Fishing for Good News: Motivated Information Acquisition *

Carl Heese †   Si Chen ‡

The motivated reasoning literature argues that people skew their beliefs to feel moral when acting selfishly. We study the information acquisition of decision-makers with a motive to form positive moral self-views and one to act selfishly. The selfish motive makes individuals dynamically "fish for good news": They are more likely to continue acquiring information having so far observed information indicating that acting selfishly is harmful to others, and more likely to stop after information indicating it is harmless. Further analysis finds no evidence the selfish motive worsens others' outcomes and suggests this is due to individuals fishing for good news.

Experimentally and theoretically, we study the information acquisition of a decision-maker for whom information might reconcile two motives that govern her utility: an egoistic motive—a desire to maximize personal gains—and a moral motive. Growing empirical evidence shows that moral motives are often belief-based. People want to "feel moral", whether their decisions are moral or not (for reviews, see Kunda, 1990; Bénabou and Tirole, 2006; Gino *et al.*, 2016).

This motive to feel moral might compete with the individual's egoistic motive if she believes maximizing her personal gain is detrimental to others. That is, she cannot behave selfishly while feeling moral. However, individuals are sometimes uncertain about whether a self-benefiting choice harms others. Under uncertainty, new information brings the chance to reconcile the egoistic and the moral motives since it may suggest that an egoistic decision is also moral. This gives incentives to

*; We are grateful for helpful discussions with Thomas Dohmen, Lorenz Götte, Lena Janys, Botond Kőszegi, Sebastian Kube, Stephan Lauermann, George Loewenstein, Florian Zimmermann, as well as comments from audiences at Bonn (2018), the EDP Jamboree (2018), the workshops on Belief-Dependent Preferences 2019 (Copenhagen), Behavioral Game Theory 2019 (Norwich), Contracts and Incentives 2019 (Munich), Social Preferences (theem 2019, Konstanz), Belief-Based Utility 2021 (Amsterdam), and other venues such as the Econometric Society World Congress 2020, and the ECBE Conference 2021. Funding by the German Research Foundation (DFG) through CRC TR 224, Project A01 (Chen) and Project B03 (Heese) is gratefully acknowledged.

†The University of Hong Kong, Department of Economics, heese@hku.hk

‡The University of Hong Kong, Department of Economics, econchen@hku.hk

seek new information. Conversely, there are incentives to avoid further information when her belief is that the motives are not conflicting (see Dana *et al.*, 2007, and the “strategic moral ignorance” literature).

There are numerous examples where an egoistic motive might conflict with a motive to feel moral: doctors receive commissions for prescribing certain drugs, but prescribing the commissioned drug might harm the patient’s health. A human resource manager may have personal preferences for job candidates of a particular ethnicity or gender, but hiring decisions based on her personal taste might harm the company’s performance and the job candidates’ careers. Consumers might find it economical to purchase fast fashion products, but doing so might support unethical production.

In these situations, systematic biases in information collection may affect the social outcome of decisions. How doctors gather information about the patient’s medical needs may affect the suitability of the prescribed drugs and, hence, the patient’s well-being. When human resource managers’ personal tastes against minority job candidates sway how they inform themselves about the candidates’ job-related qualities, the hiring outcomes are biased. What information consumers acquire about the production conditions of goods may affect the prevalence of unethical production.

This paper studies how the potential tradeoff between an egoistic and a moral motive shapes individuals’ information acquisition strategies and welfare consequences. To this end, we develop a simple theoretical model. The model will generate core predictions to guide an experimental analysis in the laboratory and help explain the observed behavior.

It meets three demands: First, to model the tradeoff, we consider an agent with two motives. When choosing between two options, she gains utility from her material payoff and the *belief* about her choice’ being harmless to others.[1] Second, the setup should be flexible enough to speak to the many real-life settings in which individuals face dynamic information acquisition decisions where single observations yield partial, inconclusive information about the future consequences of their choice. Third, the predicted behavior should only be driven by the tradeoff and not constrained or confounded by other factors such as cost considerations. For these reasons, there is no exogenous restriction on the agent’s choice of infor-

[1] The study of belief-based utility has a long tradition in economics (e.g., Loewenstein, 1987; Geanakoplos *et al.*, 1989; Akerlof and Kranton, 2000; Köszegi, 2006).

mation and no information cost. Information about which option harms others arrives in increments over time, and at every point in time, the agent can decide to either stop or continue observing the incoming information.

A controlled laboratory environment allows us to address three challenges facing an empirical investigation of information acquisition. First, exogenous variation in the motives is required to pin down the causal effect of having two potentially conflicting motives on information acquisition. Second, individuals' heterogeneous prior beliefs, access to information, and interpretation of the information can confound the observed information acquisition strategy. Third, individuals' histories must be monitored to analyze the dynamic patterns.

Our experiment addresses these challenges. First, we induce the moral motive by having the subjects make a binary dictator decision. In the dictator decision, one of the two options reduces the payoff of a receiver, while the other does not. The dictator does not know which option is harmful to the receiver. We exogenously vary the existence of an additional egoistic motive by randomly assigning the dictators into two treatments. In one treatment, one option increases the dictator's own payment, and she knows which option is self-benefiting. Thus, the dictator has an egoistic motive to choose the self-benefiting option. In the other treatment, the dictator's payment is not at stake in the dictator's decision. This treatment serves as a baseline. Second, we fix the dictator's prior belief about the likelihood of each option being the harmful one. Before making the decision, the dictator can acquire additional information about it. Information comes in pieces, is free, and the dictator can stop or continue receiving information at any time. The information has a clear Bayesian interpretation, and we provide the dictators with the Bayesian posterior beliefs after each piece of information. Third, we record all choices.

The main empirical result is that individuals exploit their ability to dynamically sample information. They "*fish for good news*" (Finding 1 and 2): They are more likely to *continue* acquiring information after having received mostly bad news so far and more likely to *stop* after having received mostly good news so far. Here, "bad news" is a piece of information indicating that behaving selfishly harms the other, and "good news" indicate the opposite.

A simple intuition comes from the model: After bad news indicating that a materially self-benefiting option is likely to harm others, an individual may be inclined not to choose this option to avoid a low belief-based utility. Then, more information comes in handy. First, if good news arrived, it may revert her decision to the self-benefiting option. Second, even if bad news arrived and she decided to forgo the self-benefit, she would be more certain that doing so spares the other from harm. Either way, she is better off acquiring further information. In contrast, when the individual has received mostly good news indicating that the self-benefiting option is likely harmless, she may be inclined to capture the self-benefits. Collecting further information would bear the risk that this becomes morally unacceptable. Theorem 1 formalizes this intuition.

Our finding that individuals stop comparatively early after good news aligns with the concept of "strategic moral ignorance" (Dana *et al.*, 2007), which notes that individuals sometimes avoid additional information to enable selfish behavior. A novel observation is that an egoistic motive can *cause* individuals to acquire *more information* than without, namely after previous unfavorable information. The general observation is that individuals engage in an asymmetric search for information, dubbed "fishing for good news;" strategic ignorance shows up as part of this strategy. We discuss the related literature in more detail momentarily.

Predicting the observation of information seeking after bad news requires a model with an imperfectly altruistic agent (one who would not acquire complete information without other incentives than the moral one). In the context of our model, this means that the belief utilities exhibit satisficing (Simon, 1955): The marginal belief utility is zero when the likelihood of harming others is below a reservation cutoff. We show that this way, the agent would stop at the reservation cutoff without other incentives.[2]

Further empirical analysis uncovers an intriguing aspect of altruistic decision-making. More egoistic incentives should bias choices to be relatively less considerate of others, which, naively, should harm them. We find that this is not necessarily true when choices have *uncertain* effects, i.e., when their harm is uncertain. Egoistic motives do *not* lead to more harm in our data (Finding 3).[3]

We explore the channels driving this observation. First, our model predicts that the effect on the receiver's welfare can be non-negative. Theorem 2 provides a characterization when it is positive and when it is negative. Second, what drives the result in our theory is that the egoistic motive causes the agent to change

[2]While there are other ways to model satisficing, incorporating it into the belief utility seemed most natural in the context of our model of Bayesian persuasion; see Section 4. There are various interpretations of the reservation cutoff, discussed in Section 1.4.

