

Objective and Subjective Compliance: A Norm-based Explanation of ‘Moral Wiggle Room’

Kai Spiekermann

London School of Economics
Department of Government
Houghton Street
London WC2A 2AE
United Kingdom
Email: k.spiekermann@lse.ac.uk

Arne Weiss^{a,b}

(Corresponding author)
^aDFG Research Group 2104 Needs-Based
Justice and Distribution Procedures
University of Oldenburg
Institute of Theoretical Philosophy
Ammerländer Heerstrasse 136
D-26129 Oldenburg, Germany
^bCenter for Social and Economic
Behavior, University of Cologne
Email: arne.weiss@uni-oldenburg.de

Abstract

We propose a cognitive-dissonance model of norm compliance to identify conditions for selfishly biased information acquisition. The model distinguishes between: (i) objective norm compliers, for whom the right action is a function of the state of the world; (ii) subjective norm compliers, for whom it is a function of their belief. The former seek as much information as possible; the latter acquire only information that lowers, in expected terms, normative demands. The source of ‘moral wiggle room’ is not belief manipulation, but the coarseness of normative prescriptions under conditions of uncertainty. In a novel experimental setup, we find evidence for such strategic information uptake. Our results suggest that attempts to change behavior by subjecting individuals to norms can lead to biased information acquisition instead of compliance.

JEL Classification: C91, D63, D83

Keywords: norm compliance, uncertainty, experiment, self-serving biases, strategic learning, dictator game

Acknowledgements

We would like to thank James Konow, Dirk Sliwka and two anonymous referees for extensive written comments. Kiryl Khalmetski, Oliver Gürtler and Marcus Pivato provided kind mathematical help, for which we are very grateful (but all remaining errors are our own). We also thank Roberto Weber, Christoph Engel, Manfred Königstein, Bettina Rockenbach, Christopher Zeppenfeld and our colleagues in Erfurt, Cologne and London for illuminating discussions, two anonymous referees for extensive and constructive comments and Anne Schielke for valuable research assistance. This paper was presented at the ESA European Conference 2012 in Cologne, the ESA 2013 World Meetings in Zurich, the EEA-ESEM 2013 in Gothenburg, and many smaller workshops; we owe thanks to all our audiences. Financial support by the University of Erfurt and the London School of Economics is gratefully acknowledged.

1. Introduction

Suppose an individual feels compelled to comply with a norm but is able to influence what she knows about the normatively relevant facts. Can she use her information acquisition choices strategically? For example, imagine a company owner having to decide the size of a voluntary employee performance bonus. The employee is about to leave the company for good and a ‘selfish’ employer would therefore give nothing. However, suppose the owner feels bound by this social norm: ‘if an employee performed well, he ought to be rewarded’. The norm compels the employer to pay a bonus, but the size of the bonus depends on her beliefs about the employee’s performance. In such a situation, the owner may want to find out that the employee performed poorly in order to save money without violating the social norm.

In this paper, we demonstrate that individuals can ‘wiggle’ their way out of normative pressure by strategically acquiring only information that can activate less costly prescriptions. According to our model, these strategic incentives arise if norms prescribe the same action for different degrees of beliefs under conditions of uncertainty. Our analysis thereby reveals the important but so far largely neglected role of uncertainty in models of norm-compliance. We also offer experimental evidence that individuals strategically select normatively relevant information.

We first provide the intuition for the theoretical mechanism we have in mind before turning to a formal analysis. To continue the introductory example, assume the employee has worked in a customer service function, and the employer, unable to observe the employee’s performance directly, has to ask clients to learn whether the employee performed well. To find out, she can contact two different clients for references: one known to be a friend and one known to be a foe of the employee. Both are honest but have a poor response rate when being asked for reference. First, if the employee’s performance was good, the friend sometimes replies with a

30 positive report and sometimes fails to respond. The foe never responds to
 avoid saying anything positive. Second, if the employee’s performance was
 bad, the friend never responds to avoid saying anything negative. The foe
 sometimes replies with a negative report and sometimes fails to respond.
 In other words, the employer has the option to choose one or two types
 35 of probabilistic signals about the employee’s performance. Each signal can
 only lead to certainty for one of the two possible states of the world as
 displayed in this table:

		employee’s performance	
		good	bad
reference from	friend	‘good’ or no response	no response
	foe	no response	‘bad’ or no response

In a formal model, we derive two jointly sufficient conditions under which
 40 asking only the foe is an attractive way to save norm-compliance costs.
 First, the employer must interpret the norm in a ‘subjective’ sense, so that
 complying with the norm means performing the action that is consistent
 with her beliefs about performance (rather than with what is factually the
 case). Only if the demandingness of the norm (the cost of complying with
 45 the norm prescription) depends on beliefs is it worthwhile to manipulate
 the beliefs by strategic information selection.¹ By contrast, if the employer
 interprets the norm in an ‘objective’ sense, the prescription does not de-
 pend on beliefs and the strategic incentives vanish. Second, for strategic
 information selection to arise, it also has to be the case that the action
 50 prescribed by the norm only changes once the employer becomes (near)
 certain about the employee’s performance.² The demandingness of the

¹We will use the terms ‘demandingness of the norm,’ ‘demand of the norm’ and ‘normative demand’ interchangeably. We also use the terms ‘demanding’ and ‘undemanding’ to compare norm compliance costs. With respect to the selfish employer, the demandingness increases if the norm prescribes a higher bonus payment. More generally, the demandingness of the norm rises if the action it prescribes is further apart from the action that would maximize utility if there was no normative prescription.

²We thank an anonymous referee for helping us to clarify this.

norm is thus quite insensitive to changes in beliefs (as long as they do not approach certainty). We later introduce the notion of ‘coarse-grained’ norms to capture this property formally. If the two conditions are met, choosing to ask only the foe is attractive: if the foe responds, it lowers the normative demand, allowing the employer to give less; if the foe does not respond, the normative demand remains unchanged. We also test the implications of our model experimentally with a normatively framed dictator game with optional information choice, similar to the bonus payment example introduced here.

The phenomenon we are interested in is often framed in terms of ‘strategic ignorance’. The seminal paper by Dana, Weber and Kuang (2007, ‘DWK’ from now on) shows that uncertainty over the receiver’s payoff increases ‘selfish’ dictator choices. Strikingly, in DWK’s experiments the uncertainty can be resolved costlessly, but most dictators with selfish behavior decide to remain ignorant, suggesting that these dictators ‘have an illusory preference for fairness’ and ‘dislike appearing unfair’ (DWK, p. 67). Grossmann and van der Weele (2013) offer a formal model for such strategic ignorance. It suggests that DWK’s dictator information avoidance is driven by self-image concerns in a Bayesian preference-signaling model, building on earlier work by Bénabou and Tirole (2006, see also 2011). An agent cares about his own payoff, to varying degrees about the payoff of others, and about his self-image as an altruistic person, signaled by the actions performed. The agent can choose whether or not to commit to a costly prosocial action with a potential welfare benefit. Grossman and van der Weele show the existence of an equilibrium in which agents with limited altruism decide to remain ignorant and refrain from acting prosocially. Crucially, among the ignorant agents some would act prosocially if they knew with certainty about the positive social benefits. The results are driven by the fact that the signal about the agent’s type is stronger if the agent fails to engage in an action with certain rather than only potential wel-

fare benefits.³ Grossman and van der Weele’s model (like all self-signaling models, starting with Bodner and Prelec (2003)), crucially hinges on the assumption of a *dual* self, such that one self is relatively more informed
85 about the agent’s type than the other self.⁴ Feiler (2014) offers an explanation of DWK’s results based on the model of moral preferences and moral constraints by Rabin (1995): Agents following a moral rule have incentives to ‘manipulate’ their beliefs to circumvent the constraints imposed by this rule. Matthey and Regner (2011) provide yet another explanation
90 for strategic ignorance. They argue that ignorance is attractive because it lowers the cost of deceiving oneself about what one ought to do.

In this paper, we take a different route. First, the agents in our model have perfect knowledge about their preferences, and uncertainty does not serve as a self-deception tool. Second, we maintain that the focus on igno-
95 rance in previous papers has been too narrow. Our evidence suggests that ignorance itself is not a particularly special or desirable state (pace DWK, Bicchieri, 2006, p. 128-9, and others). Rather, individuals will only stay ignorant if resolving uncertainty will, in expected terms, increase normative obligations and will actively *seek* information if, in expected terms, norma-
100 tive obligations will decrease. We locate the source of moral wiggle room in the nature of the norms followed. This, we will show, offers a parsimonious yet powerful account of the mechanisms underlying moral wiggling.

Our findings add to the literature on self-serving biases. It has long been recognized that norms often leave a fairly large room for interpre-
105 tation (Hechter and Opp, 2001) and create opportunities for self-serving biases. These biases in norm compliance are based on different mecha-

³Two experimental tests provide support for their model: Moral ignorance is strategically more valuable before a choice is made and some agents are willing to pay for remaining ignorant.

⁴When taking decisions, the agent weighs in and anticipates which inference the uninformed self will draw about the agent’s type. Conceptually, we can think of the uninformed self as a future self with imperfect recall about the underlying motivation behind past choices. The present self, being aware of the future self’s forgetfulness, tries to manipulate the latter’s inference. While the authors stress the importance of self-image concerns, it should be noted that the model stays formally the same when the observer is another person, turning it into a social-image model.

nisms: they can make use of uncertainty about either the behavioral rule to be applied (‘what *exactly* am I supposed to do?’) or—more relevant for us—about the state we are in (‘what *exactly* is the case?’). Konow
110 (2000) develops a model of self-serving biases due to the former and experimentally tests for these biases by studying the strategic manipulation of beliefs about the applicable norm. Related experiments are conducted by Bicchieri and Chavez (2013), Bicchieri and Mercier (2013), and Rodriguez-Lara and Moreno-Garrido (2012). By contrast, DWK find evidence for the
115 relevance of factual uncertainty. Similar results were obtained by Larson and Capra (2009); Grossman (2010); Matthey and Regner (2011); Conrads and Irlenbusch (2013); Grossman and van der Weele (2013) and Feiler (2014).

The plan for the remainder of the paper is as follows. In section 2,
120 we motivate the different notions of norm compliance. The formal model follows in section 3 and is extended beyond the core setup in section 4 (which a reader in a rush can safely skip). In section 5, we relate the model to the existing literature on strategic ignorance. We then test the model’s core predictions with a laboratory experiment in section 6. In the final
125 section, we summarize this paper’s theoretical and empirical contributions and put them into perspective.

2. Objective and subjective norm compliance

A social norm tells us what we ought to do if we find ourselves in a certain situation; more technically, a social norm provides a mapping from a state
130 to a behavioral rule, i.e., a prescribed or proscribed act (see Bicchieri, 2006, ch. 1). In short, social norms take the form: ‘if X obtains, I ought to do ϕ ’. The condition ‘if X obtains’, for which the norm prescribes a certain behavior (‘I ought to do ϕ ’), can be interpreted in different ways. On the one hand, the clause can be substituted by ‘if X is the case’. This
135 *objective* interpretation leads to prescriptions that are conditional only on

the state of the world and are thus entirely independent from the agent's beliefs. Consequently, an *objective norm complier* strives to perform the action the state of the world demands. On the other hand, the clause can be substituted by 'if I believe that X '. This *subjective* interpretation leads
140 to prescriptions that are entirely contingent on the beliefs of the agent. If the norm is understood in the subjective sense, it is only prescriptive if the agent has the specified belief. Thus, a *subjective norm-complier* strives to perform the action his beliefs demand.⁵

Depending on one's subjective or objective understanding of norms, one
145 can be subject to different forms of psychological costs when violating a norm. Following Konow (2000), these psychological costs from non-compliance can be called *cognitive dissonance*.⁶ The dissonance arises because the agent experiences an unpleasant tension between what she ought to do and what she actually does. In Konow's model, the experienced dissonance is traded off against utility from violating the norm, leading to
150 more selfish behavior than prescribed by the norm.

