on two domains of central interest to social and behavioral scientists, and which vary in moral relevance: personal (consumer) decisions and prosocial decisions. Specifically, in each domain, we examined how well people's subjective reports of their motives relate to the motives revealed through their choices (their revealed preferences). Many researchers have used similar approaches in multi-attribute choice settings, and generally concluded that introspective accuracy into preferences is low (59, 87–91). However, in choice settings featuring a wide array of attributes, it is unclear whether inaccuracy reflects people's limited capacity for introspection, or imprecision in modeling people's preferences from complex choices. Here, by contrast, we model choices within a simple, economic game framework where preferences are straightforward to capture as motives directly map onto specific behaviors (e.g., 80, 91, 92). By measuring how an individual's subjective reports deviate from the preferences directly reflected in their behavior, we can illuminate the 'best case scenario' for people's introspective accuracy. Moreover, by testing how such deviations shift across different domains (e.g., when the motives in question are morally-relevant or not), we can better understand when and how introspection into motives might become more biased and untethered from behavior. As such, we propose that computational models of simple choice behaviors, when used in conjunction with directly analogous subjective reports, can serve as a powerful tool for quantifying introspective accuracy and bias.

Across five experiments, we examined motive introspection in a simplified setting—reporting on the relative strength of two motives that map directly onto two different behavioral options. This allowed us to test how accurate motive introspection can be when the objects of people's subjective reports are clear and constrained and afforded us precise estimates of how people's subjective reports tracked with their behavior. In Experiments 1a and 1b, participants repeatedly chose between money for themselves on one of two different gift cards (personal domain), and between money for themselves or a charity (prosocial domain), and subjectively reported how much they wanted each reward relative to the other (see Methods). To quantify introspective accuracy, we examined how well people's subjective reports of motives tracked with motives revealed by formally modeling their actual choices (53, 54, 94). To quantify introspective bias, in Experiments 1a and 1b, we tested for motivated errors in subjective reports in the prosocial domain (where researchers frequently raise concern of bias) versus a strictly personal domain (where researchers raise comparatively less concern of bias). To test the contributions of introspection versus inference to the accuracy of subjective reports, we examined whether people's accuracy remains significant after controlling for non-introspective sources of motive information (e.g., people's prior beliefs about motives and behavior). In Experiment 2, we directly tested for an 'insider advantage' in introspective accuracy by having paired observers infer each decider's motives before and after observing their unique choices and compared the accuracy of deciders to observers. Finally, in Experiments 3a and 3b, we evaluated the reliability of people's introspective accuracy by examining whether their reports remain accurate after we experimentally shift motives—thus testing whether subjective reports dynamically update based on active motivational shifts which cannot be inferred from their past behavior in the task.

Results

Across five experiments, we compared participant's subjectively reported motives to the motives revealed by modeling their actual behavior (see Figure 1A-C). In each domain, we estimated the relative strength of participants' motives (or preferences) based on a set of 51 choices using a model which is widely adopted for binary choices (94, 95; see Methods). Across all experiments, the model provided a very strong fit to people's behavior—yielding revealed preference estimates which correlated almost perfectly with the proportion of choices people made for one reward option over the other option ($r_s > .96$ across domains and experiments; see Methods for model details, and Supplementary Results for model validations).
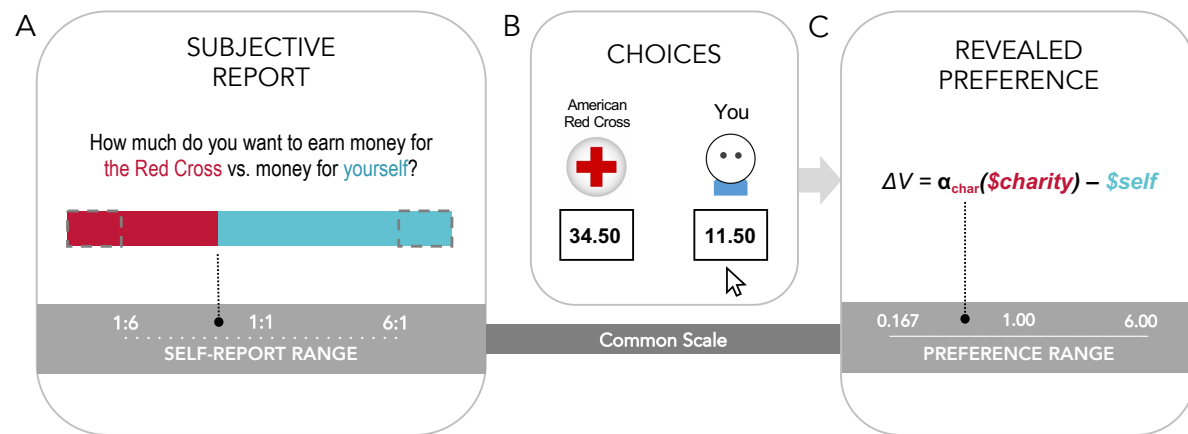

Figure 1. Illustration of key task measures within the prosocial domain. **(A)** Subjective reports were made with a slider that shifted the proportion of two colors reflecting the relative strength of one's motives for the two options. **(B)** Participants made 102 unique choices across two domains (prosocial or personal) by clicking on one of the two varying monetary options. **(C)** Revealed preference estimates (α) were derived from choices in the task.

In all experiments, participants made subjective reports about the relative strength of their motives for each reward immediately before and after making choices (in Experiments 1a, 1b, and 2, only half made reports prior to choice; see Methods). Subjective reports of motives were reported on a colored visual scale, which allowed them to express the relative strength of their motives intuitively (see Figure 1A). This scale was designed to share a common scale with estimates of revealed preferences (see Figure 1C and Methods for further details), while removing the potential influence of numeric anchors that could afford more direct comparisons

to choice behavior. As such, this approach allowed us to directly compare subjective reports against revealed preferences. Since subjective reports and revealed preferences were both nonlinear, all key inferential tests use non-parametric statistics.

Our key test of introspective accuracy was the extent to which people's subjective reports of their motives tracked with the motives revealed by modeling their actual choices. We examine the accuracy of subjective reports made after the task (post-choice reports) as well as before the task (pre-choice reports). Our approach entails comparing participants' subjective reports of motives with the corresponding choice parameter extracted from the model of their full set of choices. That is, we are *not* simply asking participants to subjectively report back examples of choices they made and comparing those reports with the choices they actually made. Rather, we are measuring in a common space participants' overall subjective perception of the strength of their preference between the two choice options, and the computational model's estimate of the strength of their preference between the two choice options based on their full set of choices. If subjective reports of motives accurately reflect the motives that guide participants' behavior in the task, then these subjective reports should be strongly correlated with revealed preferences.

The accuracy of subjective reports (Experiment 1a & 1b)

In Experiment 1a (*N* = 293) and 1b (a pre-registered replication; *N* = 499), participants chose between varying amounts of money for themselves on either a giftcard for Grubhub or Home Depot (personal domain) and chose between varying amounts of money for themselves or a charity (prosocial domain).

One key assumption arising from past work is that people exhibit low introspective accuracy. That is, people's subjective reports should tend to show at best a weak correlation with their behavior (e.g., *r* = .33; see 59, 81, 96–98). In Experiments 1a and 1b, we tested this possibility by examining the correspondence between people's post-choice subjective reports and their revealed preferences. Specifically, we used rank correlations to assess the strength of

the monotonic relationship between people's post-choice subjective reports and revealed

preferences in the personal domain and prosocial domain.

        We first examined accuracy in the personal choice domain. Contrary to the low accuracy

assumption, when we directly compared people's revealed preferences to their analogous

subjective reports, we found that post-choice reports were strongly correlated with revealed

preferences (Exp. 1a: r_s = .77, p < .001; Exp. 1b: r_s = .78, p < .001; Figure 2A). Next, we

examined accuracy in the prosocial domain, where researchers tend to be more suspicious of

people's accuracy due to self-image concerns. When we compared people's post-choice

subjective reports to their revealed social preferences, we found a similarly strong correlation

(Exp. 1a: r_s = .86, p < .001; Exp. 1b: r_s = .81, p < .001, Figure 2B). Moreover, when we directly

compared these correlations, we found that accuracy in the prosocial domain was not

consistently different from the personal domain (Exp. 1a: z = 3.08, p = .002, Exp. 1b: z = 1.14, p
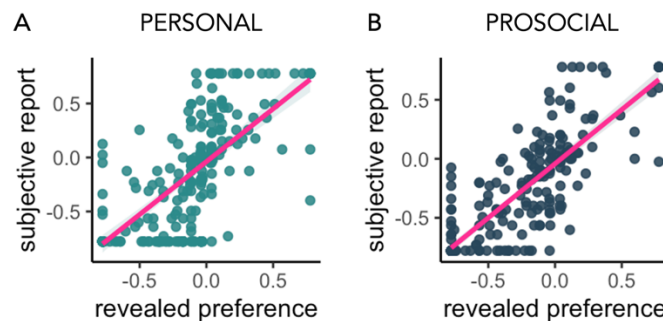
= .25).

Figure 2. Illustration of the relationship between post-choice subjective reports and the motives
revealed through people's choices in Experiment 1a in the prosocial domain **(A)** and personal
domain **(B)**. A LOESS-smoothed curve is fit to show the monotonic relationship between
subjective reports and revealed preferences.