[3]The null hypothesis of different levels of harm is rejected even at the 10% level.

her information acquisition; she fishes for good news. Without this change in information, the sole effect of the egoistic motive would be to change the agent's preferences over the two options. She chooses the option now carrying a self-benefit even when it is more likely to harm the other. Theorem 2 effectively characterizes when the effect driven by information even outweighs the negative effect related to preferences. Roughly, this typically happens when the disutility from harming others is high enough and when the agent is a strong enough satisficer. Both properties are necessary: If the agent does not care about the other at all, she would simply choose the self-benefiting option. Without satisficing, she would choose complete information and cause zero harm when she only has other-regarding incentives.

Third, we construct a *counterfactual* treatment with our data to disentangle the two effects. This analysis suggests that the empirical effect related to information is positive. In other words, the caused change in information strategy—i.e. fishing for good news—reduces the harm to others. Moreover, it offsets the negative preference effect (Findings 4 and 5).

The paper primarily contributes to two empirical literature streams. First, the literature on motivated belief formation, particularly in the moral context. Most closely related is the prior work information choices. It has observed that people choose to remain "strategically ignorant". They avoid information about their actions' negative externalities on others and then behave more selfishly (see Dana *et al.*, 2007, and the subsequent literature)[4].

The main contribution is empirically documenting that selfishly motivated individuals use a strategy dubbed "fishing for good news" where they exploit their ability to sample information, in an asymmetric way: They tend to avoid additional information after good news and seek additional information after bad news. As discussed above, strategic information avoidance shows up as one part of the general strategy but has a counterpart of strategic information seeking. Our theory shows that "fishing for good news" is a fundamental implication of belief-based social preferences (Theorem 1). This strengthens the view of the broader literature on motivated beliefs that has argued that social preferences have a belief-based component (see, e.g., Gino *et al.*, 2016).

Regarding the experimental design, we complement the literature on strategic

[4]For a review of the work following Dana *et al.* (2007), see the meta-analysis by Vu *et al.* (2023) and the surveys by Gino *et al.* (2016) and Golman *et al.* (2017)).

ignorance by considering flexible, dynamic information acquisition. This allows to speak directly to the applications where information acquisition is piece-wise and dynamic. Another novelty is that we include a control treatment where selfish incentives are removed. This allows us to study the *causal* effect of selfish incentives—such as commissions offered to doctors or personal bias in hiring decisions—on information choices, and in particular, to include an analysis of the causal welfare effect.[5]

Our welfare results add to the literature on prosocial behavior. A rich body of work in economics has identified powerful drivers of altruistic and prosocial behavior, such as equity concerns or reciprocity; see, e.g., Fehr and Schmidt (1999) and Falk and Fischbacher (2006) or Camerer (2011) for a review. We highlight a novel and potentially counterintuitive aspect that arises under incomplete information. Namely, when a decision-maker has additional selfish incentives, this sometimes does not worsen and may, in theory, even improve outcomes for others affected. The reason is that these incentives make people change how they acquire information about the externalities of their choice on others before making it.

Section 4 discusses further literature and contributions. First, we compare the size of effects in our experiment with those in the prior work on "strategic ignorance", such as Vu *et al.* (2023) and Feiler (2014), and how our control helps identify in which instances and to which extent ignorance is "strategic" (in the sense of caused by selfish incentives). Second, we discuss self-image concerns, the most prominent explanation of strategic moral ignorance, in the context of our model. Third, we relate to the broader experimental literature on motivated reasoning and belief formation, including Ditto and Lopez (1992) and Eil and Rao (2011). Fourth, we relate to the experimental literature on dynamic information acquisition, in particular, Caplin and Dean (2013) and the neuroscience work on the drift-diffusion model. Finally, we explain how our model has an interpretation as one of Bayesian persuasion (Kamenica and Gentzkow, 2011) where the sender and receiver are two selves of the same agent ("self-persuasion"); we then provide a test of our self-persuasion model by estimating the empirical distribution of the persuasion strategies across treatments.

[5] In contrast, the classic experimental paradigm of Dana *et al.* (2007) compares a treatment where subjects have the binary choice to either acquire complete information or no information with a control treatment where all subjects have complete information.

1 Theory

We propose a formal model to analyze an agent's information acquisition in a decision where she has an egoistic motive and a motive to believe that her decision is moral and does not harm others. To highlight how the egoistic motive alters strategies and outcomes, we also study the scenario in which the egoistic motive is removed.

Section 1.1 presents the model and Section 1.2 some preliminary equilibrium analysis. Section 1.3 states the main theorem about information acquisition, showing that agents "fish for good news". In Section 1.4, we make three points that are important for understanding the interplay between the egoistic and the moral motive. These three points together build the intuition for the theorem. Section 1.5 states a second theorem about the welfare of others, characterizing when the egoistic motive affects the welfare positively. Section 1.6 derives testable predictions to guide our empirical investigation.

1.1 A model of conflicting motives

An agent (she) has to decide between two options, x and y. There is an unknown binary state $\omega \in \{X, Y\} = \Omega$, and the prior belief is that the probability of X is $p_0 \in (0, 1)$. A passive agent, whom we hereafter refer to as "the other" or "the receiver" (he), can be affected by the agent's decision between x and y. When the agent chooses an option that does not match the state, i.e., x in Y or y in X, the option has a negative externality of -1 on the other and otherwise not.

Two motives govern the agent's preferences. First, if choosing x, the agent receives a state-independent remuneration $r \geq 0$, while she receives no remuneration if choosing y. When $r > 0$, the remuneration constitutes an egoistic motive to choose x.[6] The case $r = 0$ serves as the benchmark without egoistic motive. Second, the agent has a moral motive. She dislikes the belief that her decision harms the other. We model this as the agent receiving a utility $u(a, q)$ when she believes her choice a is harmless for the other agent with probability q. When the agent believes that state X holds with probability p, she believes that x is

[6]The remuneration here is a token that stands not only for monetary interests but also any private interest that the agent might have. In the example of a discriminatory human resource manager, the private interest can be the utility of her choosing a candidate of her personally preferred gender.

harmless with probability $q = p$ and that y is harmless with probability $q' = 1-p$. Then, when she chooses $a \in \{x, y\}$, her utility is given by

$$U(a, p; r) \quad = \quad \begin{cases} u(a, p) + r & \text{if } a = x, \\ u(a, 1-p) & \text{if } a = y. \end{cases} \tag{1}$$

The belief-based utility u is weakly increasing in the second argument. We let $u(x, 1) = u(y, 1) = 0$. That is, the dictator feels no disutility if she is certain that her choice does not harm the other.[7] We also call the function u the (preference) type of the agent and consider the set of all types u satisfying the above assumptions unless otherwise stated. For concreteness, we state a parametric example of the belief-based utility u:

$$u(a, q) = -\alpha(1-q)^2.$$

Here, the belief utility is parametrized by $\alpha = u(a, 0)$, the disutility from choosing an action a that harms the other with certainty, i.e., when $q = 0$.

Before deciding between x and y, the agent can acquire information about the state at no cost; in particular, the discount rate is zero. Let $\mu(\omega) = -1$ if $\omega = X$ and $\mu(\omega) = 1$ if $\omega = Y$. Time is continuous, and at every instant in time, the agent can observe an information process $(Z_t)_{t\geq 0}$ given by $dZ_t = \mu(\omega)dt + dW_t$ where $(W_t)_{t\geq 0}$ is a standard Brownian motion. The posterior probability that the agent assigns to the state X at the time t is

$$p_t = \Pr(\omega = X | (Z_s)_{s\leq t}).$$

At every point in time, the agent can decide to stop or continue observing the process $(Z_t)_{t\geq 0}$, depending on the information she has already received. When the agent stops at $t \geq 0$, subsequently, the agent chooses an action a that maximizes her payoffs, i.e., $a \in \max_{a\in\{x,y\}} U(a, p_t; r)$ and the game ends. Formally, a strategy of the agent is a real-valued stopping time τ adapted to the natural filtration generated by the information process.

For technical reasons, we impose the "coarseness condition" that the agent stops and takes a decision when $p_t \leq \epsilon$ or $p_t \geq 1-\epsilon$, for some positive but *arbitrarily* small $\epsilon \approx 0$.[8] This rules out strategies where the agent observes the

[7] All results hold if $u(x, 1) = u(y, 1) = c$ for any $c \in \mathbb{R}$. Later, we show that Nash equilibria correspond to strategies that maximize the agent's utility. Hence, if $c \neq 0$, shifting utilities by the constant c does not alter the set of equilibria. So, $c = 0$ is a normalization. We normalize to simplify the exposition and because this way $u(a, p)$ has a natural interpretation as the disutility from harming the other person with probability $1-p$.