Under certainty, objective and subjective compliance are behaviorally equivalent because the epistemic state matches the state of the world. However, the two types of compliance come apart under uncertainty. An objective
155 norm complier will suffer from dissonance under uncertainty because she cannot (at the same time) comply with what the norm prescribes for two different states. In other words: an objective norm complier feels com-

⁵This distinction is echoed in ethical theory and epistemology. For example, Zimmerman (2008) distinguishes between the objective and subjective view of moral obligations. According to the objective view, one's moral obligations are determined only by the relevant facts, not by what one knows about these facts. The objective view therefore entails that ignorance or incomplete information does not change one's moral obligations at all (though most proponents of this view would say that wrongdoing due to exculpatory ignorance may be blameless). According to the subjective view, by contrast, moral obligations are a function of one's knowledge. This entails that incomplete or false knowledge of the facts changes one's moral obligations. Zimmerman finds neither view convincing and instead defends the prospective view, which (very roughly) expresses moral obligations as determined by expected values, given one's available evidence.

⁶The concept of cognitive dissonance builds on Festinger (1957) and was introduced into formal economic theory by Akerlof and Dickens (1982) in order to explore the welfare consequences of biased beliefs.

pelled to ‘get it right’ and suffers dissonance because she (potentially) fails to do so. By contrast, for a subjective norm complier, uncertainty is just
160 another possible epistemic state to which a norm applies. Doing what the epistemic state requires suffices to meet one’s obligations. Given this difference, objective and subjective compliers have different incentives to resolve uncertainty. In our treatment of uncertainty, we assume that agents have no intrinsic preference for information and also abstract from ‘meta-norms’
165 that prescribe information acquisition. This is not to say that neither is relevant in actual decision making. Instead, abstracting from these confounding motivations to acquire information gives us a clear distinction between two very different forms of norm compliance.⁷

The example of the employee’s leaving bonus gave us a first taste how
170 subjective norm compliance and coarse-grained norms can create strategic incentives for acquiring information selectively. Similar incentives to strategically acquire information can also apply in many other settings. For instance, many people may endorse a norm against harming others. At the same time, they cause excessive greenhouse gas emissions while preserving
175 a belief that the impacts of anthropogenic climate change are unclear. They might ignore articles in a Greenpeace magazine (or the IPCC reports for that matter), but may well pay active attention to ‘climate change deniers’. By consuming information from the latter, they may expect to either find convincing evidence that there is not much to worry about one’s personal
180 lifestyle or to learn nothing substantially new at all. In the same vein, a consumer may avoid documentaries about the suffering of caged hens while lapping up news about ‘organic food’ scams in order to avoid the normative pressure to buy expensive eggs.

The behavior in such examples can be rationalized within our model by
185 assuming coarseness in the mapping from beliefs to prescriptions and the availability of signals for which a subjective norm complier has prior beliefs

⁷See Grossman and van der Weele (2013) for a model in which information acquisition becomes morally relevant.

about whether he might attain an attractive or an unattractive epistemic state after acquiring the signals. The next section will formalize these concepts and claims.

190 **3. Formal model**

In order to introduce a formal distinction between objective and subjective compliance, we model a dictator game, enriched by social norms of equity, such that the receivers are more or less deserving. The dictator is initially uncertain about the deservingness of the receiver, but can acquire signals
195 to eliminate this uncertainty. We will embed this setting into a simplified version of Konow’s (2000) model, and then extend this baseline model to fit it to our setup.

The dictator has the amount \bar{y} to distribute, such that he gives y to himself and $x = \bar{y} - y$ to the receiver ($0 \leq y \leq \bar{y}$). The payoff utility
200 derived from this decision is $v(y)$, with the usual assumption of positive but decreasing marginal utility of money, such that $v'(y) > 0$ and $v''(y) < 0$. If dictators were maximizing utility from monetary payoff only, their obvious choice would be to set $y = \bar{y}$. But besides monetary payoff, dictators also care about limiting the difference between what they think is normatively
205 required and what they actually give. The greater the difference, the higher are the non-compliance costs in the form of cognitive dissonance.

For simplicity, we assume a dichotomous state space $\Omega = \{L, H\}$, where L and H can conveniently be interpreted as ‘low’ and ‘high’, indicating the deservingness of the receiver.⁸ We call the actual state ω . For an objective
210 norm complier, the amount the dictator takes to be normatively permitted to keep (the ‘personal fair point’) is a function of the actual state ω ; this action is characterized by keeping ϕ_ω . For a subjective norm complier, by contrast, the fair point is a function of an epistemic state, characterized by the (Bayesian) probability p that state L obtains (and $1 - p$ that state H

⁸A generalization to other state space partitions should be straightforward.

215 obtains). This distinction between objective and subjective definitions of what is normatively required, expressed by ϕ_ω and ϕ_p respectively, is the crucial extension of Konow’s 2000 model. To simplify the analysis, we do not model the internal process through which agents form their personal fair points but treat them as exogenous.⁹ In the setting we consider, we assume
 220 the ϕ s to be a function of shared empirical and normative expectations of a relevant reference group.¹⁰

The dissonance cost experienced by the dictator is a function of the difference $\Delta = |y - \phi|$ between the the fair point and what she actually keeps, and the dissonance cost function $f(y, \phi) = f(\Delta)$ determines the
 225 experienced disutility. As in Konow (2000), f is a twice differentiable, strictly convex function, such that $f(0) = 0$ (that is, if $y = \phi$), $f'(\Delta) > 0$ for $\Delta \neq 0$, and $f''(\Delta) > 0$.

The dictator’s decision problem is to trade off the utility from keeping more money against the disutility from cognitive dissonance created by
 230 deviating from the perceived fair distribution:

$$\max_y E[u(y, \phi)] \equiv E[v(y) - f(y, \phi)] \text{ subject to } 0 \leq y \leq \bar{y}, 0 \leq \phi \leq \bar{\phi}.$$

The behavior of a dictator is therefore determined by the fair point ϕ and by the relative value of money and cognitive-dissonance avoidance.

We first model maximization under certainty, in which factual and epistemic states match. Let ϕ_L and ϕ_H denote the fair points given the re-
 235 spective state, assuming that $\bar{y} > \phi_L > \phi_H > 0$ (that is, it is fair to keep more in state L than in state H , but it is never fair to keep everything). Let $\hat{u}(\phi) = \max_y (v(y) - f(y, \phi))$ be the maximum utility achievable as a function of ϕ . This function is increasing in ϕ as a higher fair point al-

⁹Konow (2000) also considers self-serving biases in the belief formation about what is fair. As our research question focuses on strategic acquisition of information on the receiver’s entitlement once the dictator has formed his belief on what he ought to do under certainty, we do not model how individuals reach their normative beliefs. We therefore take them as given, even though we concur that this is another important form of self-serving bias.

¹⁰For details see Bicchieri (2006, ch. 1).

lows the dictator to keep more for himself.¹¹ Since we have assumed that
 240 $\phi_L > \phi_H$, it follows that $\hat{u}(\phi_L) > \hat{u}(\phi_H)$. When deciding on an allocation in state H , the dictator reaches the maximum utility $\hat{u}(\phi_H)$, with $y_H^* = \operatorname{argmax}_{y_H}(u(y_H, \phi_H))$. Similarly, when deciding on an allocation in state L , the maximum utility is $\hat{u}(\phi_L)$, with $y_L^* = \operatorname{argmax}_{y_L}(u(y_L, \phi_L))$.¹²

The novel aspect of our model is the treatment of uncertainty. For an
 245 objective norm complier, cognitive dissonance depends on the state and can only be $f(y, \phi_H)$ or $f(y, \phi_L)$. Thus, under uncertainty, an objective complier cannot make sure to choose the morally appropriate action that the (unknown) state requires. Therefore, the *maximum expected* utility under uncertainty for a given p is achieved when keeping the amount
 250 $y_U^*(p) = \operatorname{argmax}_{y_U}(v(y_U) - (1 - p)f(y_U, \phi_H) - pf(y_U, \phi_L))$.

For a subjective norm complier, the normatively required action is a function of his epistemic state. A norm implies a function that maps a state onto a required action, in our case the amount one is permitted to keep. We believe that social norms typically only provide a coarse-grained
 255 mapping from states to prescriptions (but will consider a different mapping in the next section). The reason is tied to the necessity of enforceability. For a social norm to exist, individuals must be able to reliably distinguish between compliant and non-compliant agents in order to form behavioral expectations and sanction transgressions. Since degrees of beliefs are not
 260 observable in detail, it is unlikely that social norms take them as argument with any great precision. A norm that distinguishes between, for instance, degrees of beliefs 0.6 and 0.7 simply cannot be enforced. This is mirrored in our everyday language regarding normative choices, in which we rarely refer to degrees of beliefs (let alone a Bayesian updating of normative pre-
 265 scriptions).

The function we consider here distinguishes between three epistemic states: knowing that the receiver is deserving, knowing that the receiver is

¹¹A formal proof is provided in the appendix.

¹²To ensure an interior solution, we assume that $v'(\bar{y}) < \frac{\partial f(\bar{y}, \phi_L)}{\partial y}$, which implies that $y_L^* < \bar{y}$ and $y_L^* > y_H^*$.

undeserving, or being uncertain about the receiver’s deservingness. In such a case, the fair amount ϕ is determined by a step function:

$$\phi_p = \begin{cases} \phi_L & \text{if } p \geq 1 - \epsilon \\ \phi_U & \text{if } \epsilon < p < 1 - \epsilon \\ \phi_H & \text{if } p \leq \epsilon. \end{cases} \quad (\text{COARSE})$$

270 As the most general case, we assume that $\phi_L > \phi_U > \phi_H$ but will later consider other orderings. The parameter ϵ ($\epsilon \geq 0$) represents the margin of tolerance for ‘near certainty’, i.e. the threshold beyond which there is no longer ‘reasonable doubt’ about the state. For now, think of ϵ as being close to zero; we will later provide more precise boundary conditions. The
 275 maximum utility under uncertainty is $\hat{u}(\phi_p)$, with $y_p^* = \underset{y}{\operatorname{argmax}}(u(y, \phi_p))$.

In order to learn about the state of the world, dictators can optionally acquire costless signals. There are two different signals available, represented by random variables S_L and S_H , and for each the dictator can choose whether she wants to receive them. When obtaining signal S_L ,
 280 the dictator has a chance to learn that the state is L . More precisely, the conditional probability $s \in (0, 1)$ of learning that the state is L , given that the state is indeed L , is $s = \Pr(S_L = L | \omega = L)$. Similarly, when obtaining signal S_H , we assume the same value s for the conditional probability of learning that the state is H , given that the state is indeed H :
 285 $s = \Pr(S_H = H | \omega = H)$. In all other cases the dictator receives a ‘null’ signal, which means the dictator remains uncertain about the state. After a null signal, dictators perform a Bayesian update on the probability p that state L obtains. If the dictator receives only signal $S_L = 0$, then she updates such that $p' = (1 - s)p / ((1 - s)p + (1 - p))$. Similarly, if the dictator
 290 receives $S_H = 0$, then she updates such that $p' = p / (p + (1 - s)(1 - p))$. We assume for both possible updates p' that receiving a null signal never removes uncertainty. This provides the following boundary conditions on ϵ : $\epsilon < p' < 1 - \epsilon$. Finally, if a dictator acquires both signals, but both are

null signals, no update is necessary, as the two signals cancel each other
295 out.

We can now state how objective and subjective norm compliers will acquire the signals on offer.

Proposition 1. *Under uncertainty, objective compliers will acquire both signal S_L and signal S_H .*

300 Both types of signals increase the dictator's chance to reach his utility maxima given the respective states. Having all available signals is best because additional free information can never reduce and will usually increase expected utility as information increases the dictator's chance to play the optimal response to the state of the world. A proof is in the appendix.

305 **Proposition 2.** *Under uncertainty, subjective compliers who follow a coarse-grained norm (COARSE) will acquire signal S_L but not signal S_H .*

For a sketch proof, recall that the maximum utility is increasing in the fair amount to keep: $\hat{u}(\phi_L) > \hat{u}(\phi_U) > \hat{u}(\phi_H)$. COARSE, together with the assumption that a null signal never removes uncertainty, ensures that
310 any update from p to p' after receiving a null signal does not change ϕ_p . It is immediately obvious that obtaining signal S_L is beneficial because there is no down-side risk but a potential gain: it only increases the probability of receiving the highest utility $\hat{u}(\phi_L)$. And it is equally obvious that obtaining signal S_H is never beneficial because there is a down-side risk but
315 no potential gain: it only increases the probability of receiving the lowest utility $\hat{u}(\phi_H)$. A proof is provided in the appendix.