        These results offer preliminary support for the idea that subjective reports of motives can

accurately reflect the motives revealed in people's actual choices. In a simple choice context,

when we compare subjective reports directly with analogous revealed preferences, we observe

very strong correspondence. Although, peoples' accuracy is still below the theoretical 'best

case' for accuracy when accounting for noise in their choice processes (see Supplementary

Results).

One challenge in interpreting these findings as evidence of introspective accuracy is that

our task might simply allow participants to infer their own motives by observing their own

behavior—much the same way an outside observer could. In other words, accurate subjective

reports in such a simple choice setting might not truly reflect introspective accuracy, but rather

an *illusion* of introspective accuracy arising from accurate inferences from observing one's own

choices. To address this issue, Experiment 2 tested how different sources of self-knowledge,

particularly introspection into current states versus inference from self-observation, contribute to

the accuracy of people's subjective reports.

Introspective versus non-introspective sources of accuracy (Experiments 1a, 1b and 2)

Even if an individual accurately reports their motives, it is nonetheless possible that they

did so by inferring those motives through observing their own choice behavior (55)—a feat that

informed observers could accomplish just as well (the self-other symmetry assumption). This

concern is especially prominent when individuals report their motives post-choice. To test

whether individuals' accuracy is merely achieved through self-observation and not introspection,

we compared the accuracy of their subjective reports to that of (i) a separate group of 'observer'

participants, and (ii) individuals' pre-choice reports.

Accuracy from social observation. In Experiment 2, we directly tested the claim that

people infer their own motives from observed covariation in their past behavior. Specifically, we

recruited a group of observers, each of whom was tasked with carefully observing and

confirming the choices of a participant, a 'decider', from Experiment 1b (see Methods; Figure 3).

This approach yielded 443 observers who were paired with a unique decider. We then

compared the accuracy of deciders' subjective reports to inferences that observers made about

the deciders' motives from observing their choices in the task. If people draw on reliable

introspective data to report their motives, we should see an 'insider advantage', such that

deciders' reports are more accurate than observers. By contrast, if deciders are simply inferring their motives from their own choices (23, 55), then an observer given the same information about decider's choices should be just as accurate in reporting their motives as deciders.
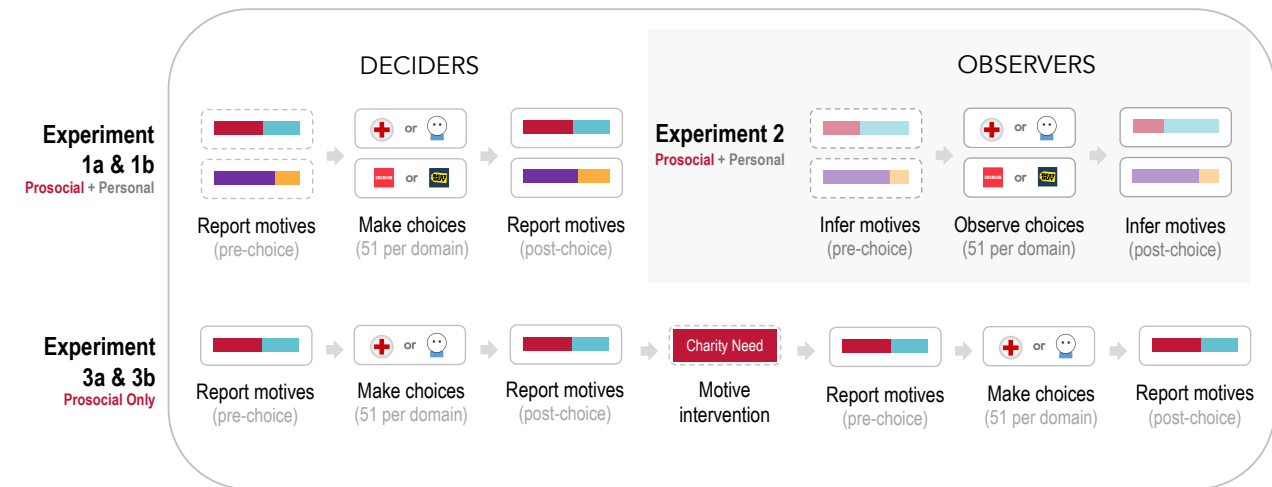

Figure 3. Overview of the task designs for Experiments 1a and 1b (which included both personal and prosocial choices), Experiment 2 (which involved observing and confirming choices made by others), and Experiments 3a and 3b (prosocial choices only–with a motive intervention midway through the task). In Experiments 1a and 1b, half of participants were randomly assigned to make (vs. not make) pre-choice subjective reports, and in Experiments 3a and 3b, half of participants were randomly assigned to the motive intervention condition (vs. a no intervention control condition).

Our results were contrary to the self-other symmetry assumption. We found that deciders were significantly more accurate at reporting the motives reflected in their choices than observers were in the personal domain ($r_{deciders}$ = .78, $r_{observers}$ = . 67, z-test: z = 3.04, p = .002) as well as in the prosocial domain ($r_{deciders}$ = .81, $r_{observers}$ = .71, z-test: z = 3.12, p = .002; see Figure 4A). Moreover, in line with this asymmetry in accuracy, we found that deciders' subjective reports and observers' inferences were only moderately correlated in either domain ($r_{personal}$ = .54, p < .001; $r_{prosocial}$ = .63, p < .001; see Figure 4A), consistent with the idea that deciders drew on distinct information to form their subjective reports.
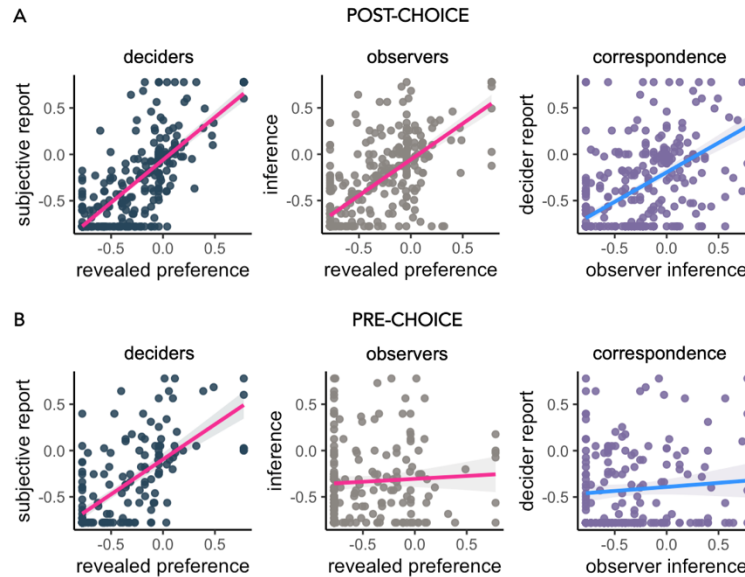
Figure 4. Comparison of how revealed preferences relate to deciders' subjective reports versus observers' inferences in the prosocial domain, as well as the correspondence between deciders' reports and observers' inferences. Even after observing deciders' choices, observers do not match deciders' level of accuracy (A). Before observing deciders' choices, observers' reports show zero correlation with deciders' behavior (whereas deciders' remain predictive; B).

These findings suggest that observers can accurately infer deciders' motives after watching their choices in simple choices settings. At the same time, our findings also suggest that people's subjective reports are more accurate than, and distinct from, inferences made by observers. This is consistent with the idea that peoples' subjective reports draw on privileged, introspective data above and beyond that which can be discerned from observing their behavior alone—even in simple choice settings.

Accuracy from self-observation. Another way to examine the difference between decider and observer reports is to examine accuracy prior to choice. If deciders and observers are both relying on behavioral information in the task, both groups should exhibit much lower accuracy prior to observing any choices. By contrast, if people are relying on introspective data in their subjective reports, their reports should be accurate even before making any choices, whereas only observers should be far less accurate. We next tested this possibility.

In Experiments 1b and 2, half of the participants additionally reported their motives *prior* to making any choices—that is, they made pre-choice subjective reports (see Methods). We confirmed that making pre-choice subjective reports did not in itself alter peoples' choices on key model parameters (i.e., revealed preferences or choice stochasticity; see Supplementary Results). As such, examining pre-choice reports allowed us to test how much behavioral observation (whether observing oneself or observing another) contributes to the accuracy of subjective reports.

To this end, we examined how accurate deciders and observers were when making pre-choice relative to post-choice reports across both domains. We found that participants' pre-choice subjective motive reports were still highly accurate, exhibiting strong correspondence with estimates of revealed preferences in both the personal domain (Exp. 1a: r_s = .71, p < .001; Exp. 1b: r_s = .69, p < .001) and the prosocial domain (Exp. 1a: r_s = .77, p < .001; Exp. 1b: r_s = .78, p < .001; see Figure 4B). By contrast, observers' inferences about deciders' motives prior to observing their choices showed zero correlation with deciders' preferences in both the personal domain (r_s = .02, p = .74) and the prosocial domain (r_s = .05, p = .52; see Figure 4B). In addition, deciders' subjective reports and observers' inferences showed no correspondence to each other in either domain ($r_{personal}$ = - .02, p = .83; $r_{prosocial}$ = .04, p = 63).