[8] In the experiment, posteriors are rounded to two decimal places, so that e.g beliefs below

information process infinitely with positive probability.

1.2 Equilibrium characterization

Lemma 1 *There are cutoffs $p_l \leq p_0 \leq p_h$ so that the following constitutes a Nash equilibrium: The agent continues to observe the information process as long as $p_l < p_t < p_h$ and stops whenever $p_t \leq p_l$ or $p_t \geq p_h$.*

The proof is in Appendix B. Lemma 1 shows the existence of an equilibrium. To prove the lemma, we leverage an insight from the analysis of Bayesian persuasion (Kamenica and Gentzkow, 2011).

Since there is no cost of observing the information process, any Nash equilibrium must maximize $\mathrm{E}(V(p_\tau))$ with

$$V(p) = \max_{a \in \{x,y\}} U(a, p; r)$$

and where p_τ is the stopped belief. This means that Nash equilibria are the agent-optimal stopping policies. It implies that all Nash equilibria are payoff-equivalent for the agent.

Lemma 2 *There is a unique Nash equilibrium in which the agent stops when indifferent between stopping and continuing.*

We prove Lemma 2 in Appendix B and show that the equilibrium in Lemma 2 is given by the belief cutoffs p_l and p_h as follows: let $\bar{V}$ be the smallest concave function with $\bar{V}(p) \geq V(p)$ for all $p \in [\epsilon, 1-\epsilon]$. If $\bar{V}(p_0) = V(p_0)$, then $p_h = p_l$. Otherwise, $I = (p_l, p_h)$ is the largest open interval in $[\epsilon, 1-\epsilon]$ with $\bar{V}(p) > V(p)$ for all $p \in I$.

Equilibrium selection. For ease of exposition, we focus on the equilibrium in Lemma 2 and use p_l and p_h to refer to it. One can show that the main result does not depend on this equilibrium selection.[9]

1% are identified with certainty, essentially implementing $\epsilon = 0.01$.

[9] In the Appendix D.2, we also show that the equilibrium in Lemma 2 is the unique Nash equilibrium that is stable concerning the introduction of minimal cost, by considering a variation of the model with cost (attention cost, time cost, search cost, etc.).

1.3 Result: fishing for good news

The critical difference between the scenarios with and without an egoistic motive is that when there is an egoistic motive ($r > 0$), the agent makes a tradeoff between the desire for the remuneration and a desire for accurate beliefs. How does this tradeoff affect the agent's information acquisition?

Our main result, Theorem 1, concerns the effect on the "intensive margin" of information acquisition, i.e., the agent's decision to continue or stop acquiring information once she has started. In Appendix D.1.1, we discuss the "extensive margin" of information acquisition, i.e., the agent's decision whether to acquire *any* information. In Theorem 1, we consider all types that plan on acquiring some information and use it in a "responsive" way, i.e., choosing y after information indicating that y is harmless to the other, and x after information indicating that x is harmless to the other.[10] [11] Formally, a type is "responsive" given some $r > 0$ if it is strictly optimal in equilibrium to choose y at $p_l(r')$ and x at $p_h(r')$ when $r' = r$ and also when $r' = 0$.

The theorem shows that when $r > 0$, the agent stops and chooses y only at a more extreme belief in y being harmless, $1-p_l(r) \geq 1-p_l(0)$. Conversely, the agent stops and chooses x at a less extreme belief in x being harmless, $p_l(r) \leq p_l(0)$.

Theorem 1 (Fishing for good news) *Let $r > 0$. For any responsive preference type u,*

$$p_h(r) \leq p_h(0), \quad \textit{and } 1 - p_l(r) \geq 1 - p_l(0).$$

Theorem 1 reveals an asymmetry. In intuitive terms, the left inequality shows that to convince herself to choose the remunerative option x, the agent needs less information supporting the innocuousness of x ("good news"). The right inequality shows that for choosing the non-remunerative option y the agent needs more information opposing the innocuousness of x ("bad news"). Taken together, the agent "fishes for good news" to choose the remunerative option.

We elaborate on the intuition for such behavior below in Section 1.4. Figure 1 illustrates. It shows an example where both inequalities are strict; this requires

[10]The theoretical analysis in this section focuses on the tradeoff between belief-based utility and the remuneration. It turns out that when $r > 0$, some types choose to acquire some information but choose x regardless of the information they receive. However, such behavior is not driven by a meaningful tradeoff between belief-based utility and remuneration. For example, concavity of u would rule out such behavior.

[11]In our data, most dictators behave responsively (88%).

additional conditions on the belief utility, stated at the end of Section 1.4.

Figure 1: Stopping and continue intervals with and without egoistic motive

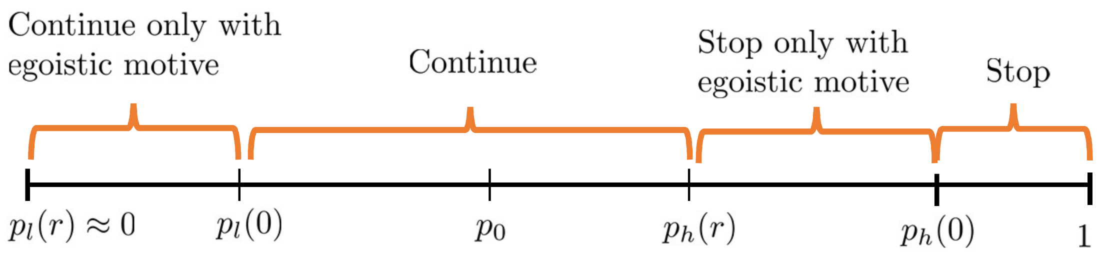

1.4 Intuition for fishing for good news

We highlight three points that are important to understand Theorem 1. Finally, we put things together to sketch the logic of fishing for good news. We first discuss each point in intuitive terms and then state a formal result.

The desired belief. The first point is that when one option is remunerative, the agent prefers higher beliefs in the state where this option is harmless. This is because, when believing that the remunerative option is harmless to the other, she can capture the reward without having a bad conscience. In contrast, when she believes that the remunerative option harms the other, she has to make a tradeoff between a clear conscience and the remuneration—there is a moral dilemma.

Formally, let $r > 0$. Recall that in equilibrium, the agent eventually either stops at p_h or p_l and that at p_h, she has a higher belief about the likelihood that x is harmless. Similar to the intuition sketched in the previous paragraph, the following result shows that, in equilibrium, the agent is better off when she stops at the higher belief p_h than when she stops at p_l. The proof is in the Appendix.

Lemma 3 *For all $r > 0$, the agent chooses x when stopping at the belief $p_t = p_h$. Further, if the agent weakly prefers to choose y when holding the belief $p_t = p_l$, then, $V(p_l) < V(p_h)$.*

In the following two parts, we analyze at which beliefs the agent stops. First, we analyze at which belief p_l the agent stops and chooses y. Second, we analyze at which belief p_h the agent stops and chooses x.

Waiting for good news. Above, we made the point that the agent prefers to believe that the likelihood of the remunerative option being harmless is high. The second point is: When she believes this likelihood to be low so that she is inclined to choose the non-remunerative option, the agent prefers to continue observing the arriving information. One intuitive reason for this behavior is that she hopes to receive "good news" so that her belief increases, making it optimal to choose the remunerative option. The second reason is that even if no good news arrived, her belief in the innocuousness of the non-remunerative option would increase, and so would her belief-based utility when choosing it. In any case, she is better off continuing. Formally, we show the following result.