In our model all dictators hold unbiased (Bayesian) beliefs and all types of dictators would prefer to be in a world in which they are paired with a low performer and would use this fact to give little. However, an objective
320 norm complier, who only cares about the state of the world, cannot change the state and is therefore always better off with more information. A subjective norm complier with coarse-grained norms, by contrast, who takes prescriptions as a function of his beliefs, has 'moral wiggle room' (DWK).

4. Model Extensions

325 4.1. Social norms with linear fairness points

While there are good reasons to take the mapping from beliefs to prescription to be coarse-grained, the model also allows for different mappings. On the other end of the spectrum from coarse to continuous, it could be assumed that the social norm is perfectly sensitive to p , stating a different
330 prescription for each epistemic state. For instance, the fair amount could be a weighted average of the prescriptions under certainty, such that

$$\phi_p = p\phi_L + (1 - p)\phi_H. \quad (\text{LINEAR})$$

Such a norm makes demands directly proportional to the Bayesian belief.¹³

With LINEAR, the predictions for information acquisition are dramatically different:

335 **Proposition 3.** *Under uncertainty, subjective compliers who follow a linear norm (LINEAR) will not acquire any signals.*

Intuitively, signals provide fair lotteries over different levels of ϕ_p . However, because of the decreasing marginal value of money the dictators are risk averse and the maximum utility $\hat{u}(\phi)$ is concave.¹⁴ Therefore all such
340 lotteries have a lower expected utility than the utility of the status quo. A proof is provided in the appendix.

Note that the difference between propositions 2 and 3 lies in the way dictators do or do not update what is normatively required when uncertainty remains after acquiring a signal. When receiving a null signal after
345 acquiring S_L , the fair point of a dictator following LINEAR will decrease such that he will feel forced to keep less for herself. In expected terms, this decrease of the fair point entirely offsets the possible increase of the fair

¹³More precisely, LINEAR minimizes the expected distance between ϕ_p and ϕ_ω , the ϕ for the true (but unknown) state of the world.

¹⁴A formal proof is provided in the appendix.

point if the dictators learns that the state is L . Based on the assumption that marginal utility from keeping money is decreasing, a dictator following LINEAR will therefore gladly decline the lotteries offered by either S_L or S_H . By contrast, a dictator following COARSE does not distinguish the normative demand arising from different levels of probability under uncertainty, and acquiring S_L either leads to the desired state of certainty (and therefore lowers his normative demands), or to no change at all—an attractive proposition.

4.2. Binary fairness points

The model can be adjusted to account for coarse-grained norms of subjective compliance that only know two different prescriptions by effectively removing a separate fairness point under uncertainty. The crucial question is whether uncertainty comes with a normatively demanding or an undemanding action. In the first case, $\phi_U = \phi_H < \phi_L$: As long as there is reasonable doubt about the state being L , the normatively demanding action shall be performed (such as keeping an underperforming employee on the payroll). Once the threshold of reasonable doubt is passed, a normatively less demanding action is permitted (laying off the worker). Within our model of dictator giving, a subjective complier would keep $y_U^*(p) = y_H^*$ as long as $p \leq 1 - \epsilon$ and $y_L^* > y_H^*$ for $p > 1 - \epsilon$, with p the belief that the state is L .

In the second case, $\phi_U = \phi_L > \phi_H$. In this case, agents are only expected to perform the normatively demanding action if there is (near) certainty about the state being H . As long as there is reasonable doubt, agents can feel entitled to perform the action that benefits themselves. In terms of dictator giving, a subjective complier would keep $y_U^*(p) = y_L^*$ as long as $p > \epsilon$ and would keep $y_H^* < y_L^*$ otherwise ($p \leq \epsilon$).

It follows straightforwardly from proposition 2 that subjective compliers have different incentives for signal acquisition in these two cases. When $\phi_U = \phi_H$, an agent who is uncertain about the state of the world will prefer

to acquire signal S_L (to possibly reach state L) and will be indifferent between acquiring S_H or not (as ϕ_H will be demanded for all outcomes of signal S_H). When $\phi_U = \phi_L$, an agent will avoid acquiring signal S_H as it can only make her worse off. This type of agent would be indifferent between acquiring signal S_L or not as there is nothing to gain (but also nothing to lose).

4.3. Heterogeneity in subjective fairness points

385 **Convinced, carefree and troubled egoists**

We will now allow for heterogeneity in personal fairness points and in the relation between marginal utility of money and marginal disutility of norm violation. This enables us to represent egalitarian and selfish behavior and suggests some novel implications.

390 The first important case is $\phi_H = \phi_L = \bar{y}$. An agent with such selfish fair points is convinced that the receiver is undeserving and will hence keep everything for himself in all possible epistemic or factual states. In the context of a dictator game, this agent behaves like a selfish ‘homo economicus’. However, such an agent does care about norms—he suffers cognitive dissonance from norm violation, i.e. if he does not get everything himself.¹⁵ By contrast, a standard ‘homo economicus’ who only cares about his own income can be modeled as suffering no dissonance over the entire action space: $f(\Delta) = 0$. We might call this type of agent a ‘carefree egoist’. A carefree egoist does not care about any deviation from the fair points and, for that matter, may not even have any relevant normative beliefs. For both convinced and carefree egoists, and regardless of whether norm compliance is understood in the objective or the subjective way, it is

¹⁵Going beyond a dictator game, such a ‘convinced egoist’ would therefore be willing to invest resources to ‘correct’ such an ‘unjust’ distribution. Hence, convinced egoists are willing to engage in ‘spiteful’ behavior to reduce someone else’s income. This implication, of course, generalized beyond this particular set of fairness points. All norm compliers in our model are in principle willing to reduce someone else’s income if this income is higher than their fair share.

meaningless whether the state of the world is H or L . Therefore, egoists are indifferent between acquiring or not acquiring any type of signals.

405 A third type of selfish agent is the ‘troubled egoist’. This agent does believe that she ought not keep everything herself but her marginal disutility from norm violation is smaller than the marginal utility of money over the entire action space. Such an agent would therefore optimally set $y^* = \bar{y}$ but nevertheless suffers cognitive dissonance from norm violation. If the
 410 agent’s fairness points are ordered such that $\phi_L > \phi_H > 0$, the troubled egoist would, in contrast to the convinced and careless egoist, strictly prefer to be in state L than in state H because of the lower disutility from norm violation. Consequently, in contrast to both previous types of egoist, propositions 1 and 2 apply.

415 **Egalitarians: pure and impure**

At the other end of the spectrum we can also model two different types of egalitarians, ‘pure egalitarians’, who religiously follow their egalitarian conviction and ‘impure egalitarians’, who fall short of their egalitarian normative beliefs. For this section we allow the dissonance function to be
 420 kinked at zero (and therefore non-differentiable at that point). Both types of egalitarians have fairness points $\phi_H = \phi_L = \bar{y}/2$. They come apart in the relative size of the marginal utility of money compared to the marginal disutility from norm violation at the equal split ($y = \bar{y}/2$). At this point (and therefore over the relevant action space of $y \in (\bar{y}/2, \bar{y})$), pure egalitari-
 425 ans have a higher marginal disutility from norm violation than the marginal utility of money: $v'(\bar{y}/2) < \frac{\partial f(\bar{y}/2)}{\partial y}$. Such an agent will therefore split the pie equally between himself and the receiver: $y^* = \bar{y}/2$. ‘Impure egalitarians’, by contrast, have an interior solution to their maximization problem (with $v'(\bar{y}) < \frac{\partial f(\bar{y})}{\partial y}$ and $v'(\bar{y}/2) > \frac{\partial f(\bar{y}/2)}{\partial y}$) such that $y_H^* = y_L^* = y_U^*(p) > \phi_H = \phi_L$.

430 These types of egalitarians are indifferent towards acquiring signals, regardless of whether they are subjective or objective norm compliers, just like the convinced and carefree egoists.

5. An Alternative Explanation of Strategic Ignorance

435 DWK (2007) provide experimental evidence that agents often prefer to remain ignorant about the negative consequences of their selfish acts. We will show that our model offers a parsimonious explanation of their results.

The DWK subjects play a binary version of the dictator game. Dictators can choose between actions A and B to determine the payoffs for them and
440 their receiver. In the baseline treatment, A results in distribution (6,1), B in (5,5). Here, 74% of dictators choose the fair and efficient option B. By contrast, in the hidden information treatment, the outcomes are assigned in two different ways with equal probability: either A causes (6,1) and B (5,5), or A causes (6,5) and B (5,1). By default, it is unclear whether A or B
445 hurts the receiver, but, importantly, the dictator can optionally resolve this uncertainty costlessly by clicking on a button. Almost half of the dictators, however, deliberately remain ignorant and 86% of those choose the payoff-dominant action A, even though there is a chance of $\frac{1}{2}$ to impose a severe negative externality on the receiver.

450 We call the state in which A causes a negative externality H and the state in which it does not L , with p as the Bayesian belief that L is the case. For subjective compliers, DWK's setup can be modeled with a social norm that has only two fairness points:

$$\text{prescription} = \begin{cases} \phi_L(\text{do A}) & \text{if } p > \epsilon \\ \phi_H(\text{do B}) & \text{if } p \leq \epsilon \end{cases}$$

If p is greater than ϵ , uncertainty comes with the normatively undemand-
455 ing prescription ϕ_L , as in the second case of section 4.2. Such a binary social norm can explain the seeming discrepancy in the behavior of the dictators in the baseline and hidden information treatment. Many dictators behave consistently with such a binary social norm and choose action

B under certainty about the state being H ($p = 0$). Under uncertainty,
460 by contrast, subjective compliers decide to remain ignorant and enjoy the
monetary benefits of choosing A. Revealing the game’s payoff structure is
akin to acquiring a perfectly informative signal that will reveal the state
of nature being either H ($p = 0$) or L ($p = 1$). Therefore, by clicking on
the button that is being offered, the dictator can only lose: either he will
465 have to choose B or he stays put where he would be under uncertainty
anyway—being permitted to choose A.¹⁶

DWK do not provide a formal model for the behavior they observe, but
they suggest that dictators have an ‘illusory preference for fairness’ and
that many of them only want to appear fair, either to themselves or to
470 others. More generally, DWK view ignorance as a particularly desirable
state for dictators, as their selfish behavior can be hidden (from others
or themselves). We offer a different explanation: In our model, ignorance
is not used to hide at all—in the DWK setup it simply turns out to be
an attractive epistemic state for subjective norm compliers because it is
475 undemanding. Our explanation of strategic information acquisition also
does not rely on the assumption of a dual self or on uncertainty as a self-
deception tool; rather, signal choices can be modeled as a deliberate and
rational process with perfect knowledge about one’s preferences.

In addition, our model offers an interesting interpretation for the sizable
480 fraction of dictators in DWK who choose to reveal: they can be modeled as
objective norm compliers, for whom the prescription demands implement-
ing the fair outcome regardless of whether they know which action pro-
duces this outcome. A ‘fair’ objective complier, who would choose fairly in
the baseline treatment with complete information, strictly prefers to know
485 which action produces the fair outcome in the hidden information treat-
ment and will act accordingly. The model also allows for a ‘selfish’ objective
complier who would choose the payoff-dominant option in the baseline con-

¹⁶Note that we do not have to assume ϵ to be particularly small in order to explain the
ignorance of dictators, as p is $\frac{1}{2}$ under uncertainty in DWK’s experiments.

dition with complete information. For this type of dictator, the dissonance from norm violation is smaller than the gain in monetary utility. Such a
490 ‘selfish’ objective complier is indifferent between revealing the game’s pay-off structure or not and would not react to the signal outcome (see the appendix for a proof).¹⁷

6. Experimental test

6.1. Hypotheses

495 Our new experimental setup is designed to test whether individuals strategically *seek* (as well as avoid) normatively relevant information. All studies known to us have only ever offered, in the terminology of our model, a signal similar to S_H but never a signal akin to S_L . This has led to the belief that ‘wiggle room’ is intimately linked to the option of remaining uncertain
500 or ‘strategically ignorant’. In marked contrast, our model suggests that getting some information, but not all information, is often the best way to ‘wiggle’.

The theoretical underpinning is provided by proposition 2, which is a generalization of claims pertaining to strategic ignorance. To test proposition
505 2, our experimental design explicitly implements a coarse-grained norm, distinguishing it from all previous experimental work in this area. The design also uses the central elements of the model in section 3: a dictator game embedded in social norms of equity such that the receivers are likely to be perceived as either deserving or undeserving by the dictators; and,
510 crucially, an opportunity for dictators to acquire signals S_H and S_L , which, in expectations, reduce the dictator’s uncertainty about the deservingness of the receiver.