These findings show that deciders' accuracy draws on sources of self-knowledge beyond behavioral data from the present situation. Indeed, we find similar levels of pre-choice introspective accuracy irrespective of whether deciders had prior experience in other experiments with making choices in these domains (see Supplementary Results).

Ruling out multiple non-introspective sources of accuracy. The results thus far provide initial support for the idea that people meaningfully draw on introspective data when reporting their own motives, which translates into greater accuracy in their subjective reports. If this is indeed the case, then people's subjective reports should be a robust predictor of their revealed preferences, above and beyond non-introspective sources of data, including not just

behavioral information from observers, but also peoples' own memory of their behavior (76, 100, 101), and their shared cultural knowledge (i.e., their intuitive theories of motivation; 59, 98–100).

To formally test this, we compared a model predicting peoples' revealed preferences in each domain based on these strictly non-introspective sources of knowledge about deciders. The first two sources of knowledge we tested in both domains were behavioral: (i) observers inferences of deciders' motives based on their observed choices, and (ii) deciders' own memories of their choice behavior recalled after the task (see Methods). A third source of knowledge, which we tested in the prosocial domain, was relevant cultural knowledge that might guide one's reports: (iii) deciders' beliefs about what the average person's motives would be in each choice domain (see Methods). We compared models in each domain containing these non-introspective sources of knowledge against a model which included these non-introspective predictors as well as people's post-choice subjective reports. Consistent with the idea that introspection makes a unique and meaningful contribution to predicting people's revealed preferences, we found that in both domains, the model which additionally included peoples' subjective reports was superior to a simpler model that relied only on non-introspective sources of data to predict peoples' preferences ($\Delta AIC_{prosocial}$ = 91.30, ΔBIC = 87.44, likelihood-ratio test: χ^2 (1) = 93.30, $p < 0.001$; $\Delta AIC_{personal}$ = 54.07, ΔBIC = 50.21, likelihood-ratio test: $\chi^2(1)$ = 56.07, $p < 0.001$; see Supplementary Tables 1 & 2). Taken together, these findings further support the idea that people's reliance on introspection distinctly and meaningfully contributes to the accuracy of their subjective reports.

Introspective error and bias as a function of violating personal standards

Above, we presented evidence highlighting the potential accuracy of subjective reports, and how the accuracy of such reports is undergirded by an introspective process. However, observing a strong correspondence between peoples' subjective reports and their revealed preferences does not rule out the possibility that those reports contained self-serving errors and biases (see 104). Indeed, a growing literature suggests that people are prone to errors and

distortions when recalling events that threaten their moral self-image (76, 100, 106, 107). To see whether such phenomena influence introspection as well, we next examined motivated biases in subjective reports. Specifically, we tested whether individuals whose motives pose a threat to their self-image might (i) exhibit greater introspective errors in general, and (ii) show a directional bias towards self-serving introspective errors in particular—such that they report stronger prosocial motivation than their revealed preferences imply.

To discern whether peoples' choices posed a threat to their self-image, we probed participants' beliefs about what was a fair amount to want to earn money for themselves versus earn money for charity in the task (see Methods). We reasoned that people whose actual choices fell short of their own fairness standards should exhibit more introspective error and bias than those whose behavior upheld their standards. To this end, in Experiments 1a, 1b, 3a, and 3b, we separately examined introspective errors in participants whose revealed social preferences upheld their own fairness standard (i.e., upholders, N_{exp1a} = 151, N_{exp1b} = 243, N_{exp3a} = 134, N_{exp3b} = 286) and those whose revealed preferences violated their own fairness standard (i.e., violators, N_{exp1a} = 112, N_{exp1b} = 200, N_{exp3a} = 135, N_{exp3b} = 314; an approach adapted from (76); see Methods). Importantly, the revealed preferences of these two groups were relatively well matched (W_{exp1a} = 7613, p = .167, d = 0.26, δ = 0.10; W_{exp1b} = 21288, p = 0.025, d = 0.25 δ = 0.12; W_{exp3a} = 8984.5, p = 0.94, d = 0.34, δ = 0.01; W_{exp3b} = 31313, p = 0.43, d = 0.28, δ = 0.04, Mann-Whitney U-tests, means and standard deviations are reported in Supplementary Table 3), suggesting that any differences between these groups should not be due to their choices alone, but whether or not those choices were subjectively at odds with their personal standards.

Introspective error. We predicted that fairness violators would tend to make larger introspective errors, exhibiting greater deviations in their reported prosocial motivation than fairness upholders. To test this hypothesis, we examined introspective errors (the absolute size of the discrepancy between subjective reports and revealed preferences) in violators versus

upholders within the prosocial domain. Consistent with our predictions, we found that violators tended to make significantly greater introspective errors than upholders (W_{exp1a} = 5588.5, p < .001, d = 0.49, δ = 0.34; W_{exp1b} = 16931.5, p < .001, d = 0.35, δ = 0.30; W_{exp3a} = 7972.5, p = 0.095, d = 0.17, δ = 0.12; W_{exp3b} = 24538, p < .001, d = 0.26, δ = 0.25, Mann-Whitney U-tests, means and standard deviations are reported in Supplementary Table 4; see Figure 5A). Next, we examined whether violators and upholders showed the same introspective error tendencies in the personal choice domain. Here we expected to see no difference between the groups as violators should not be motivated to be less precise. Consistent with our predictions, when we examined these two groups in the personal domain, we saw no difference in the size of their errors (W_{exp1a} = 8746.5, p = 0.63, d = 0.07, δ = 0.03; W_{exp1b} = 23876, p = 0.75, d = 0.02, δ = 0.02; Figure 5A).
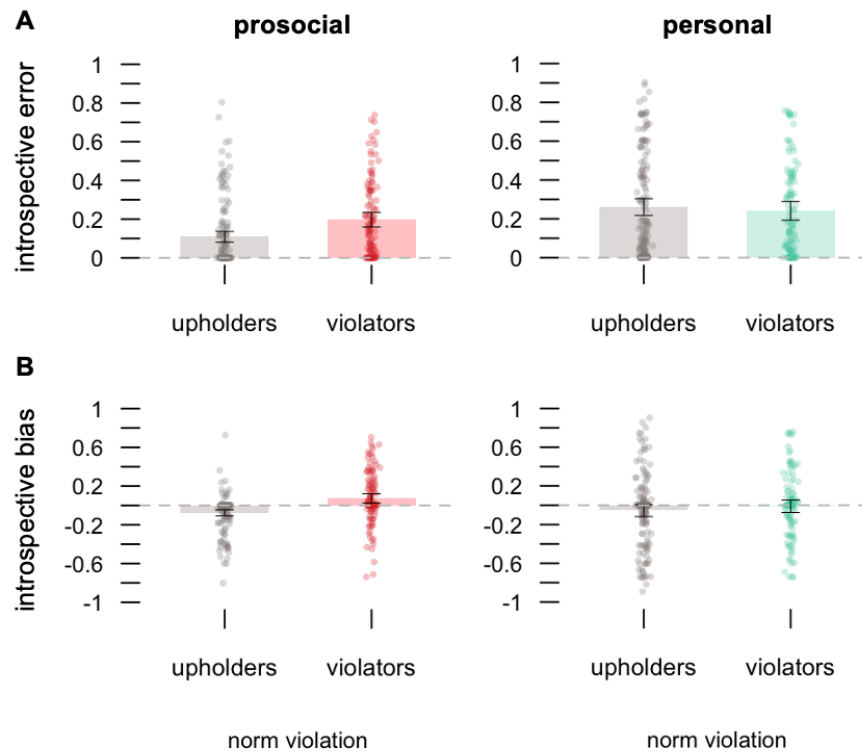

Figure 5. Here we plot discrepancies between participants subjectively-reported motive strengths, versus their revealed preferences as a function of choice domain and whether people's motives violated their fairness standard in Experiment 1a. Panel A visualizes introspective errors (the absolute size of participants' errors), and Panel B visualizes introspective bias (the size and direction of their errors), such that positive values reflect "self-

serving" introspective errors, and negative values reflect "self-deprecating" introspective errors. Error bars represent 95% confidence intervals (CIs).