Lemma 4 *For all $r > 0$: If the agent weakly prefers to choose y when holding the belief $p_t = p_l$, then $p_l = \epsilon \approx 0$.*

Proof. At each point of time $t \geq 0$, the equilibrium strategy τ^*, given by p_l and p_h, maximizes the continuation payoff $\mathrm{E}(V(p_\tau)|(Z_s)_{s\leq t})$,

$$\begin{aligned}\mathrm{E}(V(p_{\tau^*})|(Z_s)_{s\leq t}) &= \frac{p_h - p_t}{p_h - p_l}V(p_l) + \frac{p_t - p_l}{p_h - p_l}V(p_h) \\ &= u(y, 1 - p_l) + \frac{p_t - p_l}{p_h - p_l}\Big[V(p_h) - V(p_l)\Big], \end{aligned} \tag{2}$$

where, for the first equality, we used that $\mathrm{E}(p_{\tau^*}|(Z_s)_{s\leq t}) = p_t$ by Bayes-consistency.[12] For the second equality, we used that the agent chooses y at p_l so that $V(p_l) = u(y, 1 - p_l)$. We see that the continuation payoff strictly decreases in p_l since the likelihood of reaching p_h, that is $\frac{p_t - p_l}{p_h - p_l}$, decreases in p_l and since the utility $u(y, 1 - p_l)$ when reaching the lower belief p_l, also decreases in p_l. We conclude that, unless the agent is certain that y is harmless, she would like to continue observing the arriving information; thus, $p_l = \epsilon$.[13] ■

Good enough news. The third point is that when she believes that the remunerative option is likely harmless, she decides whether to stop and choose this option by making a tradeoff between her belief-based utility and the remuneration. On the one hand, if she continues, her belief in this option being harmless may increase further, allowing her to have a better conscience when choosing it. On the other hand, continuing to acquire information bears the risk of observing

[12] Given the strategy τ^*, Bayes-consistency implies $\Pr(p_{\tau^*} = p_h|(Z_s)_{s\leq t}) = \frac{p_t - p_l}{p_h - p_l}$ and $\Pr(p_{\tau^*} = p_l|(Z_s)_{s\leq t}) = \frac{p_h - p_t}{p_h - p_l}$.

[13] Recall the technical restriction that the agent has to stop if $p_t = \epsilon$ where $\epsilon \approx 0$ is arbitrarily small.

information that makes the remunerative option unacceptable, i.e., that leads her to update to a low posterior and choose the non-remunerative option.

Formally, at each point of time $t \geq 0$, the equilibrium strategy τ^*, given by p_l and p_h, maximizes the continuation payoff $\mathrm{E}(V(p_\tau)|(Z_s)_{s\leq t})$. From Lemma 4, we take $p_l \approx 0$, so that $\mathrm{E}(V(p_\tau)|(Z_s)_{s\leq t}) \approx \Pr(p_\tau = p_h|(Z_s)_{s\leq t})V(p_h)$. For expositional purposes only, let $u(x,p)$ be continuously differentiable at $q > p_0$. Using $V(p_h) = u(x,p_h) + r$, the first-order condition with respect to p_h is

$$0 = \Pr(p_\tau = p_h|(Z_s)_{s\leq t})\frac{\partial u(x,p_h)}{\partial p_h} + \frac{\partial \Pr(p_\tau = p_h|(Z_s)_{s\leq t})}{\partial p_h}(u(x,p_h) + r), \quad (3)$$

which shows that the agent makes a tradeoff between the marginal increase in belief-based utility from stopping at a higher belief p_h and the marginal decrease in the likelihood of stopping at p_h, which comes with the remuneration r. Rewriting (3),[14]

$$0 = \frac{\partial u(x,p_h)}{\partial p_h}p_h - (u(x,p_h) + r). \quad (4)$$

Recalling $u(x,1) = 0$, one sees from (4) that when the marginal increase in belief-based utility is relatively small for high beliefs $p_t \approx 1$, precisely when $\frac{\partial u(x,1)}{\partial p_h} < r$, the agent is willing to stop and choose x before she is certain that x is harmless. One may say that the agent stops when she has received "good enough news."

The logic of fishing for good news. To derive how the egoistic motive alters the information acquisition incentives, let us shortly turn to the benchmark scenario in which both options are *not* remunerative.

In the benchmark, the agent's utility depends solely on her belief about the likelihood that her action does not harm the other. For the agent types u with $\frac{\partial u(a,q)}{\partial q} > 0$ for $a = x, y$, the more certain they are that their decision does not harm the other, the higher their utility would be. For these agent types, acquiring as much information as possible is optimal. Other agent types have a threshold level of certainty. They are content when believing they are sufficiently likely to spare the other from harm. Any further certainty beyond the threshold does not increase their belief-based utility. At the threshold, such types are indifferent between continuing and stopping, so they may as well stop. This behavior mirrors that of *satisficing* as in Simon (1955).[15]

[14]Recall that $\Pr(p_{\tau^*} = p_h|(Z_s)_{s\leq t}) = \frac{p_t - p_l}{p_h - p_l} \approx \frac{p_t}{p_h}$, so that $\frac{\partial \Pr(p_\tau = p_h|(Z_s)_{s\leq t})}{\partial p_h} \approx -\frac{p_t}{p_h^2}$. Plugging this into (3) gives $\frac{p_t}{p_h}\frac{\partial u(x,p_h)}{\partial p_h} - \frac{p_t}{p_h^2}(u(x,p_h) + r) = 0$, which simplifies to (4).

[15]Caplin *et al.* (2011) find that the satisficing model of Simon (1955) describes dynamic

Formally, the threshold level of certainty is

$$l(a) = \min\{q : u(a,q) = 0\}.$$

Recall here that $u(a,q) \leq u(a,1) = 0$ for all (a,q).

To sum up, in the benchmark without egoistic motive, the agent stops acquiring information (only) when further certainty no longer increases her utility, i.e. $p_h(0) = l(x)$ and $1 - p_l(0) = l(y)$. In contrast, with an egoistic motive, the agent trades off the belief-based utility with the remuneration and is willing to forego belief utility from a more accurate belief and choose x at lower beliefs $p_h(r) \leq l$ already; formally, this is implied by (4). Further, with an egoistic motive, the agent only stops and chooses y at extreme beliefs, i.e. $1 - p_l(r) = 1 - \epsilon \approx 1$. Together, this shows the inequalities of Theorem 1.

Interpretations of the Threshold. The threshold level of certainty l and the related satisficing behavior in the benchmark may be interpreted literally (marginal utilities are zero) but also rationalized in other ways. For example, the agent may exhibit emotions of guilt. Typically, guilt is formulated as a relative notion in games (see, e.g., Battigalli and Dufwenberg, 2007). That is, the agent's guilt increases in the harm she inflicts on the receiver relative to some expectation of the receiver. This receiver expectation may be captured by a likelihood l of not being harmed. More generally, the agent may attempt to meet some moral reference point or constraint (Rabin, 1995) captured by l.

Strict Effects. Strictness of the inequalities in Theorem 1 requires additional conditions on the belief utility. The right inequality is strict whenever the agent exhibits *some* satisficing, i.e., when $l(y) < 1$: Then, $1 - p_l(0) = l(y) < 1$, and $1 - p_l(r) \approx 1$ for $r > 0$, given Lemma 4. The left inequality is strict whenever the agent's marginal belief utility is smooth at $q = l$: If $\frac{\partial u(x,q)}{\partial q}$ is continuous at $q = l$, then $\frac{\partial u(x,l)}{\partial q} = 0$ so that $p_h(r) < l$ given (4) and thus $p_h(r) < p_h(0)$ since $p_h(0) = l$.

1.5 Results on Receiver's Welfare

How does the presence of the egoistic motive affect the receiver? More egoistic incentives should bias choices to be relatively less considerate of others, which, naively, should harm them.

---

information acquisition behavior well across a range of settings.

Indeed, if the agent does not care about the other sufficiently, she simply chooses the remunerative option. If the remuneration r is larger than the disutility $\alpha = |u(x, 0)|$ from harming the other with certainty, she will choose to pocket it given *any* belief about the state. This harms the receiver; he is worse off since the agent acts responsively to her acquired information and in the receiver's interest absent the egoistic motive.

Theorem 2 shows that, maybe surprisingly, the opposite can be the case if the agents' prosocial attitude is more pronounced (the converse scenario $r < \alpha$): The agent's egoistic motive sometimes benefits the receiver. As in Theorem 1, we only consider types with "responsive" behavior, i.e., those that acquire some information and act responsively to it both with and without egoistic motive $r > 0$. Responsiveness implies

$$r < \alpha; \tag{5}$$

beyond this, it rules out that the type's prior belief about action x being harmless would be so high he chooses x immediately, and some "odd" behavior where the agent acquires information but chooses x regardless of the information she receives.[16] If the agent would not act responsively with an egoistic motive (but only without), the welfare effect must be negative, by the same reasoning as when $r > \alpha$.

For the theorem, we suppose that the label of the option that harms the other does not matter to the agent: $u(x, q) = u(y, q)$. In particular, $l := l(a)$ does not depend on the option $a \in \{x, y\}$.

Theorem 2 *Let $r > 0$, $\epsilon \approx 0$. For any responsive preference type u with $u(x, q) = u(y, q)$ for all $q \in [0, 1]$, the egoistic motive's effect on the receiver's welfare is strictly positive if*

$$1 - p_h(r) < \frac{1 - l}{1 - l + p_0},$$

with $p_h(r)$ defined after Lemma 2. It is strictly negative if the reverse inequality holds strictly.