¹⁷Our model can therefore also capture the fraction of dictators who choose to reveal and then choose the selfish option (25% of those revealing in DWK’s data). The difference to proposition 1, which stipulates a strict preference for information, stems from DWK’s design, in which dictators can only choose two options. In the more general case of inner solutions to dictator’s allocation problems, objective compliers strictly prefer more information (as captured by proposition 1).

The null hypothesis H0 states that dictators will acquire S_L and S_H equally often. H0 can be derived in the form of full information acquisition
515 from theories that assume genuine preferences for fair outcomes; it is also captured by proposition 1, based on the assumption of objective norm compliance. Note that H0 is also consistent with ‘strategic ignorance’ and with proposition 3 of our model for the case of a linear norm under uncertainty. The novel part of our experimental test is the alternative hypothesis H1,
520 derived from proposition 2, based on the assumption of subjective compliance to a coarse-grained norm. It states that dictators will seek information strategically by acquiring S_L but not S_H .

6.2. Experimental stages

We now explain the stages of the experimental study (a more expansive
525 account is provided in appendix F). The subjects in the main study (which we simply call ‘experiment’) play a normatively embedded dictator game in which the dictators decide on the distribution of 20 Euros. First, and prior to the experiment, we elicit normative beliefs within the same subject pool. Specifically, the subjects of this norm elicitation session receive a de-
530 scription of the setup of the experiment and are asked how much a dictator ought to give to a high performer and a low performer (see appendix G for details). According to the modal responses in the norm elicitation session, dictators ought to give 10 Euros to a high performer (state H in the model) and 5 Euros to a low performer (state L in the model). We thereby anchor
535 the personal fair points of all dictators at $\phi_H = 5$ and $\phi_L = 10$. This greatly reduces normative uncertainty and allows us to study the consequences of factual uncertainty, the focus of this paper.

Second, in the experiment we create two types by playing a competitive
540 knowledge quiz. All subjects answer knowledge questions taken from ‘Who Wants to be a Millionaire’ under time pressure. The best 75% performers are declared ‘high performers’, the lowest 25% ‘low performers’. All subjects are informed that doing well in the quiz (i.e., being a high performer)

makes it more likely to (i) be a dictator in a subsequent dictator game; and (ii) to win a ‘bonus’ of 20 Euro that is available for distribution between the dictator and a receiver later on. We then assign dictator and receiver roles such that all dictators are high performers, while receivers are, in equal shares, high and low performers.

Third, before the dictator-receiver pairs are formed, the dictator game is played with a strategy method. More precisely, all dictators submit a strategy of how much to give:

- (i) in case they learn they are paired with a low performer; and
- (ii) in case they learn they are paired with a high performer.

When entering the strategy, the dictators do not know whether or under which circumstances information about the type of their receiver might become available to them, but they are told that their strategy choice is binding. Before choosing the strategy, we inform the dictators that if they remain uncertain about the type of their receiver, the mean of the two stated amounts will be transferred to the receiver. After entering the strategy information, the dictators are paired with equal probability with either a high performer or a low performer, but they do not learn by default which type of receiver they are paired with.

Fourth, the dictators have an unannounced opportunity to acquire information that may inform them about the type of their receiver. To make this optional information uptake intuitively plausible, we tell the dictators that the information about the (lack of) deservingness of their receiver is contained in exactly one of four envelopes symbolically displayed on screen. If the receiver is a high performer, the information is in one of two envelopes called ‘gold envelopes’. If the receiver is a low performer, the information is in one of two envelopes called ‘silver envelopes’. The subjects can open up to one envelope of each type. More formally, the signals available are S_L and S_H , as described above. That implies four different sets of signals can be chosen: $\{\}$, $\{S_L\}$, $\{S_H\}$, or $\{S_L, S_H\}$. A dictator wishing to obtain

as much information as possible will open one envelope each, a dictator who only wants to increase the chance of learning that the receiver is a low performer will only open a silver envelope, and so on.¹⁸

Finally, the dictator game is implemented. Whether a dictator-receiver pair gets a bonus of 20 Euro for distribution depends on the type of the receiver: the bonus is always provided if the receiver is a high performer but only with probability $\frac{1}{2}$ if the receiver is a low performer. This fact, which the subjects were informed about at the start of the session, underscores the distinction between deserving and undeserving receivers: only being paired with a high performer increases the chance to win a bonus. At the end of the treatment, the bonus (if available) is distributed according to the strategy of the dictator and the information the dictator obtained about the receiver. All parameters and stages in the experiment apart from the information acquisition are common knowledge among the subjects.

6.3. Experimental strategy

Our experimental strategy differs from previous investigations of ‘strategic ignorance’ because we test different hypotheses. The possibility to reject H_0 hinges on two conditions: First, we must be able to observe the signal choices of dictators with personal fairness points $\phi_L^i > \phi_H^i$ (with i now indicating the individual dictator). Second, a coarse-grained norm of desert under uncertainty (instead of, for instance, a linear mapping of prescriptions to beliefs) must be induced successfully. In this section we explain our use of the strategy method and our method for creating or reinforcing a coarse-grained norm.

We establish a social norm of appropriate giving as a function of desert through the instruments described in the previous section. We thereby emphasize that deserving and undeserving receivers ought to be treated

¹⁸In terms of model parameters, the prior probability for the types is therefore $p = \frac{1}{2}$ and the probability of resolving uncertainty when choosing the ‘correct’ signal is $s = \frac{1}{2}$.

600 differently. But despite these efforts, there is still likely to be heterogeneity among dictators. In our setup, a social norm of equity conflicts with a social norm of equality, which suggest fairness points of $\phi_H^i = \phi_L^i = 10$. In addition, some experimental subjects may also be ‘convinced’ or ‘care-free egoists’ (see section 4.3), who believe they deserve to keep everything themselves, with $\phi_H^i = \phi_L^i = \bar{y}$.¹⁹ Therefore, we measure dictator compliance with the exogenously introduced social norm of equity by employing a variant of the strategy method (Selten, 1967) and then offer an opportunity to acquire signals. The strategy method is a standard experimental tool for an economical data-collection process, to lead subjects to make thoughtful 610 decisions and, important for our purposes, to gain better insights into the motives underlying decisions (Brandts and Charness, 2011).²⁰ In particular, the strategy method allows us to separate those with and without incentives to strategically acquire information. Two measures are taken to ensure truthful revelation of allocation intentions: First, dictator allocation 615 strategies are binding and cannot be changed after acquiring information. Second, when the dictators enter their strategies, they do not yet know about the signals that become available later. This ensures that dictators treat both epistemic states H and L as equally relevant when deciding over the allocation of money.

620 We implement the coarse-grained norm by restricting the action space of dictators under uncertainty. We inform the dictators that if they remain uncertain about the type of their receiver, the mean of the two stated amounts

¹⁹While our design gives very little reason for this belief from an ‘intellectually honest’ perspective, previous research has shown that fairness perceptions are often selfishly biased (e.g., Babcock et al., 1995; Konow, 2000).

²⁰In the model, when dictators decide over the signal acquisition, they first calculate expected utilities of signal acquisition by solving for the optimal allocation in the different epistemic states. By first eliciting the dictator giving strategy, we gain insights into part of this decision process, which would otherwise be a completely black box. While the strategy method should yield the same observations than the direct-response method according to standard game-theory, it has been criticized on the behavioral grounds that a ‘cold’ strategy might be systematically different than a ‘hot’ direct response. The Brandts and Charness (2011) meta-study largely alleviates such concerns (apart from studying punishment behavior, which is more pronounced under the direct-response method).

for states H and L will be transferred to the receiver: $y_U = (y_L + y_H)/2$.
With the amount for uncertainty fixed externally, the dictators cannot
625 change their giving continuously as a function of their belief about the
receiver type. This makes following a linear norm impossible. More impor-
tantly, it also provides a strong normative justification for giving the same
amount under uncertainty even if beliefs change, inducing the crucial flat
part of the function ϕ_p , which drives proposition 2. In addition, by setting
630 y_U in between y_L and y_H this method mirrors a symmetric norm under un-
certainty, which seems a natural starting point to test for subjective norm
compliance, as it ensures that the monetary incentives of acquiring either
signal is the same in absolute terms for all dictators. Note therefore that
our experiment is not designed to test for the prevalence of coarse-grained
635 norms (this is left for future research), nor do we focus on how much dicta-
tors give under uncertainty about the receiver’s deservingness. Instead, we
want to test whether making a coarse-grained norm salient can be a source
for both strategically *ignoring and seeking* information.

6.4. Procedures

640 Our subjects were recruited with the online recruitment system ORSEE by
Greiner (2004) from the student subject-pool of the Cologne Laboratory for
Economic Research at the University of Cologne (CLER). Subjects had not
previously participated in dictator-game or normative-choice experiments.
However, all subjects had some previous experience with laboratory ex-
645 periments. We first ran one survey session with 26 participants to elicit
normative beliefs. Subsequently, we ran three experimental sessions with
32 participants each, resulting in 48 independent observations of dictator
behavior. Subjects took part in only one session and assumed only one role.
General instructions about the experiment were provided on paper (see ap-
650 pendix H). The summary part of the instructions was also read aloud to the
subjects with two PowerPoint slides facilitating understanding. All subse-
quent interactions took place at computer terminals in cubicles, controlled

with z-Tree (Fischbacher, 2007). Anonymity was guaranteed by ensuring that subjects were randomly matched and by prohibiting communication
655 between the subjects during the experiment. Average payments in the main experiment were 14.3 Euros, close to the expected value of 14.5 Euros. In addition to the expected earnings of 15 Euros per dictator-receiver pair, each subject received 5 Euros show-up fee and 2 Euros for completing a post-experimental questionnaire. Sessions lasted on average 90 minutes.

660 **6.5. Results**

Before proceeding to the signal acquisition choices, the main variable of interest, we first look at the dictators' allocation strategies. The creation of a wedge in entitlements, in line with the model's assumptions, is successful: dictators commit to giving, under certainty, significantly more to
665 a high performer than to a low performer ($p < 0.001$, Wilcoxon signed-rank test).²¹ As can be seen from Figure 1, no dictator would give more to a low performer than to a high performer. Note that the number of horizontally clustered dots in Figure 1 indicates the frequency of the coordinate, each data-point corresponding to the transfer strategy of one dictator. On
670 average, dictators commit to giving a substantial amount to both types of receivers. The mean contribution to low performers is about 3.7 and to high performers approximately 6.6.²² However, Figure 1 also shows that there is a substantial number of dictators who do not differentiate between high and low performers (visually represented by the observations on the
675 45 degree line). These dictators are apparently not receptive to the norm of desert we tried to make salient. Among these, two types (see discussion in section 4.3), stand out: 'egoists', who keep everything for themselves, and 'pure egalitarians', who split the pie equally with any type of receiver. We

²¹All statistical tests are two-sided.

²²Note that the significantly ($p < 0.001$, Wilcoxon signed-rank test) closer alignment of dictator-giving with normative prescriptions in state L than in state H is also predicted by our model. On average, dictators fall short of what is normatively required by only 1.3 in state L but by 3.4 in state H .

asked our subjects in a post-experimental questionnaire about the motives
of their giving decision. Not all subjects stated clear reasons, but the an-
680 swers nevertheless shed some light on the difference between differentiating
and undifferentiating dictators. Specifically, most of the 11 dictators who
mentioned considerations of equality and selfishness did not differentiate.²³
By contrast, the 21 subjects who mentioned either entitlement/desert or
685 norms as motives (while neither mentioning equality nor selfishness) dif-
ferentiated more strongly in their giving decision (average difference 4.7
Euros). This shows that subjects motivated by egalitarian values or self-
ishness tended to reject the norm we instilled, while others were receptive
to it.