Introspective bias. We next examined introspective bias among participants in the prosocial domain. We quantified introspective bias by examining the *direction* of participants' introspective errors (subtracting their subjective reports from their revealed preferences). On this measure, positive values reflected self-serving errors (reporting stronger prosocial motivation than one's revealed preferences imply), whereas negative values reflect self-deprecating errors (reporting weaker prosocial motivation than one's revealed preferences imply). In all experiments, we found that participants generally showed a systematic bias not toward self-serving errors, but self-deprecating errors. That is, the reported strength of their motives to give to charity were significantly lower than their choices suggested (V_{exp1a} = 13850, p = .004, d = .05, δ = -0.38; V_{exp1b} = 36020, *p* < .001, *d* = 0.07, δ = -0.44; V_{exp3a} = 15167, *p* = 0.019, *d* = 0.06, δ = -0.32; V_{exp3b} = 56593, *p* = 0.008, *d* = 0.05, δ = -0.29, Wilcoxon signed-rank test; means and standard deviations are reported in Supplementary Table 5). To test whether this bias changed as a function of whether one violated their standards, we separately assessed errors in violators versus upholders. As predicted, violators showed a systematic bias towards self-serving introspective errors, reporting being significantly more prosocially motivated than they actually were (V_{exp1a} = 4074, *p* = 0.008, *d* = 0.27, δ = 0.05; V_{exp1b} = 11938, *p* = 0.021, *d* = 0.23, δ = -0.06; V_{exp3a} = 5829, *p* = 0.007, *d* = 0.24, δ = 0.05; V_{exp3a} = 23463, *p* < .001, *d* = 0.27, δ = 0.15; see Figure 5B). Upholders, by contrast, made errors in the opposite direction in the prosocial domain, showing a bias towards self-deprecating errors (V_{exp1a} = 1899, *p* < .001, *d* = 0.39, δ = -0.70; V_{exp1b} = 4223, *p* < .001, *d* = 0.42, δ = -0.76; V_{exp3a} = 1693, *p* < .001, *d* = 0.46, δ = -0.69; V_{exp3b} = 4368, *p* < .001, *d* = 0.48, δ = -0.77). However, when we examined whether violators and upholders made similar deviations in the personal domain in Experiment 1a and 1b, we did not observe a consistently significant tendency to make introspective errors in either direction (violators: V_{exp1a} = 3002, *p* = 0.64, *d* = 0.04, δ = -0.18, V_{exp1b} = 7075, *p* < .001, *d* =

0.24, δ = -0.25 ; upholders: V_{exp1a} = 4691, p = 0.052, d = 0.14, δ = -0.21, V_{exp1b} = 12258, p =

0.019, d = 0.11, δ = -0.23; see Figure 5B). Since revealed preferences were generally no

different between violators and upholders, these differences cannot simply be explained by

'regression to the mean', which would predict that both groups should make errors in the same

direction.

   Taken together, these findings support the idea that people whose motives violate their

fairness standards show lower introspective precision in general and show a bias towards self-

serving introspective errors in particular. However, they also suggest people who uphold their

own fairness standards may exhibit a bias too, but in the other direction (a self-deprecating

bias).

   We next sought additional evidence that the introspective errors and biases we observed

are in fact motivated (i.e., driven by a desire to preserve one's moral self-image), and not a non-

motivated, cognitive bias simply arising from due to the types of choices that violators and

upholders made. To arbitrate between a motivated and a non-motivated account of bias, we

examined whether observers (Experiment 2) would make the same errors after observing

deciders' choices. If individuals merely observing the decisions of violators and upholders tend

to make the same errors about deciders' motives as deciders themselves, then we cannot

conclude that this bias is motivated. On the other hand, finding that only deciders exhibit this

bias would be consistent with a motivated account of introspective bias.

   We first examined the size of introspective errors in observers within the prosocial

domain. We found that observers of violators also made significantly greater introspective errors

than observers of upholders (W = 12929.5, p = 0.012, d = 0.13, δ = 0.16; Wilcoxon signed-rank

test, means and standard deviations are reported in Supplementary Table 4; see

Supplementary Figure 1A), however the effect size was smaller than that seen in paired

deciders (d = 0.31, δ = 0.29). Next we examined directional introspective errors. Consistent with

our hypothesis that the introspective biases we observed are motivated, we found no evidence

of a bias towards overreporting motives among observers of violators ($V = 6778$, $p = 0.14$, $d =$

0.18, $\delta = -0.12$), no evidence of a self-deprecating bias among observers who witnessed

upholders ($V = 9908$, $p = 0.96$, $d = 0.09$, $\delta = -0.26$), and no difference in the reports made by

observers who were witnessing violators versus upholders ($W = 15065.5$, $p = 0.79$, $d = 0.10$, $\delta =$

0.02; Mann-Whitney U test, means and standard deviations are reported in Supplementary

Table 5; see Supplementary Figure 1B).

Taken together, these results suggest that while individuals who violate their moral

standards might exhibit a self-serving bias—overestimating their motives to give to charity

relative to what their behavior suggests, it appears that those who uphold their standards also

exhibit a self-deprecating bias—underestimating their motives to give to charity relative to their

behavior. Moreover, we find that this bias is only seen in the prosocial domain, and observers

do not exhibit this same bias, alluding to its motivated nature. Importantly, these findings

suggest a more nuanced picture of introspection, whereby some individuals may exhibit

motivated biases in introspection, while overall people still demonstrate a remarkable degree of

introspective accuracy. This suggests that these biases, rather than being grounds for

dismissing introspective reports, can be studied in hand with the informative data produced

through subjective reports on motives.

Shifting motives to test the reliability of introspection (Experiment 3a and 3b)

A key aim of the present research has been to test the contributions of introspection

versus inference to the accuracy of subjective reports on motives, and chiefly, the influential

idea in psychology that people's introspective accuracy is achieved by inferring their motives

from their past behavior. Experiment 1b and Experiment 2 together showed that introspection

contributes to the accuracy of subjective reports above and beyond knowledge gleaned from

one's observable choices. However, an individual's knowledge of their prior behavior might

extend well beyond that which can be observed in the task, and beyond that which we could

adequately measure. To rule out the possibility that people's introspective accuracy relies on

knowledge of past behavior beyond that which present observers can access, in Experiment 3a

(N = 299) and Experiment 3b (N = 600; a pre-registered replication), we *intervened* on people's

social preferences.

In Experiments 3a and 3b, participant's exclusively made choices in the prosocial

domain (choosing between money for themselves or the American Red Cross), and at the

midpoint of the task, half of the participants received a motive intervention (motive intervention

group), and half did not (control group). We conducted each experiment during an active

humanitarian crisis: the 2021 Texas Winter Storm (Experiment 3a), and the 2022 Invasion of

Ukraine (Experiment 3b). To boost peoples' social preferences, at the midpoint of the task, the

motive intervention group received information about how the Red Cross was actively

supporting victims of each crisis (see Methods for details), whereas the control group received

no information. This intervention proved effective: the intervention group's choices expressed

significantly stronger social preferences after the intervention (preference change: V_{exp3a} = 3632,

$p < .001$, $d = 0.41$, $\delta = 0.23$; V_{exp3b} = 17431, $p < .001$, $d = 0.47$, $\delta = 0.39$; V_{exp3a} = 15167, $p =$

0.019, $d = 0.06$, $\delta = -0.32$; Wilcoxon signed-rank tests; means and standard deviations are

reported in Supplementary Table 6), whereas the control group did not (preference change:

V_{exp3a} = 2602, $p = 0.68$, $d = 0.11$, $\delta = 0.02$; V_{exp3b} = 10535, $p = 0.33$, $d = 0.05$, $\delta = 0.03$).

If people draw on knowledge of their prior behavior to infer their motives in the moment,

the accuracy of their subjective reports should decrease when their prior behavior is no longer

an informative source of data. By contrast, if people's accuracy stems from introspection, their

accuracy should remain even after the strength of their motives change. Consistent with the

latter account, we found that subjective reports among the motive intervention group were

strongly correlated with their revealed preferences even after we experimentally shifted people's

preferences in Experiment 3a and 3b (Exp. 3a: r_s = .86, $p < .001$; Exp. 3b: r_s = .79, $p < .001$),

and were no less accurate than their reports prior to the motive intervention (Exp. 3a: r_s = .82, p

$< .001$; Exp. 3b: r_s = .82, $p < .001$; p_{z-test}=.34; see Supplementary Figure 2). Moreover, the

accuracy of the motive intervention group was no different from individuals in the control group (Exp. 3a, p = .09; Exp. 3b: p = .64; $p_{z\text{-test}}$=.25), suggesting that people dynamically updated their subjective reports alongside changes to their revealed preferences.

Overall, the data across these five experiments offer clear evidence against two common assumptions about motive introspection. The first assumption is that people tend to exhibit low accuracy in reporting their motives, which predicts a weak correlation between subjective reports of motives and the motives revealed by their actual choices. In contrast with this assumption, across four experiments, we find that peoples' subjective reports of motives were strongly correlated (rs between .78-.86) with motives revealed by their actual behavior.

The second assumption is that people's introspective accuracy, even if high, is simply a consequence of deploying the right causal inference strategy—one which informed observers should be just as accurate at deploying. At odds with this assumption, our data are more consistent with the view that people's subjective report accuracy stems in part from conscious introspective data about internal states, as opposed to inferences from prior behavior or causal theories. People's post-choice subjective reports were significantly more accurate than those made by observers of their choices, and they remained a robust predictor of behavior even when we controlled for observers' inferences of their motives, deciders' own memories of their choice behavior, and data from their intuitive theories. Moreover, people's subjective reports remained accurate even after we shifted their social preferences, such that their past behavior was no longer an accurate guide for making inferences about their own motives.

Discussion

Social and behavioral scientists often regard subjective reports of motives with suspicion. In one of the most influential reviews on this topic, Nisbett & Wilson wrote "The accuracy of subjective reports is so poor as to suggest that any introspective access that may exist is not sufficient to produce generally correct or reliable reports" (23, p. 233). This idea has

echoed through psychological science in subsequent decades (42, 57, 60). Popular accounts concur, suggesting people have little to no direct access to their motives (27, 47, 64), casting further doubt on the fidelity of subjective reports.