The result highlights the role of satisficing: For example, when $u(x, q)$ is weakly convex for $q \leq l$, the condition holds if and only if $l < 1$.[17] Similarly, when $u(x, q)$

[16] As noted in footnote 10, such is not driven by a meaningful tradeoff between the belief-based utility and the remuneration and thus cannot be rationalized meaningfully.

[17] Given $l < 1$ and convex u, the agent acquires more information in a Blackwell sense when

is strictly concave for $q \leq l$, for a wide range of specifications (the curvature of the belief utility, the reward r, the prior p_0), there is a cutoff $0 < l^* < 1$ so that the condition cannot hold when $l \geq l^*$; see Appendix B.7.

To conclude, for a wide range of specifications, the welfare effect can only be positive if the agent's disutility from harming the other is sufficiently high ($r < \alpha$) and if she is a strong enough satisficer, as measured by a sufficiently low l.

Satisficing implies that she acquires only partial information without an egoistic motive. The egoistic motive alters her information acquisition; she "fishes for good news", and this *indirectly* affects final choices and the resulting harm to the other. Without this change in information acquisition, the sole effect of the egoistic motive would be the negative one sketched above for the scenario $r > \alpha$: The egoistic motive might make her choose the option now carrying a remuneration even at the low stopped belief p_l where it more likely harms the other than the other option does. This would harm the receiver.

Formally, fix the equilibrium information strategy $\tau(0)$ without the egoistic motive and suppose the agent would not alter it when facing an egoistic motive. The egoistic motive then only alters the preferences over the option. This "preference effect" on the receiver's welfare is

$$\mathrm{PE} = \mathrm{E}\Big[-1_{a(p_{\tau(0)},r)\neq\omega}\Big] - \mathrm{E}\Big[-1_{a(p_{\tau(0)},0)\neq\omega}\Big] \leq 0, \tag{6}$$

where $p_\tau(0)$ is the realized belief at which the agent stops given $\tau(0)$, and $a(p_\tau(0), r')$ is the agent's optimal action at the realized belief with and without egoistic motive, $r' = r > 0$ and $r' = 0$.[18] The effect of altering the information strategy ("information effect") is the residual difference in the receiver's welfare between both scenarios, with and without egoistic motive.

Theorem 2 characterizes when the information effect implies an overall positive welfare effect and entirely offsets the negative preference effect. Appendix B provides the proof.

To conclude the discussion of the theorem, we show why $1 - p_h(r) < \frac{1-l}{1-l+p_0}$ is the relevant condition. Recall the equilibrium analysis from Section 1.4: The likelihood of harming the other is $1-l$ in the scenario without egoistic motive. The

$r > 0$. This is because the stopped beliefs $p_l(r)$ and $p_h(r)$ are more extreme than $p_l(0)$ and $p_r(0)$: namely, $p_l(r) \approx 0 < 1 - l = p_l(0)$ and $p_h(r) = l = p_h(0)$.

[18]The expression $a(p, r) \neq \omega$ slightly abuses notation; it describes the cases when $a(p, r) = y$ and $\omega = X$, or when $a(p, r) = x$ and $\omega = Y$.

likelihood of harming the other with egoistic motive is approximately $\Pr(x|Y)(1-p_0)$—since Lemma 4 applies when the type is responsive; so, the agent almost never chooses y in X, and the other is essentially only harmed when x is chosen in Y. Therefore, the overall welfare effect is positive if

$$\Pr(x|Y)(1-p_0) < 1-l \Leftrightarrow 1-p_h(r) < \frac{1-l}{1-l+p_0}; \tag{7}$$

for the equivalence we used that $\Pr(x|Y)(1-p_0) = 1-l$ is equivalent to $1-p_h(r) = \Pr(Y|x) = \frac{\Pr(x|Y)(1-p_0)}{\Pr(x|Y)(1-p_0)+\Pr(x|X)p_0} \approx \frac{1-l}{1-l+p_0}$. The overall welfare effect is negative if the reverse inequality holds.

1.6 Testable predictions

We derive testable predictions from Theorem 1 and 2 for our experimental setup. These predictions will guide our data analysis, which is presented in Section 3.

In the experiment, we invite a pool of subjects to individually play the information acquisition game of the theory part. To make the setting natural for the experimental subjects, we let them acquire information piece-wise via binary signals that are i.i.d. conditional on the state. We collect the data on each subject's stopping and continuing decision after each draw of an information piece. The collected observations are on the individual-draw level.

People fish for good news. We derive two predictions for the collected data from the theorem about "fishing for good news" (Theorem 1). The first concerns behavior after fixed sequences of information pieces, e.g., one piece of good news or two pieces of bad news, etc.

Prediction 1. Having observed a fixed information sequence that leads to a posterior $p_t < p_0$ ($p_t > p_0$), individuals are weakly more likely to *continue* (*stop*) acquiring information when they have an egoistic motive, compared to when they have not.

The prediction rules out strict differences in the wrong direction after any fixed information sequence. It stems from the observation that, for all preference types, the continue interval below the prior is weakly larger, and the continue interval above the prior is weakly smaller with an egoistic motive; see Figure 1 for an illustration. Thus, the likelihood that p_t falls into the interval is larger and smaller, respectively.

The second prediction relies on all observations on the individual-draw level.

Prediction 2. When individuals have an egoistic motive, conditional on having observed information so that $p_t < p_0$ ($p_t > p_0$), they are strictly more likely to *continue* (*stop*) acquiring further information, compared to when they do not have an egoistic motive.

The prediction requires that a share of the individuals in the population exhibit the mild properties implying strict differences in the stopping cutoffs with and without egoistic motive (we discussed in Section 1.4 which properties imply strictness of Theorem 1's inequalities). By considering the stopping decisions and pooling together all individual-draw observations, we will pick up these strict differences in the stopping cutoffs predicted by the theory.

The Receiver's welfare. Theorem 2 predicts heterogeneous effects of the egoistic motive on the welfare of others: the effect may be positive or negative, depending on the prosocial attitude. This means, averaging across individuals, the egoistic motive need not reduce the other's welfare.

Prediction 3. The egoistic motive need not lead to a reduction of the other's welfare on the population level.

In the theory part, we provided some initial discussion on how such a possible non-negative effect could only be driven by the agent changing her way of acquiring information, i.e., by her fishing for good news. We continue discussing this channel in Section 3 of the empirical part.

2 A laboratory experiment

We conduct a laboratory experiment with modified binary dictator games. All participants have the same initial endowment. Contingent on an unknown state, one of the two options that the dictator has to choose from reduces the receiver's payoff, while the other option does not reduce the receiver's payoff. Before deciding, the dictator can acquire costless information about which option has a negative externality on the receiver.

2.1 The treatment variations

Our experiment has a 2×2 design. The key treatment variation in our experiment is whether one option in the dictator game generates more payoff for the dictator than the other. In the *Tradeoff* treatment, one option increases the dictator's payoffs while the other does not. In the *Control* treatment, neither option affects the dictator's payoffs. The comparison between *Tradeoff* and *Control* pins down the causal effect of having a self-benefiting option on the dictator's information acquisition behavior and the welfare consequences for the receiver. We describe the details of this treatment variation below when we present the dictator game.

The secondary treatment variation helps us to address potential (dynamic) self-selection into the information acquisition process, as will be discussed in Section C.4 of the appendix. It distinguishes the *Force* and *NoForce* treatments. In *NoForce*, the dictators can proceed to the dictator game without observing any additional information about the externalities of their options, while in *Force*, all dictators receive at least one piece of information. The *Force* and *NoForce* treatments are identical except for whether the dictators are forced to receive the first piece of information or not. We detail the design of the information acquisition procedure when we outline the main stages of the experiment in Section 2.4.

Consequently, there are four treatments in our experiment: *NoForce-Tradeoff*, *NoForce-Control*, *Force-Tradeoff*, and *Force-Control*. In the main text, we report all findings pooling the respective *Force* and *NoForce* treatments. For the sake of exposition, we will there refer to *NoForce-Tradeoff* and *Force-Tradeoff* as *Tradeoff*; *NoForce-Control* and *Force-Control* as *Control*. In the appendix, we provide additional findings from comparing *Force-Tradeoff* and *Force-Control*.