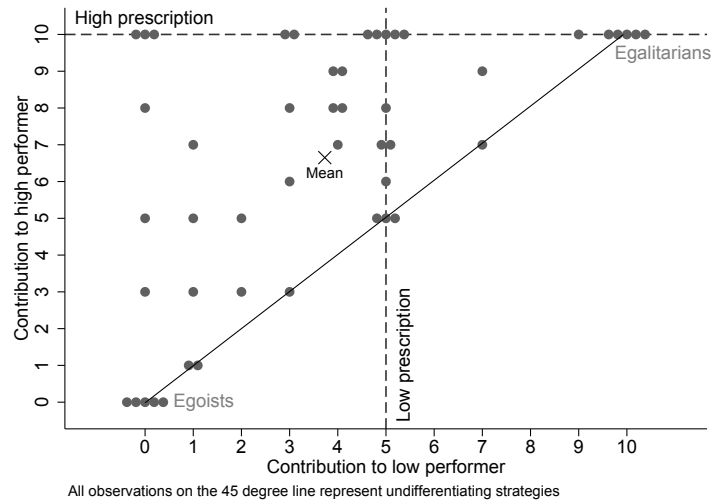


Figure 1: Dictator transfer strategies.

690 When analyzing dictators' signal choices, we distinguish between dicta-
tors with differentiated and undifferentiated giving strategies. The former
give more when they learn that they are paired with a high performer than
a low performer, the latter give the same amount. Figure 2 depicts the
information acquisition choices of the dictators with differentiated strate-
695 gies in light gray and of the dictators with undifferentiated strategies in

²³6 subjects mentioned considerations of equality, 4 selfishness, and 1 both equality
and selfishness. Among these 11 subjects, 9 had completely non-differentiated giving
strategies, the other 2 gave only 2 Euros more in case they learn they are paired with
a high performer.

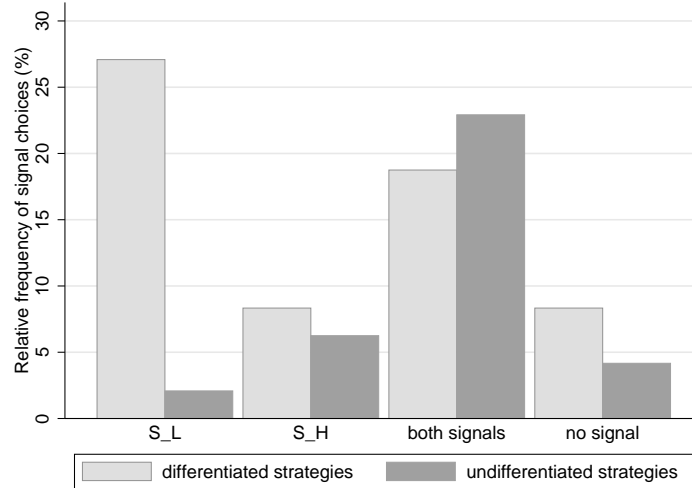


Figure 2: Distribution of information acquisition choices for subjects with differentiated and undifferentiated strategies.

dark gray. The signal choice of one self-reportedly confused subject is excluded, which leaves us with 30 dictators with differentiated and 17 with undifferentiated strategies.²⁴

As the hypotheses are based on the assumption of a wedge in entitlements, we first focus on the dictators with differentiated strategies. For this group of dictators, S_L is the modal choice of signals (43.3%), in line with H1. Acquiring both signals accounts for 30% of types of signal choices. Based on descriptive statistics, subjective compliance therefore organizes the data better than objective compliance or notions of ‘genuine fairness preferences’. The null hypothesis of equally frequent choices of S_L and S_H is rejected at a significance level of $p=0.029$ (Wilcoxon signed-rank test).

The rejection of H_0 is driven by the marked difference in selective information acquisition: There are three times as many dictators who only chose S_L than dictators who only chose S_H . As acquiring both signals is the second most frequent choice, this type of behavior cannot be dismissed easily. At first sight, this seems to indicate the presence of objective compliers among the dictators. However, Figure 2 also clearly shows that acquiring

²⁴The subject stated in the post-experimental questionnaire of having mixed up her or his choices.

	(1)	(2)	(3)	(4)
VARIABLES	S_L and S_H	S_L	S_H	no signal
y_H^*		0.434*** (0.16)	0.0953 (0.21)	0.112 (0.19)
y_L^*		-0.705** (0.29)	-0.278 (0.25)	-0.209 (0.23)
Constant		-0.746 (0.84)	-0.510 (0.86)	-1.002 (0.98)
Observations	47	47	47	47

Robust standard errors in parentheses

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 1: Multinomial logistic regression to explain signal choices.

both signals becomes the overwhelming choice for dictators with undifferentiated strategies. While this group of dictators is not a randomly selected control group, their behavior nevertheless suggests that getting as much information as possible is the default choice when no economic incentives are at stake.²⁵ The motivation behind this can be compliance with an epistemic norm to acquire as much information as possible or, plainly, curiosity.²⁶ The marked and statistically significant ($p=0.008$, Mann-Whitney U test) jump in the choice of S_L when comparing dictators with and without economic stakes makes the evidence for subjective compliance in line with the alternative hypothesis 1 even stronger: strategic information acquisition is virtually non-existent for the latter, but makes up the largest share of information choices for the former.

The role of incentives for selective information acquisition is also evident

²⁵Let us remind the reader, however, that the group of non-differentiating dictators is very heterogeneous and consists of pure ‘egoists’ as well as ‘egalitarians’.

²⁶There is some evidence, however, that the choice of acquiring both signals follows different processes when comparing those with and without economic incentives: decision times are, on average, considerably and weakly significantly longer for the former than the latter (25.9 vs 14.6 seconds, $p=0.073$, Mann-Whitney U test). This suggests that acquiring both signals is the result of a deliberative process for the group of dictators with differentiated strategies; curiosity alone might not be a good explanation for this group’s choice of signals. Future research may try to further distinguish objective compliance from other motivations to acquire maximum information.

from the multinomial regression in table 1, which takes the choice of both signals as the reference category. A dictator is more likely to choose only S_L the more differentiated the giving strategy is: The probability to choose only S_L rises in the contribution to the high performer and decreases in
730 the contribution to the low performer. Recall that a dictator can, in monetary terms, potentially gain $y_L - y_U = y_L - \frac{y_L + y_H}{2}$ by acquiring S_L and could potentially lose $y_U - y_H = \frac{y_L + y_H}{2} - y_H$ by acquiring S_H . Qualitative responses in the post-experimental questionnaire are in line with the choice analysis: 7 out of the 10 differentiating dictators who choose only
735 S_L and provided reason for the signal choices mention a monetary motivation (one dictator also directly states that choosing only S_L allows a larger payoff without having a bad conscience). Neither the regression nor the post-experimental questionnaire shed any light on the reasons behind the remaining two types of information acquisition, choosing only S_H and
740 choosing no signal. The latter type of behavior (‘complete ignorance’)—while not predicted for our experimental setup—*could* be captured within the model by assuming that these dictators follow a linear norm under uncertainty and were consequently not receptive to the coarse-grained norm we tried to implement. The choice of only S_H by differentiating dictators
745 (anything can happen with undifferentiating strategies) is outside any theoretic framework. The inconclusive statements in the post-experimental questionnaire as well as their low frequency makes mistakes or behavioral noise the most likely reason for these types of signal choices.

To take stock, we ask: How well does our and some relevant competing
750 theories explain the results? To begin with, models of purely outcome-based preferences account for the results badly, as they are inconsistent with both core findings. First, most subjects differentiate between the different types of receivers. Second, strategic information acquisition is inconsistent with all fairness preference models, and that includes context-sensitive theories
755 of entitlement and desert. If subjects want their *own* allocation strategies to

be realized, they should choose both signals.²⁷ The strategic signal choices, by contrast, provide support for our model’s core prediction.

7. Summary and discussion

Depending on one’s objective or subjective interpretation of norms, one
760 can be subject to different forms of psychological costs when violating a
norm. In line with Festinger (1957) and Konow (2000), we model these
costs as *cognitive dissonance* that arises when acts do not match what the
norm requires. Following the subjective interpretation, individuals expe-
765 rience cognitive dissonance if their beliefs and the norm together imply
prescriptions they violate. When an agent is a subjective complier, she can
strategically choose the sort of information that might render selfish actions
morally appropriate. By contrast, an objective norm complier is better off
with more information about the state of the world as this improves his
chance to choose the morally appropriate action.

770 Our model is widely applicable and can, for instance, explain strategic
ignorance in dictator games, as found by Dana, Weber and Kuang (2007–
DWK) and others. DWK interpret their results as evidence for an ‘illusory
preference for fairness’. They view ignorance as a desirable state that comes
with lower normative demands or allows to hide selfishness (from others or
775 oneself). In our model, subjective compliers do not strive for ignorance
as such; instead they decide with their signal choice whether and which
lotteries to play over their beliefs about the state of the world. The signals
offered by DWK (and others) just weren’t very attractive, as normative
demands could only increase. Contrary to previous explanations, we do
780 not need to assume self-deception or a dual self. Our model locates the
source of the ‘moral wiggle room’ in the strategic use of coarse-grained
norms under conditions of uncertainty.

²⁷See Feiler (2014) for an elaboration of why outcome-based expected utility models are inconsistent with strategic information acquisition.

A new experimental test provides evidence for strategic information acquisition in line with subjective norm compliance. By giving our subjects
785 the chance to selectively avoid and acquire information, we find empirical support for our hypothesis that subjects engage in self-serving information acquisition (not just avoidance) to reduce norm compliance costs. In light of this, recent proposals to use social norms as a subtle yet effective policy tool are perhaps not quite as promising as one might think. Rather than
790 making individuals comply with what the norm prescribes, they can make individuals look for information in a selfishly biased way to make the norm subjectively less applicable to them.

Which debiasing tools are available? Our model suggests two interventions that would work but come with significant drawbacks. First, one
795 could promote norm compliance in the ‘objective’ sense. This eliminates the ‘wiggle room’. However, a purely ‘objective view’ is difficult to defend because it implausibly expects individuals to do what the state of the world demands, even if they have no opportunity to know what this state is. A second, equally unpromising intervention would promote subjective compli-
800 ance with a norm that is linearly responsive to beliefs. As we have shown, a linear social norm indeed eliminates the bias in information acquisition. However, it is difficult to imagine how such a norm could be enforced since degrees of beliefs are rarely observable in detail. A linear social norm also comes at the considerable cost of eliminating all incentives for information
805 acquisition.

Fortunately, there are interventions that hold more promise. One such intervention consists of emphasising ‘meta-norms’ of information acquisition. According to one such meta-norm, individuals are expected to use
810 all information that is obtainable with reasonable effort before taking important decisions (such as, following our running example, asking both the foe and the friend for reference before deciding on the bonus payment). One can understand these meta-norms as rules to determine how diligently individuals have to gather evidence before subjective compliance becomes

permissible. Last but not least, one could also aim to make norms more
815 demanding under uncertainty. In particular, if a selfish action carries so-
cial risk, social norms could prescribe a pro-social action unless there is
sufficiently strong evidence that no such risk is present. If so, individu-
als have no incentive for selfishly biased information acquisition as further
information can only reduce compliance costs.

820 Finally, our model and the supportive experimental results suggest some
directions for future research. We believe that the investigation of ‘moral
wiggle room’ should move away from focusing on ‘ignorance’ only. In addi-
tion, the recent emphasis on self-deception or dual selves, fascinating as the
results certainly are, might have diverted attention from a simpler expla-
825 nation: the imperfect incentives created by social norms under conditions
of uncertainty, which some smart individuals are bound to exploit.

References

- Akerlof, G. A. and Dickens, W. T. (1982). The Economic Consequences of
Cognitive Dissonance. *The American Economic Review*, 72(3):307–319.
- 830 Babcock, L., Loewenstein, G., Issacharoff, S., and Camerer, C. (1995).
Biased judgments of fairness in bargaining. *The American Economic
Review*, 85(5):1337–1343.
- Bénabou, R. and Tirole, J. (2006). Incentives and Prosocial Behavior. *The
American Economic Review*, 96(5):1652–1678.
- 835 Bénabou, R. and Tirole, J. (2011). Identity, Morals, and Taboos: Beliefs
as Assets. *The Quarterly Journal of Economics*, 126:805–855.
- Bicchieri, C. (2006). *The Grammar of Society: The Nature and Dynamics
of Social Norms*. Cambridge University Press, Cambridge, MA US.
- Bicchieri, C. and Chavez, A. K. (2013). Norm Manipulation, Norm Evasion:
840 Experimental Evidence. *Economics and Philosophy*, 29(02):175–198.