In the present research, we highlight the promise of formalizing a direct, introspective approach to measuring the relative strength of motives. Across five experiments, we find consistent evidence that people display robust introspective accuracy. Subjective reports on motives are not only more accurate that classic work implies, but also exceed inferences made by outside observers of their choices. Moreover, subjective reports accurately update when preferences change. Together, these findings challenge the assumption that people exhibit low introspective accuracy, as well as influential ideas about peoples' introspective accuracy being on par with informed observers (42, 59). What is more, we see such robust accuracy despite also observing motivated biases in people's subjective reports within the prosocial domain (but not the personal domain). Individuals whose motives violate their own fairness standards tend to distort their level of prosocial motivation upward (a self-serving bias), and individuals whose motives uphold their own standards tend to distort their level of prosocial motivation downward (a self-deprecating bias). By developing a precise method for probing the accuracy and biases of motive introspection, this work challenges assumptions about its inherent unreliability— highlighting the value of utilizing subjective reports for investigating human motives.

Why do these findings depart from prior research suggesting that introspection is inaccurate? We believe three features of our task (reporting constraints, temporal constraints, and choice constraints) encouraged reliance on introspection, and a meaningful 'objective' comparison, thus maximizing people's ability to demonstrate introspective accuracy. First, we constrained subjective reports to a simple, continuous response format ("how much do you want X relative to Y?"). By contrast, research showing failures of introspection often rely on open-ended probes about one's reasons for their behavior (e.g., 27, 112), which may lead individuals to generate irrelevant responses or rely on non-introspective information (13, 86). Second, we

constrained subjective reports to mental states that were temporally proximate and thus still

introspectively accessible (active motives). Research traditions which argue that people can

reliably report their inner states tend to focus on subjective reports of current, conscious

experiences (86, 109, 110). By contrast, standard self-report trait measures, which show lower

accuracy (e.g., measures of empathic concern; 115), tend to ask people to generalize across

temporally distant mental states (e.g., how often they tend to feel concerned feelings for less

fortunate people), which obscures the connection between subjective reports and behavior (86,

110, 112). Finally, people made simple (binary) choices with real incentives, which should

maximize the likelihood that people's choices would reflect their real motives (113), and

minimize the potential for noise and imprecision in the 'objective' process of modeling choices

due to choice complexity (114, 115). In this way, our work differs from multi-attribute choice

research showing low introspective accuracy, where deciders often make (i) hypothetical

choices and/or (ii) choose between options that vary on a wide array of attributes (59, 87–89,

91). Adopting our approach, we found that subjective reports predict people's behavior with

strong quantitative precision. Although future work will need to shed further light on how to

quantify introspection in more complex choice environments, our approach highlights how

robust introspective accuracy can be when studied in simplified choice environments with real-

world outcomes. Moreover, it offers a method for testing how accuracy shifts based on different

motives and social situations.

We believe concerns of bias levied against motive introspection also warrant revisiting.

Concerns of bias surrounding subjective reports may be overstated. The empirical literature on

self-enhancement and social desirability biases suggests the ability of such biases to shift

subjective reports is often small and bounded by reality (76, 116–118). Moreover, in some

cases, these biases are only exhibited by people whose true behavior deviates from relevant

personally-held or social norms (76). The evidence presented above is consistent with both

ideas: people display introspective biases, but these biases do appear to not greatly undermine

their overall introspective accuracy. Moreover, these biases depend on how (and whether) one's motives deviate from their moral standards. As such, rather than dismissing subjective reports of motives as biased, researchers can and should systematically study when such biases emerge and to what extent they hinder predictive accuracy.

Importantly, the approach presented here could prove fruitful for deepening our understanding of the limits of human introspection, as well as how we measure the efficacy of interventions that seek to address pressing societal issues. Much work remains, but this work takes a small step towards charting the territory of where and when introspection can shed light on the workings of the human mind.

Methods

Participants: All experiments

293 participants completed Experiment 1a (153 female, 130 male, 7 non-binary or third gender, 3 did not specify; mean age = 34.1), 499 participants completed Experiment 1b (255 female, 229 male, 13 non-binary or third gender, 2 did not specify; mean age = 37.4), 500 participants completed Experiment 2 (260 female, 228 male, 8 non-binary or third gender, 4 did not specify; mean age = 39.2), 299 participants completed Experiment 3a (148 female, 141 male, 7 non-binary or third gender, 3 did not specify; mean age = 32.9), and 600 completed Experiment 3b (321 female, 261 male, 14 non-binary or third gender, 4 did not specify; mean age = 37.5). Of the 500 participants in Experiment 2, our procedure for matching observers to deciders yielded 443 unique observer-decider pairs. In all experiments, we recruited US participants from Prolific. In Experiment 3b, we recruited participants from the US, Canada, and the UK. All data and code are available at https://osf.io/y7ksj/.

We pre-registered three classes of participant exclusion criteria (Experiment 1b; https://aspredicted.org/HQ4_GQC; and Experiment 3b; https://aspredicted.org/V9R_SVP), though our results are robust to whether we include all participants in our analyses (see

Supplementary Results). First, we excluded participants who were extreme outliers (i.e., >3 standard deviations from the mean) in the size of their introspective errors ($N_{exp.1a}$ = 12; $N_{exp.1b}$ = 21; $N_{exp.3a}$ = 10, $N_{exp.3a}$ = 24), or observational errors in the case of Experiment 2 ($N_{exp.2}$ = 19). This exclusion criteria only captured participants whose motive reports were at the opposite extreme as their behavior, suggesting that they used our novel subjective report scale in the reverse direction (e.g., those who reported exclusively wanting a Grubhub gift card and then exclusively pursued a Best Buy gift card). In addition, we excluded participants who indicated suspicion or confusion about key aspects of the study across a series of open-ended and Likert-scale measures ($N_{exp.1a}$ = 11; $N_{exp.1b}$ = 26, $N_{exp.2}$ = 22; $N_{exp.3a}$ = 14, $N_{exp.3a}$ = 39; see believability and suspicion probe measures below). Lastly, we excluded participants who failed to ever move the subjective report measure from its random starting position across our key measures ($N_{exp.1a}$ = 10; $N_{exp.1b}$ = 9, $N_{exp.2}$ = 10; $N_{exp.3a}$ = 10, $N_{exp.3a}$ = 24). After these exclusions, the total N was 263 participants in Experiment 1a, 443 participants in Experiment 1b, 449 participants in Experiment 2, 269 participants in Experiment 3a, and 511 participants in Experiment 3b. To test the robustness of our results, we confirmed that our findings were consistent not only when we include all participants in our analyses, but also when we additionally exclude people who always made self-interested decisions in the prosocial domain (see Supplementary Results).

The Yale University Human Investigation Committee approved the procedures for all experiments (HSC: 2000022385). The experiments complied with all relevant ethical regulations for work with human participants and all participants provided informed consent.

Procedures: Experiments 1a and 1b

After consenting to participate, participants were informed they would be completing a choice task involving two choice domains (personal and prosocial choices) and would be using a novel scale to indicate how much they wanted one reward relative to the other in the task (subjective reports). After learning about the task and completing a tutorial on how to use the

subjective report scale, half of the participants were randomly assigned to immediately report the relative strength of their motives in each domain (pre-choice subjective reports). The other half of participants received no additional information and proceeded directly to the choice task.

In the choice task, participants made 102 choices between different types of monetary rewards of varying amounts (see Figure 1 and 3). In one block of trials (prosocial domain, 51 trials), these decisions were between earning money for themselves or a charitable organization the–American Red Cross, and in another block of trials (personal domain, 51 trials), these decisions were between earning money on a GrubHub or Best Buy gift card. Each choice trial was preceded by a fixation cross which appeared for one of the following intervals which was selected at random on each trial: 750ms, 800ms, 850ms, 900ms, 950ms, 1000ms. After the fixation, participants selected between two whole dollar amounts positioned below an icon symbolizing each reward type (we randomized across subjects which reward was positioned on left versus the right side of the screen during their choices; see Figure 1B and 1D). These dollar amounts reflected how much of each reward could be earned on that trial. The dollar amounts ranged from a proportion of 6:1 to 1:6 in value for one option over the other. Then present two icons indicating the reward type with a whole dollar amount directly below the icon. For instance, participants might encounter a choice in the prosocial domain in which the left side of the screen had the option to earn $36 for themselves, and the right side of the screen had an option to earn $6 for the Red Cross. The order of the prosocial and personal blocks was randomized across subjects as well. Moreover, we informed participants that for a subset of participants, one of their choices would be randomly selected, and thus they should treat all choices as real.

Immediately after completing the choice task, we probed the strength of participants' motives to earn each reward relative to the reward it was pitted against in the choice task. To do this in a way that is both intuitive but minimizes numeric anchors which might influence subsequent choices, subjective reports of motives were made using a colored visual scale (see

Figure 1A). This scale had two colors, one reflecting each reward type (e.g., the amount of **red** reflected one's desire to earn money for charity and the amount of **blue** reflected one's desire to earn money for themselves). We randomized the color assigned to earning money for oneself (blue or beige) across subjects, but always matched the colors representing other rewards to their respective logo (e.g., American Red Cross was always red).

Participants indicated their relative motive for one reward over the other by shifting the relative proportion of the colors using two different keyboard keys. We designed this scale explicitly to share a common scale with the revealed preference parameter that could be recovered in a model of behavior in the choice task—as such, the subjective report scale had 51 points, each corresponding to a ratio reflected in the choice task. This allowed us to directly compare participant's subjective reports against their behavior.