2.2 The dictator game

At the beginning of the experiment, all subjects receive 100 experimental points as an endowment. Each experimental point is equivalent to 5 Euro cents. With this endowment, the subjects play the dictator game. Table 1 presents the payment scheme of the dictator game in *Tradeoff* and *Control*, respectively. In both treatments, the dictator chooses between two options, x and y. There are two states of the world, *x harmless* or *y harmless*. Depending on the state, either option x or option y reduces a receiver's payment by 80 points, while the respective other option does not affect the receiver's payment. Note that each option harms the

Table 1: Dictator decision payment schemes

(a) *Control* Treatments

	Good state (x harmless)	Bad state (y harmless)
x	$(\mathbf{0}, 0)$	$(\mathbf{0}, -80)$
y	$(0, -80)$	$(0, 0)$

(b) *Tradeoff* Treatments

	Good state (x harmless)	Bad state (y harmless)
x	$(+\mathbf{25}, 0)$	$(+\mathbf{25}, -80)$
y	$(0, -80)$	$(0, 0)$

These tables present the dictator games in the *Control* and *Tradeoff* treatments. The number pairs in the table present (dictator's payment, receiver's payment), denoted in experimental points. Each point is worth 5 cents.

receiver in one of the states. This design ensures that the dictator cannot avoid the risk of harming the receiver without information about the state. In *Control*, the dictator receives no additional points regardless of her choice and the state. In *Tradeoff*, x is self-benefiting for the dictator: She receives 25 additional points when choosing x but no additional points when choosing y.

***Good state* vs *Bad state*.** For ease of exposition, we hereafter refer to state *x harmless* as the *Good state*, and state *y harmless* as the *Bad state*. In state *x harmless*, option x increases in *Tradeoff* the dictator's payments without disadvantaging the receiver. Believing she is in state x, the dictator can choose the self-benefiting option x without feeling immoral. In contrast, in state *y harmless*, option x increases the dictator's payment at the cost of reducing the receiver's payment: There is a moral dilemma. Although this labeling is not meaningful in the *Control* treatments, we will stick to it throughout for consistency.

The dictator starts the experiment without knowing the state that she is in. She only knows that in every twenty dictators, seven are in the *Good state*, and thirteen are in the *Bad state*. That is, the dictator starts the experiment with a prior belief of 35% on that she is in the *Good state* and of 65% on that she is in the *Bad state*. A high prior belief in the *Bad State* strengthens the moral dilemma: choosing the self-benefiting option x without further information most likely harms the receiver. The prior belief is the same in *Control* and *Tradeoff*. Hence, the comparison between *Tradeoff* and *Control* is not driven by the asymmetric prior.

Before making the decision, the dictator can draw additional information and obtain more accurate beliefs about the state that she is in. We describe the information in the following subsection.

2.3 The noisy information

We design a noisy information generator for each state. The information is easily interpretable. Specifically, each piece of information is a draw from a computerized box containing 100 balls. In the *Good state*, 60 of the balls are white, and 40 are black; in the *Bad state*, 40 balls are white, and 60 are black (see Figure 5 in Appendix C.1). The draws are with replacement from the box that matches each dictator's actual state. After each draw, we display the Bayesian posterior about the likelihood of each state on the dictator's computer to reduce the cognitive cost of interpreting the information and to prevent non-Bayesian updating.

Good news vs. bad news For ease of exposition, we refer to a white ball as a piece of *good news* and a black ball as a piece of *bad news*. This is because, in the *Good state*, dictators draw a white ball with a higher probability. Hence, the draw of a white ball leads to an increase in the dictator's belief about the likelihood of the *Good state*—the state where in the *Tradeoff* treatments the dictator can choose x and gain the additional payment without reducing the payment of the receiver. Reversely, in the *Bad state*, dictators would draw a black ball with a higher probability. A black ball is evidence for the *Bad state*. In *Control*, we will still refer to a white ball as good news and a black ball as bad news for consistency, although the dictators in *Control* should not have a preference over the two states, hence also not over the color of the balls.

2.4 The experimental procedure

The experiment consists of three parts: the preparation stage, the main stage, and the supplementary stage.

The preparation stage. (i) The dictator reads paper-based instructions about the dictator decision and the noisy information. (ii) In these instructions, we also describe Bayes' rule and tell the dictator that later in the experiment, we will help them interpret the information by showing them the Bayesian posterior beliefs after each ball they draw. (iii) The instructions also specify that each experiment participant starts the experiment with 100 points of an endowment. (iv) We also inform them that option x is harmless for 7 out of 20 dictators and y for 13 out of 20. That is, the dictator's prior information is that the Good state has a likelihood of 35% on the *Good state* and that the *Bad state* has a likelihood of 65%.

After reading the instructions, the dictators answer five control questions designed to check their understanding. They keep the paper-based instructions for reference throughout the experiment.

The main stage. In the main stage, (i) the dictators can acquire information about the state that they are individually in; (ii) they choose between x and y in the dictator game.

Specifically, the dictator can acquire a piece of information by clicking a button that makes the computer draw a ball randomly from the box matched to their actual individual state (see Figure 5 in C.1). After each draw, the screen displays the latest ball drawn and the Bayesian posterior beliefs about the *Good state* and the *Bad state* given all the balls drawn so far (rounded to the second decimal, see Figure 6 in Appendix C.1). The screen has two buttons: one to draw an additional ball and the other to stop drawing and proceed to the dictator game. In *NoForce*, the dictators can proceed to the dictator game without drawing any ball, while in the *Force*, the dictators must draw at least one ball. After the first draw, the dictators have complete autonomy regarding when to stop drawing. We do not enforce any maximum number of balls that they can draw.

Throughout the experiment, the dictators have no other way to learn about their true state besides drawing balls. It is common knowledge that the receiver does not learn the information acquired by the dictator.

The draws do not impose any monetary cost on the dictator. The time cost of acquiring information is limited: Between draws, there is a mere 0.3-second time lag to allow the ball and the Bayesian posterior to appear on the computer screen. A dictator can acquire 100 balls within 30 seconds, yielding almost certainty.

Having ended the information acquisition, the dictator chooses between x and y in the dictator game in Table 1a (in the *Control* treatments) or Table 1b (in the *Tradeoff* treatments). After that, the dictators' choices are implemented in the implementation stage, and the payments are calculated.

The supplementary stage. (i) We elicit the dictator's posterior beliefs about the state after the dictator game. The belief elicitation is incentivized.[19] We compare the elicited and the Bayesian posterior beliefs in Appendix C.8. We find that, for a majority of dictators, their elicited posterior beliefs and their Bayesian posterior beliefs coincide, and the deviation is not significantly different between *Tradeoff* and *Control* (two-sided Mann-Whitney-U test, p=0.29). (ii) The subjects take part in the Social Value Orientation (SVO) slider test, which measures "the magnitude of concern people have for others" and categorizes subjects into altruists, prosocials, individualists, and the competitive type (Murphy *et al.*, 2011). The test provides baseline evidence that prosocial concerns matter to our subjects. (iii) The subjects answer a questionnaire surveying their socio-demographics, e.g., the gender and age. They also answer a 5-item Raven's progressive matrices test (Raven *et al.*, 1998), which measures cognitive ability. We report the details of the supplementary stage in Appendix C.8.

Treatment assignment and implementation. We randomize within each laboratory session: (i) the *Tradeoff* and *Control* treatments, (ii) the states: We randomly assign 35% of the laboratory terminals to the *Good state*, and 65% to the *Bad state*. The subjects are then randomly seated and randomly matched in a ring for the dictator game. The subjects are told that their decisions would affect the payment of a random participant in the same experimental session other than themselves. After all the subjects have made their dictator decisions, the experiment moves on to the implementation stage, where we inform the subjects that the dictator game decisions are being implemented and their payments are affected according to another participant's dictator game decision. Each subject plays the dictator game only once.

We conducted the experiment in October and December 2018 at the BonnEconLab. 496 subjects took part (250 in *Tradeoff* and 246 in *Control*). Among the subjects, 60% are women, and 93% are students. They are, on average, 24 years old, with the youngest being 16 and the oldest being 69. The subjects are balanced between treatments concerning gender, student status, and age (see Appendix C.8). We used z-tree (Fischbacher, 2007) to implement the experiment and hroot (Bock *et al.*, 2014) to invite subjects and to record their participation. Instructions and interfaces on the client computers were written in German, as all subjects were native German speakers.

[19]We incentivize the belief elicitation using the randomized Quadratic Scoring Rule adapted from Drerup *et al.* (2017) and Schlag *et al.* (2013). For the stated belief that the likelihood of the good state is $b\%$, we calculate the following value

$$M = \begin{cases} \frac{(b-100)^2}{100} & \text{if} \quad x \text{ is harmless,} \\ \frac{b^2}{100} & \text{if} \quad y \text{ is harmless.} \end{cases} \tag{8}$$

Then, the computer draws a random number $A \sim U[0,100]$ and the dictator receives 30 points if $A > M$.