- Bicchieri, C. and Mercier, H. (2013). Self-serving biases and public justifications in trust games. *Synthese*, 190(5):909–922.
- Bodner, R. and Prelec, D. (2003). Self-Signaling and Diagnostic Utility in Everyday Decision Making. In Brocas, I. and Carillo, J. D., editors, *The Psychology of Economic Decisions*, pages 105–126. Oxford University Press, Oxford.
- Brandts, J. and Charness, G. (2011). The strategy versus the direct-response method: a first survey of experimental comparisons. *Experimental Economics*, 14(3):375–398.
- Conrads, J. and Irlenbusch, B. (2013). Strategic ignorance in ultimatum bargaining. *Journal of Economic Behavior & Organization*, 92(1956):104–115.
- Dana, J., Weber, R. A., and Kuang, J. X. (2007). Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness. *Economic Theory*, 33(1):67–80.
- Feiler, L. (2014). Testing models of information avoidance with binary choice dictator games. *Journal of Economic Psychology*, 45:253–267.
- Festinger, L. (1957). *A Theory of Cognitive Dissonance*. Stanford University Press, Palo Alto.
- Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178.
- Greiner, B. (2004). An online recruitment system for economic experiments. In Kremer, K. and Macho, V., editors, *Forschung und wissenschaftliches Rechnen 2003*, number 63 in GWDG Bericht, pages 79–93. Ges. für Wiss. Datenverarbeitung, Göttingen.
- Grossman, Z. (2010). Strategic Ignorance and the Robustness of Social Preferences. Available at: <http://escholarship.org/uc/item/60b93868>.

- Grossman, Z. and van der Weele, J. J. (2013). Self-Image and Strategic Ignorance in Social Decisions. Available at: <http://ssrn.com/abstract=2237496>.
870
- Hechter, M. and Opp, K.-D. (2001). Introduction. In Hechter, M. and Opp, K.-D., editors, *Social Norms*, pages xi–xx. Russell Sage, New York.
- Konow, J. (2000). Fair Shares: Accountability and Cognitive Dissonance in Allocation Decisions. *The American Economic Review*, 90(4):1072–1091.
875
- Larson, T. and Capra, C. M. (2009). Exploiting moral wiggle room: Illusory preference for fairness? A comment. *Judgment and Decision Making*, 4(6):467–474.
- Matthey, A. and Regner, T. (2011). Do I Really Want to Know? A Cognitive Dissonance-Based Explanation of Other-Regarding Behavior.
880 *Games*, 2:114–135.
- Rabin, M. (1995). Moral preferences, moral constraints, and self-serving biases. Available at: <http://escholarship.org/uc/item/97r6t5vf>.
- Rodriguez-Lara, I. and Moreno-Garrido, L. (2012). Self-interest and fairness: self-serving choices of justice principles. *Experimental Economics*,
885 15:158–175.
- Selten, R. (1967). Die Strategiemethode zur Erforschung des eingeschränkt rationalen Verhaltens im Rahmen eines Oligopolexperiments. In *Beiträge zur experimentellen Wirtschaftsforschung*, pages 136–168. Mohr, Tübingen.
890
- Varian, H. E. (1992). *Microeconomic Analysis*. W. W. Norton, New York, 3rd edition.
- Zimmerman, M. (2008). *Living with uncertainty: the moral significance of ignorance*. Cambridge University Press Cambridge, Cambridge.

Appendix

A. Proof that \hat{u} is increasing and concave

To show that $\hat{u}'(\phi) > 0$, note that, by construction of the optimal giving choice, $y^* \geq \phi$. Applying the envelope theorem we obtain

$$\begin{aligned}\hat{u}'(\phi) &= \left. \frac{\partial}{\partial \phi} u(y, \phi) \right|_{y=y^*} \\ &= \frac{\partial}{\partial \phi} (v(y^*) - f(y^* - \phi)) \\ &= -\frac{\partial}{\partial \phi} \Delta^* \frac{\partial}{\partial \Delta^*} f(\Delta^*), \text{ with } \Delta^* = y^* - \phi \\ &= \frac{\partial}{\partial \Delta^*} f(\Delta^*) = f'(\cdot) > 0\end{aligned}$$

Note that $\frac{\partial}{\partial \phi} \Delta^* = -1$ because $\frac{\partial}{\partial \phi} y^* = 0$.

900 To see that \hat{u} is concave, we calculate the second derivate:

$$\begin{aligned}\hat{u}''(\phi) &= \frac{d\Delta^*}{d\phi} \frac{d}{d\Delta^*} f'(\Delta^*) \\ &= \left(\frac{dy^*}{d\phi} - 1 \right) f''(\cdot) < 0\end{aligned}$$

The first factor is negative because of the decreasing marginal utility of money, which means that an increase in ϕ by one unit leads to an increase in y^* by less than one unit. This shows that \hat{u} is concave.

B. Proof of proposition 1

905 To show that objective compliers acquire both signals, we calculate the maximum expected utility the agent can achieve as a function of belief p . Call this $u^*(p, y_U^*(p)) = \max_y v(y) - pf(|y - \phi_L|) - (1 - p)f(|y - \phi_H|)$,

where $y_U^*(p)$ is the maximizing argument for a given p . According to the envelope theorem, the total derivative at point $y_U^*(p)$ is equal to the partial
910 derivative:

$$\frac{du^*}{dp} = \left. \frac{\partial u}{\partial p} \right|_{y=y_U^*} = f(|y_U^*(p) - \phi_H|) - f(|y_U^*(p) - \phi_L|).$$

The second derivative with respect to p is

$$\begin{aligned} \frac{d^2u^*}{dp^2} &= \text{sign}(y_U^*(p) - \phi_H) y_U'^*(p) f'(|y_U^*(p) - \phi_H|) \\ &\quad - \text{sign}(y_U^*(p) - \phi_L) y_U'^*(p) f'(|y_U^*(p) - \phi_L|). \end{aligned}$$

Note that by the assumption of an interior solution to the dictator's maximization problem, it must be the case that $y_U^*(p) > \phi_H$, that is, the dictator would keep, at all p , more than the fair amount of giving in state H.
915 Depending on the level of p and the relationship between marginal utility from money and marginal disutility from cognitive dissonance, $y_U^*(p)$ can lie between the two fair points, at ϕ_L or will be larger than ϕ_L . Consider first the case $\phi_H < y_U^*(p) < \phi_L$:

$$\frac{d^2u^*}{dp^2} = y_U'^*(p) f'(y_U^*(p) - \phi_H) + y_U'^*(p) f'(|y_U^*(p) - \phi_L|) > 0.$$

In the case that $y_U^*(p) = \phi_L$ we find:

$$\frac{d^2u^*}{dp^2} = y_U'^*(p) f'(y_U^*(p) - \phi_H) > 0.$$

920 Finally, if $y_U^*(p) > \phi_L$ we find:

$$\frac{d^2u^*}{dp^2} = y_U'^*(p) f'(y_U^*(p) - \phi_H) - y_U'^*(p) f'(y_U^*(p) - \phi_L) > 0.$$

Since $\phi_H < \phi_L$ it is true that $y_U^*(p) - \phi_H > y_U^*(p) - \phi_L$ and since, by assumption, $f' > 0$, it is the case that $f'(y_U^*(p) - \phi_H) > f'(y_U^*(p) - \phi_L)$. $y_U'^*(p) > 0$ because the more likely state L is, the more the dictator will

keep. Since $\frac{d^2 u^*}{dp^2}$ is therefore larger than 0 in all cases, we have established
 925 that $u^*(p)$ is convex. Finally, note that each signal offers a fair gamble over
 this convex function, which will always be accepted, so that both signals
 will be taken by the objective complier.

C. Proof of Proposition 2

To show that subjective compliers following a coarse-grained norm acquire
 930 only signal S_L , we compute the expected utilities of acquiring signals S_L ,
 S_H , both signals, or none. Note that, by assumption, receiving a null
 signal always yields $\phi_p = \phi_U$. This simplifies the expected utilities of the
 respective signal acquisitions:

$$\begin{aligned} Eu_L &= ps\hat{u}(\phi_L) + (1 - ps)\hat{u}(\phi_U) \\ Eu_H &= (1 - p)s\hat{u}(\phi_H) + (1 - s + ps)\hat{u}(\phi_U) \\ Eu_{LH} &= ps\hat{u}(\phi_L) + (1 - p)s\hat{u}(\phi_H) + (1 - s)\hat{u}(\phi_U) \\ Eu_0 &= \hat{u}(\phi_U) \end{aligned}$$

Since $\phi_L > \phi_U > \phi_H$ it follows that $\hat{u}(\phi_L) > \hat{u}(\phi_U) > \hat{u}(\phi_H)$. It is
 935 now obvious that $Eu_L > Eu_H$, $Eu_L > Eu_{LH}$ and $Eu_L > Eu_0$. Therefore
 acquiring only S_L is the best choice.

D. Proof of Proposition 3

To show that subjective compliers following a linear norm, note first that
 when no signal is acquired, ϕ_p remains unchanged at $p\phi_L + (1 - p)\phi_H$
 940 between ϕ_L and ϕ_H . Acquiring signal S_L either leads to ϕ_L with proba-
 bility ps or to an increase in ϕ_p after the Bayesian update to $p' = (1 -$
 $s)p / ((1 - s)p + (1 - p))$. Acquiring signal S_H either leads to ϕ_H with
 probability $(1 - p)s$ or to a decrease in ϕ_p after the Bayesian update to
 $p' = p / (p + (1 - s)(1 - p))$. When acquiring both signals, ϕ_p can go to

945 ϕ_L , ϕ_H , or stay the same (in case of two null signals). As can be checked easily, the expected value $E\phi_p$ for all these (fair) lotteries is just the initial $\phi_p = p\phi_L + (1-p)\phi_H$. However, since \hat{u} is a strictly concave function, $E\hat{u}(\phi_p) < \hat{u}(E\phi_p)$ because of Jensen's Inequality.²⁸ Therefore, the expected utility of acquiring a signal is always lower than the utility obtained when
 950 not acquiring a signal, and thus not acquiring a signal is always preferred.

E. Proof of objective compliers' preference for information in DWK

In DWK's hidden information treatment, dictators are uncertain about the game's payoff structure but know that both possible states of the world
 955 are equiprobable. The utility of an objective complier who does not know which state she is in is:

$Eu_U = v(y_U^*) - 0.5f(y_U^* - \phi_L) - 0.5f(y_U^* - \phi_H)$, with ϕ_L prescribing the undemanding (payoff-dominating) action A and ϕ_H prescribing the demanding action B and the star denoting, as before, the optimal amount to
 960 keep (in this case $y_U^* = \operatorname{argmax}_y(u(y, \phi_L, \phi_H))$).

As revealing the game will resolve uncertainty entirely, the expected utility of doing so is:

$$\begin{aligned} Eu_R &= 0.5 * \hat{u}_H + 0.5 * \hat{u}_L \\ &= 0.5 * (v(y_H^*) - f(y_H^* - \phi_H)) + 0.5 * (v(y_L^*) - f(y_L^* - \phi_L)) \\ &= 0.5 * (v(y_H^*) + v(y_L^*)) - 0.5 * (f(y_H^* - \phi_H) + f(y_L^* - \phi_L)), \end{aligned}$$

with the hat denoting, as before, the maximum utility in the respective state. An objective complier will prefer to reveal if doing so increases
 965 expected utility:

²⁸Let X be a non-degenerate random variable and $f(X)$ be a strictly concave function of this random variable. Then $Ef(X) < f(EX)$ (see e.g. Varian, 1992, 182).

$$\begin{aligned}
Eu_R - Eu_U &= 0.5 * (v(y_H^*) + v(y_L^*)) - v(y_U^*) - 0.5 * (f(y_H^* - \phi_H) \\
&\quad + f(y_L^* - \phi_L)) + 0.5(f(y_U^* - \phi_L) + f(y_U^* - \phi_H)). \quad (1)
\end{aligned}$$

Because of the experimental design, y can only take on two values, depending on which option is chosen. The optimal choice under uncertainty is obviously y_A (as the expected norm violation is the same for both options but A dominates B in terms of payoffs for the dictator).

970 There are therefore 2 meaningful cases:

1. ‘Selfish’ dictator: $y_L^* = y_H^* = y_A$

Expression (1) then simplifies to:

$$\begin{aligned}
Eu_R - Eu_U &= 0.5 * (v(y_A) + v(y_A)) - v(y_A) - 0.5 * (f(y_A - \phi_H) \\
&\quad + f(y_A - \phi_L)) + 0.5(f(y_A - \phi_H) + f(y_A - \phi_L)) \\
&= 0.
\end{aligned}$$

A ‘selfish’ dictator is therefore indifferent between revealing the game structure or remaining uncertain.