After all participants completed the post-choice introspective report, they completed a series of additional social and demographic questions, which we highlight ahead.

Procedure: Experiments 2 (observer study)

Experiment 2 was structured similarly to Experiments 1a and 1b, except participants instead observed choices—allowing us to test a self-inference account of introspective accuracy, whereby people come to know their own motives by observing their own behavior (55). Specifically, each observer participant was asked to observe 102 choices made by a "yoked" prior participant who made decisions in Experiment 1b (see Figure 1D). Thus, these decisions likewise reflected two domains (prosocial and personal). To mirror the experience that deciders had, the experimental design (including the order of choice domains and specific choices) was identical to what deciders witnessed. To ensure observer participants paid attention during the task, we instructed them to select the same choice that the decider participant had made after that choice was highlighted in green, and the task would not advance until the observer participant selected that same choice to confirm they observed it.

Procedures: Experiments 3a and 3b

Experiment 3a and 3b were similar to Experiment 1a and 1b, with several exceptions. First, participants only made choices in the prosocial domain. Second, to test introspective accuracy when motives change, we used an intervention to shift people's charitable motives midway through the task. Lastly, we probed people's motives before and after each of the two choice blocks, allowing us to compare pre-choice and post-choice reports before and after a motivational and behavioral shift (see Figure 1D). The motive interventions varied for each experiment, however both highlighted how donations could help those affected by a current humanitarian crisis.

Experiment 3a intervention: Aid for victims of the 2021 Texas Winter Storm. In Experiment 3a, participants were either randomly assigned to a control condition in which they received no information between choice blocks, or an intervention condition aimed at boosting their motivation to earn money for charity. To shift motives in Experiment 3a, midway through the choice task, participants in the intervention condition read about the efforts of the target charity (The American Red Cross) in the then ongoing 2021 Texas winter storm and power crisis:

> Why give to the Red Cross right now? The state of Texas was hit by a historic winter storm two weeks ago. As of Monday, approximately 400,000 residents are still without clean running water. The American Red Cross is working with government and community partners to coordinate water distribution, and is providing water for people at emergency shelters. Donations made in our study that are selected will be fully donated to the Red Cross disaster relief fund, which supports people such as those suffering in Texas right now.

After reading this passage, participants made a subjective motive report and made their remaining 51 choices between money for themselves or the American Red Cross.

Experiment 3b intervention: Aid for victims of the 2022 Invasion of Ukraine. To shift motives in Experiment 3b, we instead highlighted the efforts of the target charity in the 2022 war in Ukraine mid-way through the task:

Why give to the Red Cross right now? Nearly 12 million people have been displaced by the ongoing conflict in Ukraine. The Red Cross has reached 1 in 10 people impacted by this crisis and is committed to assisting the most vulnerable. The American Red Cross has sent 10 million dollars (and committed an additional 32 million dollars) to Ukraine crisis relief efforts which address vital humanitarian needs of those affected, including food, shelter, critical care items and first aid. Donations made in our study that are selected will be fully donated to the Red Cross disaster relief fund, which supports people such as those displaced in Eastern Europe right now.

After reading this passage, participants again made a subjective motive report and made their remaining 51 choices between money for themselves or the American Red Cross.

Choice model: All experiments

In all experiments, we compared participants' subjective reports of relative motive strength to a parameter capturing that same quantity in a model of their choices. Specifically, we used a utility model (95, 96) that estimates each participant's revealed preference between two types of rewards. We used identical, separate models for each choice domain (personal and prosocial), but illustrate the model with the prosocial domain:

$$\Delta V = s - \alpha(c)$$

where ΔV is the difference in subjective value between the two options, and s and c are the objective amounts of money for self and money for charity, respectively. ΔV depends on the weighting of the parameter α, which determines how valuable the gains for charity are relative to gains for self—thus acting as an *exchange rate* for peoples' preference for the two reward types. Trial-by-trial value differences were then transformed into choice probabilities using a softmax function:

$$P(choose\ alternative) = \left(\frac{1}{1 + e^{-\beta \Delta V}} \right)$$

where γ is a subject-specific inverse temperature parameter that captures how sensitive their choices are to ΔV. We used participants' 51 choices from each choice block in the task to estimate their most likely exchange rate (α) and inverse temperature (β) for a given choice

domain using a nonlinear maximum likelihood estimation function in MATLAB. Crucially, we designed the choice trials and subjective report measures to allow for transforming these a parameters into the same continuous space as participants' subjectively-reported motives. Moreover, this model provided a strong fit for people's behavior across all experiments (see Supplementary Methods for further model details and validations). Since α reflected a ratio, this parameter asymmetrically scales for one option versus the other—thus we log10 transformed α parameters for analysis and visualizations reported in the Results section. Due to this transformation, we refer to exchange rate parameters (α) throughout the manuscript as estimates of people's 'revealed preferences', or the motives revealed through their choices. Due to space limitations, we examine the relationship between introspective accuracy and inverse temperature (or choice stochasticity) in the Supplementary Results.

Other measures: All experiments

Subjective confidence. Each time participants reported their motives, we also asked them: "How confident are you that you accurately reported what you want" on a 7-point Likert scale ranging from 1 (Not at all) to 7 (Extremely).

After completing their final post-choice subjective report and confidence rating, participants responded to the following questions in the order they appear. Due to space restrictions, participants subjective confidence ratings and their relationship to introspective accuracy are reported in the Supplementary Materials.

"Fair" motives. To determine how personal moral violations might factor into introspective accuracy and bias, in the prosocial domain, we probed people's personal standards surrounding their motives. Specifically, we asked them to use the same colored subjective report scale to answer the following in each choice domain: At minimum, how much is a fair amount to want to earn money for yourself versus charity?

"Average" motives. In addition to probing people's sense of fairness surrounding levels of prosocial motivation, we also probed their beliefs about descriptive norms to measure their

intuitive theories about motives in the task. In this case, we asked people to consider the average person, and use the same colored subjective report scale to report how much they think the average person would want to earn money for themselves versus charity.

Memory for behavior in each domain. To address a separate research question, we examined how memory for one's behavior relates to their post-choice subjective motive reports. To this end, participants were presented with a surprise memory test on their choices in each choice domain. Specifically, across two questions, we asked participants to recall what percent of choices they made which were choices to earn money for the Red Cross [Grubhub]. We informed participants that they would receive a monetary bonus that depends on their accuracy, such that reporting their actual choice percent would result in 100% of the money, and for each 1% off their response was, their bonus would decrease by 2%. Responses were made by reporting a whole number between 0 and 100 reflecting the percentage in a response box.

Choice experience in each domain. To examine how a history of choice experience might guide one's pre-choice or post-choice subjective motive reports, participants were asked how frequently they had completed studies in which they (i) had to choose between money for themselves or a charity and (ii) had to choose between two different gift cards with one of the following 5 responses (1 = "Never", 2 = "A few", 3 = "More than a few", 4 = "Many", 5 = "Countless").

Social or personal pressure. As a general probe for potential personal or social pressure in the subjective report scale, we asked participants whether they felt any personal or social pressure when asked about what they wanted in the study on a 7-point scale from 1 (Not at all) to 7 (Very much so). We confirmed in our initial experiments that participants reported relatively low pressure to respond a certain way (Exp. 1a: M = 2.77, SD = 1.89; Exp. 1b: M = 2.71, SD = 1.87).

Demographics. Participants reported their age and gender (reported above), race (see Supplementary Table 7), as well as a series of other demographic variables that are reported separately elsewhere (Carlson & Crockett, in prep).

Believability and comprehension probes. Finally, to ensure people understood and were not suspicious about key aspects of the study, we included the following questions. First, an open-ended question: "*We asked you to report how much you wanted one option over another using colored bars. We would appreciate any feedback on your experience responding this way. Please use this space to share any thoughts about this*". Second, we asked a closed question: "*In today's study, you made choices involving various incentives including money for yourself or charity, and two gift cards for yourself. You were informed that for a subset of participants in this study, one of these choices really counts. Did we communicate the fact that you could earn real money, gift cards, or money for charity, clearly and transparently enough in our instructions?*" on a 5-point scale from 1 (No, not at all) to 5 (Yes, fully). Third, we asked another open-ended question: "If anything was unclear during the study or you had other doubts or remarks about any part of our instructions, please let us know why". Lastly, we asked a final open-ended question: "Apart from this, do you have any other remarks or comments about today's study?". We excluded people who indicated suspicion or confusion in the open-ended probes and excluded anyone who reported a "1" on the third, close-ended question about believability.

References

1.   C. D. Batson, *Altruism in humans* (Oxford University Press, USA, 2011).

2.   C. S. Dweck, From needs to goals and representations: Foundations for a unified theory of motivation, personality, and development. *Psychol. Rev.* **124**, 689–719 (2017).

3.   K. Lewin, *Principles of topological psychology* (McGraw-Hill, 1936).

4.   B. Weiner, *Human motivation: Metaphors, theories, and research* (Sage Publications, Inc, 1992).

5.   E. Pronin, "Chapter 1 The Introspection Illusion" in *Advances in Experimental Social Psychology*, (Academic Press, 2009), pp. 1–67.

6.   L. Ross, A. Ward, "Naive realism in everyday life: Implications for social conflict and misunderstanding" in *Values and Knowledge*, The Jean Piaget symposium series., (Lawrence Erlbaum Associates, Inc, 1996), pp. 103–135.