Payments. In the experiment, payments are denoted in points. One point equals 0.05 Euro. At the end of the experiment, the individual computer screens display the details of the points and the equivalent payments earned. The subjects received payments in cash before leaving the laboratory. The total earnings of a subject were the sum of the following components: an endowment of 5 Euro, an additional 1.25 Euro if the subject was in treatments *Tradeoff* and chose x, a 4 Euro reduction if the subject's randomly assigned dictator made a decision that reduces her payments, a random payment of either 1.5 or 0 Euro for revealing their posterior beliefs, a payment ranging from 1 to 2 Euro depending on the subject's decisions in the SVO slider test, a payment ranging from 0.3 to 2 Euro depending on the decisions in the SVO slider measure of another random subject in the same laboratory session, and a fixed payment of 3 Euro for answering the questionnaire. A laboratory session lasted, on average, 45 minutes, with an average payment of 11.14 Euro.

Discussion of experimental parameter choices. Setting a high reduction amount (4 EUR), we aimed to ensure the presence of an altruistic motive. We set the 1 EUR payment for choosing x to be comparably low so that the incentive to choose this self-benefiting option would not dominate the altruistic motive easily. The asymmetric prior of 35% plays a similar role. At a symmetric prior, both choices are symmetric in terms of expected harm on the other. This symmetry could easily nudge the subjects to choose the self-benefiting option x right at the start, making it hard to observe and analyze (dynamic) information acquisition.

3 Findings

The median number of information pieces acquired by the dictators is 6 in *Tradeoff* and 5 in *Control* (two-sided Mann-Whitney, p=.24). We provide further summarizing statistics in Appendix C.2 and proceed below with the analyses of the dictators' information acquisition behavior on the intensive margin, considering all data in the experiment.

3.1 Individuals fish for good news

The first main finding from the experiment is that the egoistic motive causes individuals to exploit their ability to sample information as follows. The dictators in *Tradeoff* "fish for good news": Compared to the dictators in *Control*, having received more bad news, the dictators in *Tradeoff* are more likely to *continue* acquiring information; having received more good news, the dictators in *Tradeoff* are more likely to *stop* acquiring information.[20]

In the following sections, we provide a series of findings that document this pattern of fishing for good news. In Section 3.1.1, we investigate Prediction 1. For this, we compare between *Tradeoff* and *Control* the dictators' decisions to continue acquiring information after the first draw and after the second draw. In Section 3.1.2, we investigate Prediction 2. For this, we analyze the entire information acquisition histories, leveraging statistical tools from survival analysis.

3.1.1 Behaviour after the first pieces of information

To provide evidence for "fishing for good news," we first consider the dictators' decision to continue or stop acquiring information after one piece of information and two pieces of information.

Finding 1 *The dictators' decision to continue acquiring information after the first piece of information differs from the decision in the Control baseline. Differences depend on the information that has been acquired. The same holds after the second piece of information. (i) Having received more bad news, weakly more dictators continue acquiring information in Tradeoff than in Control. (ii) Having received more good news, weakly fewer dictators in Tradeoff continue acquiring information than in Control.*

First, as predicted by "fishing for good news", the treatment effect on whether to continue acquiring information depends on the previous information. Specifically, in a logistic regression, the interaction effect of being in *Tradeoff* and having acquired more good news is significantly negative (after the first piece of information: $p = .00$ and after the second piece of information: $p = .049$; for details see Table 3 in Appendix A.1).[21]

[20] Please note that fishing for good news predicts individuals' behavior on the intensive margin of information acquisition, i.e., the information acquisition decisions after the individual has received at least one piece of information. Therefore, our analyses below will also focus on the intensive margin. On the extensive margin, 15% and 7% dictators proceed to the dictator game without acquiring any information in *NoForce-Tradeoff* and *NoForce-Control* respectively (Chi-Square $p = 0.02$).

[21] The first signal is exogenous in the Force treatments. When only looking at individuals in these treatments, the results after the first signal also hold, and the effect size is slightly larger.

Specifically, the effects summarized in Finding 1 confirm Prediction 1: Compared to the *Control* baseline, facing a self-benefiting option makes individuals more likely to *continue* acquiring further information when the previous information indicates that this option harms others. Conversely, when the previous information suggests that the self-benefiting option is harmless, individuals are more likely to *stop* acquiring information. Figure 2 presents the proportions of dictators who continue acquiring information right after the first piece of information and the first two pieces of information.

Figure 2: Proportion of dictators continuing after the first draws

(a) After the first piece (b) After the second piece

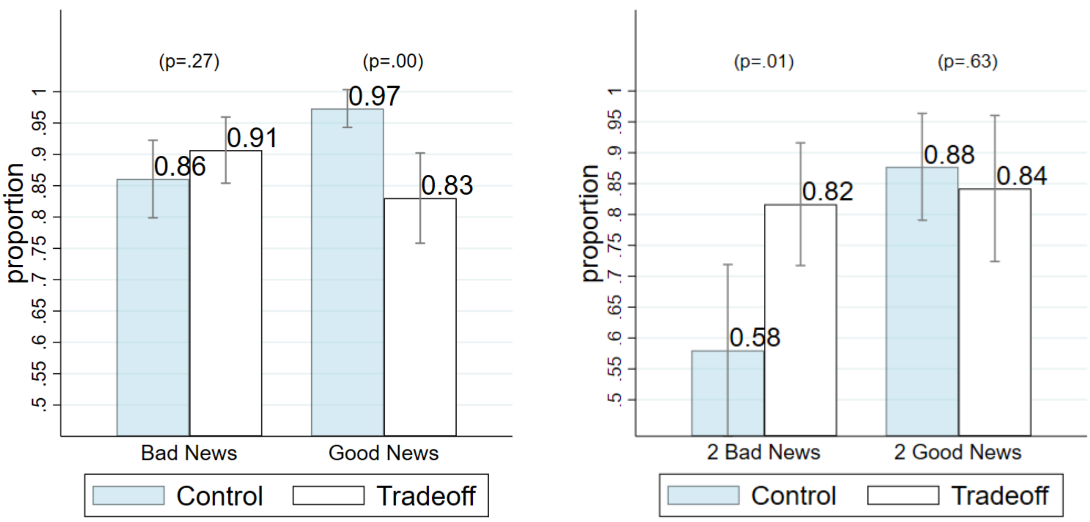

The figure presents the proportion of dictators who continue acquiring information after the first piece of information (2a) and the first two pieces of information (2b). In the parentheses, we present the p-values of the Chi-square test. In *Control*, the within treatment difference given different first news is due to the asymmetric prior belief of 35% in the *Good* state.[22] *Control* serves as the baseline and controls for the prior's effect on the information acquisition strategy.

Prediction 1 rules out strict differences in the wrong direction after any information sequence. All four comparisons made in Figure 1 are consistent with this, including the similar effect after one piece of bad news and the similar effect after two pieces of good news.

[22]In Control, behavior is likely driven by altruistic motives. Here, people might already be content at somewhat extreme beliefs to choose the action that quite certainly causes no harm ("satisficing"). The posteriors after one or two "bad news" are 0.26 and 0.19 and quite extreme. The posteriors after one or two "good news" are 0.45 and 0.55 and close to the uniform belief. Thus, satisficing would predict a higher proportion of individuals to acquire another signal after one good news than after one bad news, and a stronger such effect when comparing two good news with two bad news. Figures 2a and 2b report this.

The theory also predicts *strict* between-treatment comparisons when considering the stopping decisions after *each* draw, not just the first and second information piece (provided that a significant share of the dictators exhibits the mild properties implying strictness); see Prediction 2. We will investigate this prediction in the Section 3.1.2.

3.1.2 The entire information histories

We only provide simple comparative statics for the first two pieces of information since the sample size and the statistical power shrink as the information process unfolds and some dictators stop acquiring information.[23] We leverage tools from survival analysis to jointly estimate the effect of having received more good news or more bad news when considering all information histories.