- 975 2. ‘Fair’ dictator: $y_L^* = y_A$ and $y_H^* = y_B$

Expression (1) then simplifies to:

$$\begin{aligned}
Eu_S - Eu_U &= 0.5 * (v(y_A) + v(y_B)) - v(y_A) - 0.5 * (f(y_B - \phi_H) \\
&\quad + f(y_A - \phi_L)) + 0.5(f(y_A - \phi_H) + f(y_A - \phi_L)) \\
&= 0.5 * (v(y_B) - v(y_A)) + 0.5 * (f(y_A - \phi_H) - f(y_B - \phi_H)) \\
&= 0.5 * [v(y_B) - f(y_B - \phi_H) - (v(y_A) - f(y_A - \phi_H))].
\end{aligned}$$

By assumption, a ‘fair’ dictator prefers to choose B in state H. Consequently, the expression in square brackets is positive.

This proves that an objective complier will weakly prefer to reveal the game
980 in DWK’s hidden information treatment. More precisely, a dictator who
would choose selfishly (option A) in the baseline condition with complete
information is indifferent between revealing or remaining ignorant. By contrast,
a ‘fair’ dictator who would choose the fair option B in the baseline
treatment strictly prefers to reveal and will choose according to the signal.

985 **F. Details of experimental design**

F.1. Creation of social norms about receiver-entitlements

In order to test for subjective and objective norm compliance, we create
(or make salient) a set of social norms of entitlement whose violation would
create cognitive dissonance. We take several measures to achieve this:

990 *First*, the receivers’ entitlements were manipulated by linking the performance
in the quiz to the chance of a dictator-receiver pair winning the bonus
of 20 Euros. Subjects are informed that a low-performing receiver does
not contribute to the pair’s chance of winning the bonus, whereas a high-
performing receiver contributes as much as the dictator to the pair’s win-
995 ning chances. In the experiment, this was represented by high-performers
contributing a ‘winner lot’, but low-performers only contributing a ‘blank
lot’ to each pair’s bonus draw. The existence of the 20 Euro fund would
be determined by a draw of one lot from the pair playing, with a winner
lot providing funds, a blank lot no funds. A pair of two high-performing
1000 players would always have 2 winner lots, therefore always draw a winner lot
and consequently always have 20 Euros to distribute. A pair of one high-
performing and one low-performing player would win funds with probability
0.5.

Second, to turn this wedge in entitlements into a set of social norms on

1005 dictator-giving, we aim at creating social norms, i.e., shared beliefs of what
others expect one ought to do (Bicchieri, 2006). To that effect, the subjects
in the main study are informed about the results of an earlier survey regard-
ing the normatively appropriate behavior in the experiment. We reported,
truthfully, that the mode of respondents in the earlier survey thought that
1010 a dictator should give 10 (out of 20) Euros to a high performer, and that
the mode of respondents thought a dictator should give 5 (out of 20) Euro
to a low performer. More precisely, in this survey, 18 of 26 subjects judged
10 to be the appropriate payment to a high entitlement receiver, while 12 of
26 thought that a payment of 5 was appropriate for a receiver with low enti-
1015 tlement (and 5 was the mode for that question). As these numbers suggest,
the variance in the survey regarding the latter question was significantly
higher. This is unsurprising because a dictator playing against a receiver
with high entitlement is in a clear situation of symmetry with the receiver:
they are both in the same performance class, and have equally contributed
1020 to the funds for distribution, strongly suggesting an equal distribution. By
contrast, it is less obvious how much a dictator (who always has a high
entitlement) should give a receiver with low entitlement. Reassuringly, 18
of 26 respondents in the preliminary study stated that a lower amount for
a low performer is appropriate, in line with our expectations.

1025 *Third*, all subjects are made aware that each receiver will be informed
about the matched dictator’s transfer decisions in order to make receiver
expectations salient.

Fourth, to assess and further strengthen the relevance of the social norms,
we asked all subjects before performing the quiz how they valued, on a scale
1030 from 1 to 4, the correctness and personal relevance of the announced nor-
mative expectations elicited in the prior norm induction session. Precisely,
subjects were asked to rate for both modes of normative expectations to
what extent they consider the 26 students’ opinion in the previous survey
a) to be ‘right’ and b) to be ‘important’ for them personally.

1035 Based on the model’s assumption of ϕ as exogenous parameters set by

social norms, the fair points are estimated as $\phi_H = 10$ and $\phi_L = 5$.

This fourfold treatment makes the mentioned fair points salient, communicates a normative expectation as to what one should do, and, since all subjects are informed, makes it *ceteris paribus* more likely that other dictators will also comply, creating or reinforcing a social norm of equity. Compliance can plausibly be conceived in a subjective or an objective sense, as the wording did not refer to the epistemic state of the dictator.

Subsequently, all subjects played a competitive quiz consisting of 15 questions in the style of the well-known TV-show ‘Who Wants to be a Millionaire’ under time pressure. Correct answers were rewarded with positive points, incorrect or missing answers with point deductions. Points increased with difficulty. The performers in the top three quartiles were assigned the name ‘Gold Quiz-Players’ (high entitlement), the last quartile ‘Silver Quiz-Players’ (low entitlement). In each session, all 16 dictators were drawn from the 24 Gold Quiz-Players, while the receivers were constituted by the remaining 8 Gold and 8 Silver Quiz-Players (see assignment as Screen 1 in Appendix I).

F.2. Dictators’ transfer strategy

The dictators’ giving strategy was elicited in order to have an empirical estimation of the individual responsiveness to the entitlement norm. At that stage dictators know that they are a high performer themselves, but they do not know which type of receiver they will face. Dictators bindingly state two amounts: the amount they give to the receiver if they learn the receiver is a high performer (x_H), and the amount they give if they learn the receiver is a low performer (x_L). Dictators do not directly choose the amount to be transferred if they do not learn which type the receiver is. Instead, this amount (x_U) is calculated as $(x_H + x_L)/2$, of which dictators are explicitly made aware of (see Screen 2 in Appendix I). We refrained from letting the dictators choose x_U in order to create a coarse-grained normative system and test proposition 2. Thus, we inform the subjects

in the written instructions that giving $(x_H + x_L)/2$ under uncertainty is appropriate because this is equi-distant from x_H and x_L . This has the added benefit that, in absolute terms, the monetary consequences of acquiring either signal are the same. When entering the strategy, the dictators do not
1070 know under which circumstances the information about their receiver might become available. The written instructions informed the subjects that more information on this would become available onscreen. This seems to have successfully inhibited questions from subjects on this delicate element as no subject asked the experimenter about the conditions under which the
1075 dictator would know which type of receiver she would be matched with.

F.3. Information acquisition and payoffs

The optional opportunity to acquire information about the receiver is the crucial design element to test for objective and subjective norm compliance. After each dictator has been matched with a receiver, they have the options
1080 (i) to receive no signal, (ii) signal S_L , (iii) signal S_H or (iv) both signals. The prior probability that the state is L is $p = \frac{1}{2}$ and the probability of removing uncertainty when choosing the ‘correct’ signal is $s = \frac{1}{2}$. The four options are implemented by presenting dictators with four envelopes on their computer screens (see Screens 3–5 in Appendix I). We call two of
1085 these envelopes ‘gold envelopes’, the other two ‘silver envelopes’. Exactly one of these four envelopes contains (represented electronically) the lot of the receiver (which in turn reveals the performance of the receiver). If the receiver has a winner-lot, it will always be in one of the two gold-envelopes, while a blank lot will always be in one of the two silver-envelopes.
1090 The dictator can now choose to open one gold-envelope; open one silver-envelope; open one gold- and one silver-envelope; or open no envelope. We subsequently (electronically) ‘open’ the chosen envelopes and show the results. If the dictators find the receiver lot, they have certainty about the receiver’s type. If they do not find the lot, uncertainty about the type
1095 remains. Opening a ‘silver’ envelope means receiving S_L and opening a

‘gold’ envelope receiving S_H .

As the dictator’s contribution strategy depends on what the dictator learns about the type of the receiver, the information received determines which amount will be transferred if the bonus is won in the final stage. Dictators are informed about the monetary consequences of the information acquisition. In fact, dictators have to run through a series of test question in order to minimize errors on their part. Note that receiving only S_L (only S_H) never leads to certainty about the receiver being a high performer (low performer); x_H (x_L) will therefore never be transferred. If the bonus is won, the dictator receives y_H if he has found a ‘winner lot’ (and therefore has learned that the receiver is a high performer), y_L if he has found a ‘blank lot’ (and has therefore learned that the receiver is a low performer) and y_U if he has not found the receiver’s lot (and therefore has not learned whether the receiver is a low or a high performer).²⁹

1110 **G. Instructions in the norm elicitation session**

The decisive part of the description of our experiment for the participants in the norm-elicitation session reads as follows (translated from German):

1115 ‘All participants take part in a knowledge quiz. The questions come from the TV game show ‘Who Wants to Be a Millionaire’. As in the real game show, the questions become more and more difficult. On the basis of their performance in the quiz, participants are divided into two groups: a ‘gold group’ and a ‘silver group’. Silver Quiz-Players are the 25 % with the lowest number of points in the quiz. The remaining 75% are Gold Quiz-Players.

²⁹We deliberately left open whether the receivers would be informed about the signal-choice of the dictators (in fact, receivers are not informed), as we did not want the dictators to wonder about information acquisition norms or create an experimenter demand effect towards acquiring or not acquiring signals.

All players receive a fixed amount of 5 Euros for participation in the quiz.

1125 At the end of the experiment, all participants take part in pairs in a lottery. In the lottery each pair has a chance to win a bonus of 20 Euros. A pair consists of two participants. One participant (the ‘allocator’) in a pair decides how to allocate the bonus in case that they win the bonus.

1130 **Composition of pairs**

The allocator is always a Gold Quiz-Player. The recipient is either a Gold Quiz-Player or a Silver Quiz-Player. One half of the receivers are Gold Quiz-Players and another half are Silver Quiz-Players. The chance of a pair to win the bonus depends on the performance of both participants in the quiz.

1135 For every Gold Quiz-Player one lot is drawn. There are always two lots in the pot and one of them is a winning lot. This means that if the allocator is matched with another Gold Quiz-Player, the pair wins the bonus with certainty. If only the allocator is the a Gold Quiz Player, the pair wins the bonus with 50% chance. A Gold Quiz-Player brings his/her 50% winning chance, whereas a Silver Quiz-Player cannot contribute to the chance of winning the bonus.

1140 The allocator receives the compensation for participation of 5 Euros plus, if the pair wins the bonus, the amount the allocator assigned to himself/herself. The receiver receives the compensation for participation of 5 Euros plus, if the pair wins the bonus, the amount the allocator assigned to him/her.’

1150 After these instructions are read, all participants answer on-screen test questions to make sure that the setup is understood. We then ask what the allocator should give to ‘Gold Quiz-Players’ and ‘Silver Quiz-Players’ (the precise wording is: ‘what do you think is the right amount to give to

the...'). In our experiment, we report the respective modal answers 5 and 10 to make that norm salient.

H. Written instructions for all participants in the (main) experiment

Experiment Instructions (translated)

Thank you for agreeing to participate in this decision-making experiment. Please read this description of the experiment carefully. For the entire duration of the experiment any communication with other experiment participants is prohibited. Please turn off your cell phones now. It is a mandatory requirement for participation in this experiment to comply with these rules.

If some points remain unclear, please read the experiment instructions again. For any remaining questions, please raise your hand. We will come to your desk and answer your questions in person. All important details of the experiment are also shown on screen. In addition, test questions help to ensure that all participants understand the experiment correctly.

You can earn money in this experiment. You will always receive a compensation of 5 Euros for participation in today's experiment. For the completion of questionnaires at the end of the experiment you will receive an additional 2 Euros. How to earn any further money will be explained in these instructions. Your data and decisions are anonymous. Your answers and decisions cannot be linked to your identity and no person-identifying data will be saved.

Part 1: Quiz

All participants take part in a knowledge quiz.