7.   S. Vazire, E. N. Carlson, Self-knowledge of personality: Do people know themselves? *Soc. Personal. Psychol. Compass* **4**, 605–620 (2010).

8.   K. Christoff, Z. C. Irving, K. C. R. Fox, R. N. Spreng, J. R. Andrews-Hanna, Mind-wandering as spontaneous thought: a dynamic framework. *Nat. Rev. Neurosci.* **17**, 718–731 (2016).

9.   D. E. Melnikoff, R. W. Carlson, P. E. Stillman, A computational theory of the subjective experience of flow. *Nat. Commun.* **13**, 2252 (2022).

10.   D. Rahnev, *et al.*, Consensus Goals in the Field of Visual Metacognition. *Perspect. Psychol. Sci. J. Assoc. Psychol. Sci.* **17**, 1746–1765 (2022).

11.   L. F. Barrett, Solving the Emotion Paradox: Categorization and the Experience of Emotion. *Personal. Soc. Psychol. Rev.* **10**, 20–46 (2006).

12.   J. E. LeDoux, Coming to terms with fear. *Proc. Natl. Acad. Sci.* **111**, 2871–2878 (2014).

13.   M. D. Robinson, G. L. Clore, Belief and feeling: Evidence for an accessibility model of emotional self-report. *Psychol. Bull.* **128**, 934–960 (2002).

14.   J. A. Russell, Core affect and the psychological construction of emotion. *Psychol. Rev.* **110**, 145–172 (2003).

15.   D. L. Paulhus, S. Vazire, The self-report method. *Handb. Res. Methods Personal. Psychol.* **1**, 224–239 (2007).

16.   D. J. Bem, Self-perception: An alternative interpretation of cognitive dissonance phenomena. *Psychol. Rev.* **74**, 183–200 (1967).

17. E. G. Boring, A history of introspection. *Psychol. Bull.* **50**, 169–189 (1953).

18. A. Freud, *The ego and the mechanisms of defence* (International Universities Press, 1946).

19. G. Hein, Y. Morishima, S. Leiberg, S. Sul, E. Fehr, The brain's functional network architecture reveals human motives. *Science* **351**, 1074–1078 (2016).

20. C. L. Hull, *A behavior system; an introduction to behavior theory concerning the individual organism* (Yale University Press, 1952).

21. D. C. McClelland, *Studies in motivation* (Appleton-Century-Crofts, 1955).

22. H. A. Murray, Thematic apperception test. (1943).

23. R. E. Nisbett, T. D. Wilson, Telling more than we can know: Verbal reports on mental processes. *Psychol. Rev.* **84**, 231 (1977).

24. R. S. Peters, *The concept of motivation* (Humanities Press, 1958).

25. P. J. Silvia, G. H. E. Gendolla, On introspection and self-perception: Does self-focused attention enable accurate self-knowledge? *Rev. Gen. Psychol.* **5**, 241–269 (2001).

26. J. B. Watson, Psychology as the behaviorist views it. *Psychol. Rev.* **20**, 158 (1913).

27. T. D. Wilson, *Strangers to ourselves* (Harvard University Press, 2004).

28. K. M. Carlsmith, The roles of retribution and utility in determining punishment. *J. Exp. Soc. Psychol.* **42**, 437–451 (2006).

29. M. J. Crockett, Y. Özdemir, E. Fehr, The value of vengeance and the demand for deterrence. *J. Exp. Psychol. Gen.* **143**, 2279 (2014).

30. P. W. Eastwick, E. J. Finkel, "Speed-dating: A powerful and flexible paradigm for studying romantic relationship initiation" in *Handbook of Relationship Initiation*, (Psychology Press, 2008), pp. 217–234.

31. I. H. Franken, P. Muris, Individual differences in decision-making. *Personal. Individ. Differ.* **39**, 991–998 (2005).

32. B. Gawronski, J. De Houwer, "Implicit measures in social and personality psychology" in *Handbook of Research Methods in Social and Personality Psychology, 2nd Ed*, (Cambridge University Press, 2014), pp. 283–310.

33. K. R. Menary, M. G. Kushner, E. Maurer, P. Thuras, The prevalence and clinical implications of self-medication among individuals with anxiety disorders. *J. Anxiety Disord.* **25**, 335–339 (2011).

34. G. Murphy, E. Loftus, R. H. Grady, L. J. Levine, C. M. Greene, Misremembering motives: The unreliability of voters' memories of the reasons for their vote. *J. Appl. Res. Mem. Cogn.* **9**, 564–575 (2020).

35. K. O'Connor, G. S. Adams, Affective antecedents of revenge. *Behav. Brain Sci.* **36**, 29–30 (2013).

36. P. M. Podsakoff, S. B. MacKenzie, N. P. Podsakoff, Sources of method bias in social science research and recommendations on how to control it. *Annu. Rev. Psychol.* **63**, 539–569 (2012).

37. K. White, D. R. Lehman, Culture and social comparison seeking: The role of self-motives. *Pers. Soc. Psychol. Bull.* **31**, 232–242 (2005).

38. D. G. Winter, O. P. John, A. J. Stewart, E. C. Klohnen, L. E. Duncan, Traits and motives: toward an integration of two traditions in personality research. *Psychol. Rev.* **105**, 230 (1998).

39. J. M. Nolan, P. W. Schultz, R. B. Cialdini, N. J. Goldstein, V. Griskevicius, Normative social influence is underdetected. *Pers. Soc. Psychol. Bull.* **34**, 913–923 (2008).

40. P. Carruthers, *The Opacity of Mind: An Integrative Theory of Self-Knowledge* (Oxford University Press, 2011).

41. U. Frith, F. Happé, Theory of mind and self-consciousness: What is it like to be autistic? *Mind Lang.* **14**, 82–89 (1999).

42. A. Gopnik, How we know our minds: The illusion of first-person knowledge of intentionality. *Behav. Brain Sci.* **16**, 1–14 (1993).

43. A. I. Jack, A. Roepstorff, Introspection and cognitive brain mapping: from stimulus–response to script–report. *Trends Cogn. Sci.* **6**, 333–339 (2002).

44. R. Saxe, S. Offen, "Seeing ourselves: What vision can teach us about metacognition" in *Metacognition and Severe Adult Mental Disorders: From Research to Treatment*, (Routledge/Taylor & Francis Group, 2010), pp. 13–29.

45. J. Beshears, J. J. Choi, D. Laibson, B. C. Madrian, How are preferences revealed? *J. Public Econ.* **92**, 1787–1794 (2008).

46. C. Jerolmack, S. Khan, Talk is cheap: Ethnography and the attitudinal fallacy. *Sociol. Methods Res.* **43**, 178–209 (2014).

47. N. Chater, *The mind is flat: the illusion of mental depth and the improvised mind* (Penguin UK, 2018).

48. W. T. Harbaugh, U. Mayr, D. R. Burghart, Neural responses to taxation and voluntary giving reveal motives for charitable donations. *Science* **316**, 1622–1625 (2007).

49. M. Pessiglione, *et al.*, How the Brain Translates Money into Force: A Neuroimaging Study of Subliminal Motivation. *Science* **316**, 904–906 (2007).

50. Y. Li, J. Hu, C. C. Ruff, X. Zhou, Neurocomputational evidence that conflicting prosocial motives guide distributive justice. *Proc. Natl. Acad. Sci.* **119**, e2209078119 (2022).

51. R. Custers, H. Aarts, The Unconscious Will: How the Pursuit of Goals Operates Outside of Conscious Awareness. *Science* **329**, 47–50 (2010).

52. O. C. Schultheiss, *Implicit motives.* (The Guilford Press, 2008).

53. E. Fehr, A. Rangel, Neuroeconomic Foundations of Economic Choice–Recent Advances. *J. Econ. Perspect.* **25**, 3–30 (2011).

54. R. C. Wilson, A. G. Collins, Ten simple rules for the computational modeling of behavioral data. *eLife* **8**, e49547 (2019).

55. D. J. Bem, "Self-Perception Theory11Development of self-perception theory was supported primarily by a grant from the National Science Foundation (GS 1452) awarded to the author during his tenure at Carnegie-Mellon University." in *Advances in Experimental Social Psychology*, L. Berkowitz, Ed. (Academic Press, 1972), pp. 1–62.

56. G. Ryle, *The concept of mind* (Barnes & Noble, 1949).

57. J. Haidt, The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychol. Rev.* **108**, 814–834 (2001).

58. C. S. Dweck, *Self-theories: Their role in motivation, personality, and development* (Psychology Press, 1999).

59. T. D. Wilson, J. I. Stone, Limitations of self-knowledge: More on telling more than we can know. *Rev. Personal. Soc. Psychol.* **6**, 167–183 (1985).

60. F. Cushman, Rationalization is rational. *Behav. Brain Sci.* **43** (2020).

61. A. N. Hornsby, B. C. Love, How decisions and the desire for coherency shape subjective preferences over time. *Cognition* **200**, 104244 (2020).

62. H. H. Kelley, The processes of causal attribution. *Am. Psychol.* **28**, 107–128 (1973).

63. R. E. Nisbett, L. Ross, Human Inference: Strategies and Shortcomings of Social Judgment (1980) (February 14, 2023).

64. K. Simler, R. Hanson, *The elephant in the brain: Hidden motives in everyday life* (Oxford University Press, 2017).

65. A. G. Greenwald, The totalitarian ego: Fabrication and revision of personal history. *Am. Psychol.* **35**, 603–618 (1980).

66.   Z. Kunda, The case for motivated reasoning. *Psychol. Bull.* **108**, 480–498 (1990).

67.   R. W. Robins, O. P. John, "The quest for self-insight: Theory and research on accuracy and bias in self-perception" in *Handbook of Personality Psychology*, (Elsevier, 1997), pp. 649–679.

68.   C. Sedikides, A. P. Gregg, Self-enhancement: Food for thought. *Perspect. Psychol. Sci.* **3**, 102–116 (2008).

69.   T. A. Wills, Downward comparison principles in social psychology. *Psychol. Bull.* **90**, 245 (1981).

70.   A. L. Edwards, *The social desirability variable in personality assessment and research* (Dryden Press, 1957).

71.   R. F. Baumeister, A self-presentational view of social phenomena. *Psychol. Bull.* **91**, 3–26 (1982).

72.   R. W. Carlson, Y. E. Bigman, K. Gray, M. J. Ferguson, M. J. Crockett, How inferred motives shape moral judgements. *Nat. Rev. Psychol.* **1**, 468–478 (2022).

73.   G. Reeder, Mindreading: Judgments about intentionality and motives in dispositional inference. *Psychol. Inq.* **20**, 1–18 (2009).

74.   L. Young, L. Tsoi, When mental states matter, when they don't, and what that means for morality. *Soc. Personal. Psychol. Compass* **7**, 585–604 (2013).

75.   C. D. Batson, *What's wrong with morality?: a social-psychological perspective* (Oxford University Press, 2016).

76.   R. W. Carlson, M. A. Maréchal, B. Oud, E. Fehr, M. J. Crockett, Motivated misremembering of selfish decisions. *Nat. Commun.* **11**, 2100 (2020).

77.   N. Epley, D. Dunning, Feeling "holier than thou": Are self-serving assessments produced by errors in self- or social prediction? *J. Pers. Soc. Psychol.* **79**, 861–875 (2000).

78.   B. Monin, A. Merritt, Moral hypocrisy, moral inconsistency, and the struggle for moral integrity. (2012).

79.   G. E. Bolton, E. Katok, R. Zwick, Dictator game giving: Rules of fairness versus acts of kindness. *Int. J. Game Theory* **27**, 269–299 (1998).

80.   C. A. Hutcherson, B. Bushong, A. Rangel, A Neurocomputational Model of Altruistic Choice and Its Implications. *Neuron* **87**, 451–462 (2015).

81.   R. O. Murphy, K. A. Ackermann, Social value orientation: Theoretical and measurement issues in the study of social preferences. *Personal. Soc. Psychol. Rev.* **18**, 13–41 (2014).

82. A. Waytz, J. Zaki, J. P. Mitchell, Response of dorsomedial prefrontal cortex predicts altruistic behavior. *J. Neurosci.* **32**, 7646–7650 (2012).

83. P. Johansson, L. Hall, S. Sikström, A. Olsson, Failure to detect mismatches between intention and outcome in a simple decision task. *Science* **310**, 116–119 (2005).

84. K. Lewin, *The conceptual representation and the measurement of psychological forces* (Duke University Press, 1938).

85. J. Kagan, Motives and development. *J. Pers. Soc. Psychol.* **22**, 51–66 (1972).

86. K. A. Ericsson, H. A. Simon, Verbal reports as data. *Psychol. Rev.* **87**, 215–251 (1980).

87. E. M. Caruso, D. A. Rahnev, M. R. Banaji, Using conjoint analysis to detect discrimination: Revealing covert preferences from overt choices. *Soc. Cogn.* **27**, 128–137 (2009).

88. R. M. Dawes, The robust beauty of improper linear models in decision making. *Am. Psychol.* **34**, 571 (1979).

89. R. L. Keeney, H. Raiffa, R. F. Meyer, *Decisions with multiple objectives: preferences and value trade-offs* (Cambridge university press, 1993).

90. R. E. Nisbett, N. Bellows, Verbal reports about causal influences on social judgments: Private access versus public theories. *J. Pers. Soc. Psychol.* **35**, 613 (1977).

91. P. Slovic, S. Lichtenstein, Comparison of Bayesian and regression approaches to the study of information processing in judgment. *Organ. Behav. Hum. Perform.* **6**, 649–744 (1971).

92. J. Li, *et al.*, The effect of preceding self-control on prosocial behaviors: the moderating role of awe. *Front. Psychol.* **10**, 682 (2019).

93. Y. Y. Teoh, Z. Yao, W. A. Cunningham, C. A. Hutcherson, Attentional priorities drive effects of time pressure on altruistic choice. *Nat. Commun.* **11**, 3534 (2020).

94. N. D. Daw, Trial-by-trial data analysis using computational models. *Decis. Mak. Affect Learn. Atten. Perform. XXIII* **23** (2011).

95. E. Fehr, A. Rangel, Neuroeconomic Foundations of Economic Choice–Recent Advances. *J. Econ. Perspect.* **25**, 3–30 (2011).

96. R. C. Wilson, A. G. Collins, Ten simple rules for the computational modeling of behavioral data. *Elife* **8**, e49547 (2019).

97. I. W. Eisenberg, *et al.*, Uncovering the structure of self-regulation through data-driven ontology discovery. *Nat. Commun.* **10**, 2319 (2019).

98. R. Frey, A. Pedroni, R. Mata, J. Rieskamp, R. Hertwig, Risk preference shares the psychometric structure of major psychological traits. *Sci. Adv.* **3**, e1701381 (2017).

99. B. A. Murphy, S. O. Lilienfeld, Are self-report cognitive empathy ratings valid proxies for cognitive empathy ability? Negligible meta-analytic relations with behavioral task performance. *Psychol. Assess.* **31**, 1062–1072 (2019).

100. M. Kouchaki, F. Gino, Memories of unethical actions become obfuscated over time. *Proc. Natl. Acad. Sci.* **113**, 6166–6171 (2016).

101. C. Saucet, M. C. Villeval, Motivated memory in dictator games. *Games Econ. Behav.* **117**, 250–275 (2019).

102. A. Barasch, E. E. Levine, J. Z. Berman, D. A. Small, Selfish or selfless? On the signal value of emotion in altruistic behavior. *J. Pers. Soc. Psychol.* **107**, 393 (2014).

103. R. W. Carlson, J. Zaki, Belief in Altruistic Motives Predicts Prosocial Actions and Inferences. *Psychol. Rep.*, 00332941211013529 (2021).

104. J. Jara-Ettinger, L. E. Schulz, J. B. Tenenbaum, The Naïve Utility Calculus as a unified, quantitative framework for action understanding. *Cognit. Psychol.* **123** (2020).

105. D. Dunning, E. G. Helzer, Beyond the correlation coefficient in studies of self-assessment accuracy: Commentary on Zell & Krizan (2014). *Perspect. Psychol. Sci.* **9**, 126–130 (2014).

106. K. N. Rotella, J. A. Richeson, Motivated to "forget" the effects of in-group wrongdoing on memory and collective guilt. *Soc. Psychol. Personal. Sci.* **4**, 730–737 (2013).

107. L. L. Shu, F. Gino, M. H. Bazerman, Dishonest deed, clear conscience: When cheating leads to moral disengagement and motivated forgetting. *Pers. Soc. Psychol. Bull.* **37**, 330–349 (2011).

108. T. D. Wilson, R. E. Nisbett, The accuracy of verbal reports about the effects of stimuli on evaluations and behavior. *Soc. Psychol.* **41**, 118–131 (1978).

109. R. T. Hurlburt, C. L. Heavey, Telling what we know: describing inner experience. *Trends Cogn. Sci.* **5**, 400–403 (2001).

110. B. R. Newell, D. R. Shanks, Unconscious influences on decision making: A critical review. *Behav. Brain Sci.* **37**, 1–18 (2014).

111. M. H. Davis, Measuring individual differences in empathy: Evidence for a multidimensional approach. *J. Pers. Soc. Psychol.* **44**, 113–126 (1983).

112. J. Dang, K. M. King, M. Inzlicht, Why Are Self-Report and Behavioral Measures Weakly Correlated? *Trends Cogn. Sci.* **24**, 267–269 (2020).

113. O. FeldmanHall, *et al.*, What we say and what we do: The relationship between real and hypothetical moral choices. *Cognition* **123**, 434–441 (2012).

114. N. K. Malhotra, Information load and consumer decision making. *J. Consum. Res.* **8**, 419–430 (1982).

115. G. A. Miller, The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychol. Rev.* **63**, 81–97 (1956).

116. N. Mazar, O. Amir, D. Ariely, The Dishonesty of Honest People: A Theory of Self-Concept Maintenance. *J. Mark. Res.* **45**, 633–644 (2008).

117. S. Shalvi, F. Gino, R. Barkan, S. Ayal, Self-Serving Justifications: Doing Wrong and Feeling Moral. *Curr. Dir. Psychol. Sci.* **24**, 125–130 (2015).

118. T. Stafford, Evidence for the rationalisation phenomenon is exaggerated. *Behav. Brain Sci.* **43**, e48 (2020).