After the completion of the quiz, participants are divided into two groups:

- "Gold Quiz-Players": the 75% of all participants with the highest relative quiz performance
- "Silver Quiz-Players": the 25% with the lowest relative quiz performance
- You will not be informed of your individual performance in the quiz.
- Gold Quiz-Players get a **winner lot** for their performance. Silver Quiz-Players get a **blank lot**. These lots are used in a 20-Euros-draw, which is explained below.
-

Part 2: 20-Euros-draw

- Two players each form a participant-pair and take part in the 20-Euros-draw. A pair always consists of an allocator and a recipient.
- Which receiver and which allocator form a participant-pair is determined at random.
- Both the allocator and the recipient bring their lot from the quiz to the draw.
- All allocators are recruited from the Gold Quiz-Players and therefore always bring a winner lot to the draw.
- Half of the recipients are recruited from the Gold Quiz-Players and half from the Silver Quiz-Players. So half of the recipients bring a winner lot and half bring a blank lot to the draw.
- Neither the recipient nor the allocator knows if the recipient brings a winner lot or a blank lot to the draw. However, depending on the further course of the experiment, this information might be revealed to you. More information on this will be available on the screen.
- One of the two lots provided by the participant-pair will be drawn.
- If a winner lot is drawn: the allocator has to distribute the 20 Euros between himself and the recipient.
- If a blank lot is drawn: the 20 Euros are not won.

Task for the Allocator:

- The allocator decides on the allocation of the 20 Euros before he gets matched with a recipient. At this point it is still unknown which type of lot the recipient will bring into the lottery and if the 20 Euros will be won.

- The allocator decides on the allocation **in advance and bindingly**. **The distribution depends on the information the participant-pair will receive about the lot of the recipient:**
 - The allocator determines in advance the amount that will be transferred to the recipient in case the pair learns that the recipient is a **Gold Quiz-Player** and brings a **winner lot** into the lottery. This amount is called **GOLD-transfer**.
 - The allocator determines in advance the amount that will be transferred to the recipient in case the pair learns that the recipient is a **Silver Quiz-Player** and brings a **blank lot** into the lottery. This amount is called **SILVER-transfer**.
- These amounts determine:
 - **GOLD/SILVER-transfer:** the amount that will be transferred to the recipient if the pair **does not learn** which lot the recipient brings to the draw.
 - **GOLD/SILVER-transfer = average of the GOLD-transfer and the SILVER-transfer.** This amount is of equal distance to the GOLD-transfer and the SILVER-transfer. This amount is chosen because it is either possible that the recipients brings a winner lot, or that he brings a blank lot.
- In case of winning the 20 Euros:
 - The recipient gets, depending on the three cases, either the GOLD-transfer, or the SILVER-transfer, or the GOLD/SILVER-transfer.
 - The allocator gets 20 Euros minus the corresponding transfer. So either 20 minus the GOLD-transfer, or minus the SILVER-transfer, or minus the GOLD/SILVER-transfer.

Information for the recipient at the end of the experiment:

- The amounts of the GOLD-transfer, SILVER-transfer, and GOLD/SILVER-transfer, which were determined by the allocator;
- If the 20 Euros have been won;
- Which kind of information the allocator got regarding the recipient's lot;
- In case of winning the 20 Euros, whether the GOLD-transfer, the SILVER-transfer, or the GOLD/SILVER-transfer was transferred to him.

Chronological sequence:

1. Quiz
2. Players with the highest relative performance become Gold Quiz-Player (with a winner lot) and players with the relatively lowest quiz performance become Silver Quiz-Player (with a blank lot).
3. Division into allocators and recipients. Allocators are Gold Quiz-Player. Recipients are, distributed randomly, half Gold Quiz-Players and half Silver Quiz-Players.
4. Allocators decide on GOLD-transfer and SILVER-transfer. From these amounts the GOLD-SILVER-transfer results (The average of GOLD-transfer and SILVER-transfer).
5. Each allocator is matched with a recipient.
6. The participant-pair might or might not get the information about the type of the recipient's lot.
7. The draw decides whether the pair wins 20 Euros or not. Exactly one of the two lots of the players is picked at random.
8. In case of winning: Distribution corresponding to allocator's transfer decisions.
9. Allocators and recipients get information about the results of the experiment and their own payoff.

This sequence is represented graphically below:

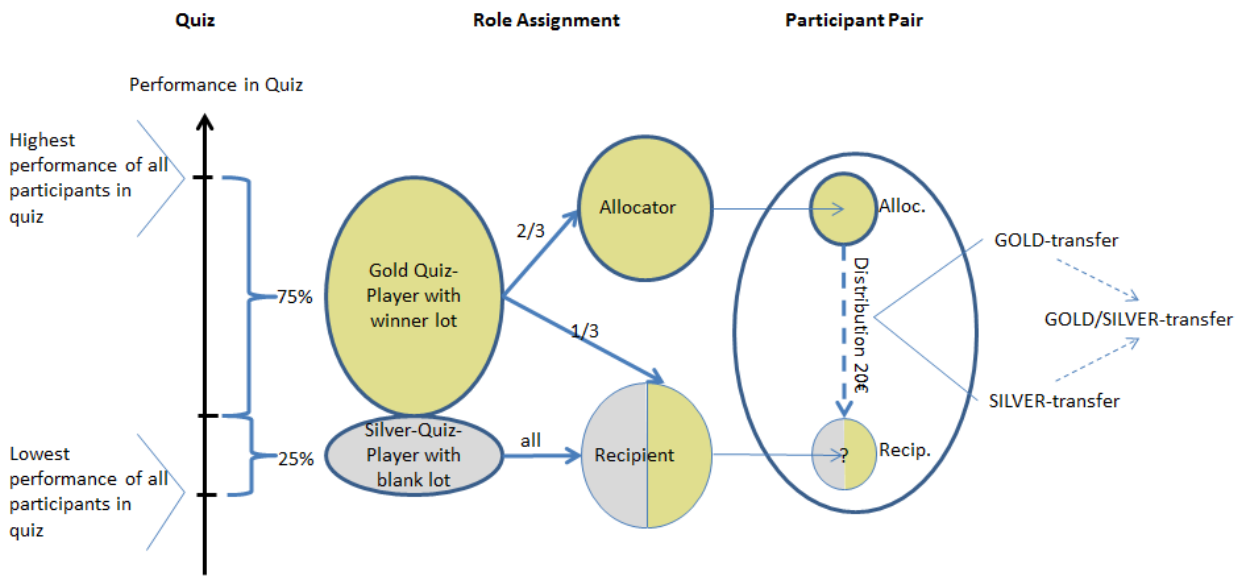


Figure 1

Taking Part as a Participant-Pair in the Draw

Case 1: Recipient brings blank lot to the draw

Case 2: Recipient brings winner lot to the draw

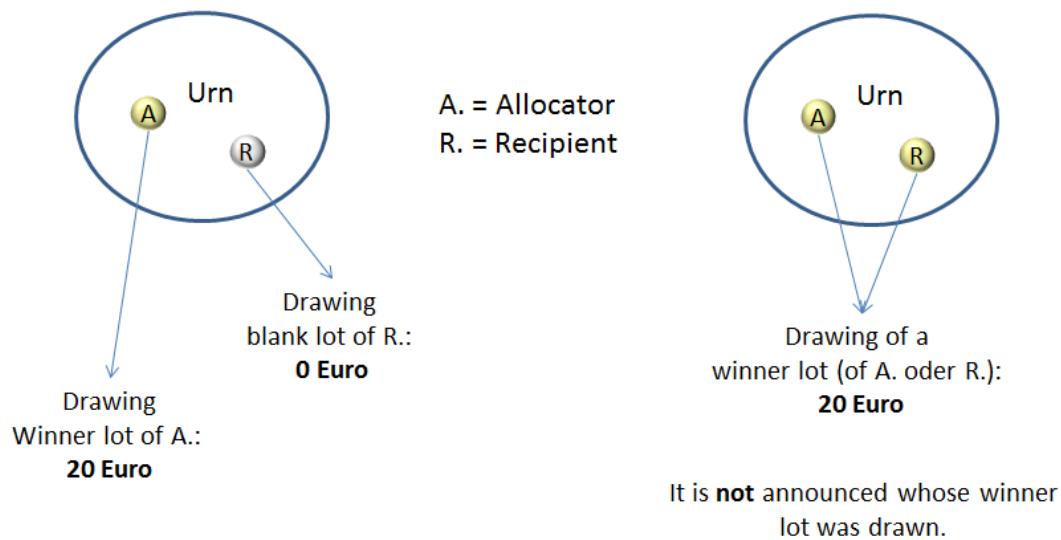


Figure 2

I. Important treatment screens (translated)

Screen 1

1 - Wine desktop

Periode 1 von 1

Verbleibende Zeit [sec]: 42

Allocation decision:

You are a Gold Quiz-Player and you received a winner lot for your performance.

2/3 of all Gold Quiz-Players were randomly chosen to be an allocator.

You are an allocator.

Soon you will be matched with a recipient. He is either a Gold Quiz-Player (with a winner lot) or a Silver Quiz-Player (with a blank lot). The probability for both cases is the same.

Continue

Screen 2

1 - Wine desktop

Please state now which part of the 20 Euros you transfer if you will win the money.

If you do not learn which type of lot the recipient brings into the lottery, the average of GOLD-transfer and SILVER-transfer will be transferred.

This amount is called GOLD/SILVER-transfer.

Possible cases:		You transfer	So you keep
1. Your participant-pair gets the information that the recipient is a Gold Quiz-Player and brings a winner lot into the lottery.	GOLD-transfer	<input type="text" value="10"/>	10
2. Your participant-pair gets the information that the recipient is a Silver Quiz-Player and brings a blank lot into the lottery.	SILVER-transfer	<input type="text" value="5"/>	15

Confirm Input

This results in the third case:

		Transfer	So you keep
3. You don't get the information which type of lot the recipient brings in and if he is a Gold Quiz-Player or a Silver Quiz-Player.	GOLD/SILVER-transfer	7.5	12.5

Continue

Screen 3

1 - Wine desktop

Periode 1 von 1

Verbleibende Zeit [sec]: 52

Explanation: Information Period

- **Now you have the chance to possibly see the lot the recipient brings into the lottery, before it is put into the urn for the 20 Euro draw. You DON'T influence the probability of winning the 20 Euros with this decision.**
- **Your choice influences what you will learn about the the recipient's lot and whether the GOLD-transfer, SILVER-transfer or GOLD/SILVER-transfer will be transferred, if you win the 20 Euros.**
- You now see 4 envelopes. **Exactly one of these envelopes contains your recipient's lot.**
- All other envelopes are empty.
- **If your recipient is a Gold Quiz-Player and has a winner lot, it is in one of the gold envelopes. If your recipient is a Silver Quiz-Player and has a blank lot, it is in one of the silver envelopes.**

Continue

Envelope 1 Envelope 2 Envelope 3 Envelope 4

Screen 4

1 - Wine desktop

Periode 1 von 1

Verbleibende Zeit [sec]: 49

Explanation: Information Period

- **Only if you find the recipient's lot in a gold envelope you know for sure that he brings a winner lot into the lottery and hence the GOLD-transfer of [var] Euros will be transferred.**
- This also means: if you do not open a gold envelope you will never find out whether the recipient brings in a winner lot and hence the GOLD-transfer will never be transferred.
- **Only if you find the recipient's lot in a silver envelope you know for sure that he brings a blank lot into the lottery and hence the SILVER-transfer of [var] Euros will be transferred.**
- This also means: if you do not open a silver envelope you will never find out whether the recipient brings in a blank lot and hence the SILVER-transfer will never be transferred.
- **The GOLD/SILVER-transfer of [var] Euros will be tranferred if you do not know which lot the recipient brings in. This is the case if you do not find the recipient's lot in one of the envelopes.**

Continue

Envelope 1 Envelope 2 Envelope 3 Envelope 4

Screen 5

1 - Wine desktop

Periode: 1 von 1

Verbleibende Zeit [sec]: 28

Decision: Lot Revelation


Please make your decision now.
You can open from 0 to 2 envelopes, but not more than one of the same kind.

Which envelope(s) do you want to open?

- Envelope 1 (Gold-Envelope)
- Envelope 2 (Gold-Envelope)
- Envelope 3 (Silver-Envelope)
- Envelope 4 (Silver-Envelope)
- none


[Continue](#)

One of these 4 envelopes contains the recipient's lot




Gold-Umschlag

Envelope 1




Gold-Umschlag

Envelope 2



Silber-Umschlag

Envelope 3



Silber-Umschlag

Envelope 4