Model specification. We carry out our analysis in the framework of the Cox proportional hazards model (Cox, 1972). The Cox proportional hazards model is often used for studying what influences individuals' hazards of choosing an exit option when they face exit decisions repeatedly (e.g., in the literature on unemployment, see Card *et al.*, 2007; Michelacci and Ruffo, 2015). We use the Cox model to investigate the dictators' hazards to stop acquiring information when they repeatedly decide whether to acquire further information. The Cox model has the advantage that it can address the dynamic selection that happens as observations drop out from the observed process. We will discuss this in detail in Section C.4. Another advantage of the Cox model is that the coefficient estimates have a direct interpretation as hazard ratios, which we will explain momentarily when interpreting our estimation results.[24]

Taking *Control* as the baseline, we analyze the dependence of the *Tradeoff* dictators' decision to continue or to stop acquiring information on whether, up to that point, they have received more good news or more bad news. The model specification is the following:

$$h(t|X) = h_0(t) \cdot exp(\beta_1 \text{Tradeoff} + \beta_2 \text{Info} + \beta_{12} \text{Tradeoff} \times \text{Info} + \alpha z_t). \quad (9)$$

$h(t|X)$ denotes the dictator's hazard rate to stop acquiring information after the t-th piece of information, given the set of covariates X; the baseline hazard

[23] In Section C.4, we discuss how our empirical framework addresses potential issues with self-selection, explaining that our estimates are lower bounds for the effects.

[24] We report a robustness check using a logistic model in Appendix C.3.

function $h_0(t)$ captures the hazards over the draws at covariate vector 0.[25]. The three covariates of interest are the treatment dummy "Tradeoff"; "Info", the categorical variable denoting information histories that have more pieces of bad news, good news, or an equal number of bad and good news, with bad news dominance as the baseline; and the interaction of the two.

Model assumptions. The model is correctly specified if (a) the covariates shift the baseline hazard proportionally so that the hazard rate $h(t|X)$ is multiplicative in the covariates ("proportional hazards assumption"), and (b) there are no omitted variables.

The proportional hazards assumption can be violated when some sample subgroups have different baseline hazards, $h_0(t)$. Stratification on the characteristics that might affect the hazard rate is often employed (see, e.g., Blossfeld *et al.*, 2019) to make sure this assumption is not violated. Stratification allows the baseline hazards $h_0(t)$ to vary on the strata while it estimates the aggregate effect of the covariates across all the stratified groups. We stratify our estimation on the following characteristics that can affect the baseline hazard: gender, cognitive ability (measured by the score in a Raven's matrices test), and prosocial types (categorized by the SVO measure of Murphy *et al.*, 2011). After the stratification, the proportional hazards assumption of the Cox model is not violated, whereas without stratification, it is.[26] [27] We also control for the accuracy of the individual belief after each ball drawn.[28]

It has been shown that omitting variables in the Cox model would only lead to underestimating the effects of interest (see Bretagnolle and Huber-Carol, 1988). To conclude, since the proportional hazards assumption is not violated in our estimation and since omitted variables can only lead to underestimation of the

[25]The Cox model naturally includes no constant term, since $h_0(t)$ already captures the hazard rate at covariate vector 0 (see for example Cleves *et al.*, 2010).

[26]See Table 4 in Appendix A.2.

[27]We report a robustness check using a logistic model in Appendix C.3. The logistic model's results are in line with those of the Cox model. The logistic model can be viewed as a hazards model with a proportional odds ratio assumption (Cox, 1975). However, unlike the Cox model, it does not allow for the baseline hazard to vary with the covariates. That is, it makes stronger assumptions than the stratified Cox model. Details are in the Appendix.

[28]In the experiment, the prior belief in the Good state is 0.35, a belief smaller than 0.5. Therefore, the posterior belief is more accurate after an information history with k more pieces of bad news than good news relative to one with k more pieces of good news than bad news. We control for information accuracy to prevent this difference from being picked up by the Info dummy. For this, we use the (expected) Brier score (Brier, 1950) of the beliefs as a proxy for the accuracy of beliefs: $\text{belief}_{\text{Good}} \times \text{belief}_{\text{Bad}}^2 + \text{belief}_{\text{Bad}} \times \text{belief}_{\text{Good}}^2$.

effects of interest, the significant results we find are lower bounds of the effects.

Data. To test the dependency of the treatment effect on the previous information, we first need to obtain a crucial independent variable: a factor variable denoting whether, after a draw, the information history has more good news or more bad news. Within an individual, this variable can vary after each draw. That is, the variable is time-varying. To obtain time-varying covariates, we follow the survival analysis literature and split each dictator's information history at the unit of individual draws (see Blossfeld *et al.*, 2019, pp 137-152). The resulting data set consists of "pseudo-observations" at the person-draw level. For every draw of each dictator, the pseudo-observation records the dictator's information history up to that draw and whether the dictator chooses to stop or continue acquiring information directly after that draw. For each pseudo-observation, we distinguish between information histories with more pieces of good news, more pieces of bad news, or the same number of good and bad news.

Finding 2 *Compared to the Control baseline, (i) having received more bad news than good news, the dictators in Tradeoff are more likely to* ***continue*** *acquiring information; (ii) while they are more likely to* ***stop****, having received more good news than bad news.*

We find that given information histories with more bad news than good news, being randomly assigned to *Tradeoff* has a significantly *negative* effect on the dictators' hazards to stop acquiring information ($\hat{\beta}_1 = -.29$, $p = .02$). The interaction term between the treatment and having acquired more good news is significantly *positive* ($\hat{\beta}_{12|\text{good}} = .43$, $p = .03$) instead. We explain these results below and report the details of the Cox model estimation in Table 4 of the appendix.

These findings show that the dictators fish for good news. First, given information histories with more bad news than good news, the between-treatment comparison of the stopping hazard can be expressed by the following hazards ratio:

$$\begin{aligned} \text{HR}_{bad} = \frac{h(t|\text{bad},\ \mathit{Tradeoff})}{h(t|\text{bad},\ \mathit{Control})} &= \frac{exp(\beta_1 \cdot 1 + \beta_2 \cdot 0 + \beta_{12} \cdot 1 \cdot 0 + \alpha z_t)}{exp(\beta_1 \cdot 0 + \beta_2 \cdot 0 + \beta_{12} \cdot 0 \cdot 0 + \alpha z_t)} \\ &= \frac{exp(\beta_1 + \alpha z_t)}{exp(\alpha z_t)} \\ &= exp(\beta_1); \qquad (10) \end{aligned}$$

Recall that $\hat{\beta}_1 = -.29 < 0$. This means that when dictators have acquired more bad news, the hazard to stop is *lower* in *Tradeoff* than in *Control*. Specifically, the ratio between the two is $exp(-.29) \approx .75$—i.e. in *Tradeoff* the hazard to stop acquiring information is 25% lower than in *Control*.

Second, given information histories with more good news than bad news, the between-treatment comparison of the stopping hazard can be expressed by the following hazards ratio:

$$\begin{aligned} \text{HR}_{\text{good}} = \frac{h(t|\text{good, } \textit{Tradeoff})}{h(t|\text{good, } \textit{Control})} &= \frac{exp(\beta_1 \cdot 1 + \beta_{2|\text{good}} \cdot 1 + \beta_{12|\text{good}} \cdot 1 \cdot 1 + \alpha z_t)}{exp(\beta_1 \cdot 0 + \beta_{2|\text{good}} \cdot 1 + \beta_{12|\text{good}} \cdot 0 \cdot 1 + \alpha z_t)} \\ &= \frac{exp(\beta_1 + \beta_{2|\text{good}} + \beta_{12|\text{good}} + \alpha z_t)}{exp(\beta_{2|\text{good}} + \alpha z_t)} \\ &= exp(\beta_1 + \beta_{12|\text{good}}). \end{aligned} \tag{11}$$

Since $\hat{\beta}_1 = -.29$ and $\hat{\beta}_{12} = .43$, $exp(\beta_1 + \beta_{12|\text{good}}) = exp(-.29 + .43) \approx 1.15$. So, the hazard to stop acquiring information is 15% *larger* in *Tradeoff* than in *Control* when the dictators have received more good news than bad news up to that point.

3.2 The effect of selfish incentives on the others' welfare

How does it affect the welfare of others if the decision-makers have a selfish incentive to choose one option over the other? Such incentives should bias choices to be relatively less considerate of others, which, naively, should harm them.

We investigate the causal effect by comparing the *Tradeoff* treatment to the baseline treatment where no option carries a self-benefit (*Control*).

Finding 3 *There is no significant effect on the receivers' welfare: Having a self-benefiting option does not make the dictators harm the receivers more often.*

The proportion of harmed receivers does not significantly differ between *Tradeoff* and *Control* (32% compared to 27%, Chi-Square $p = 0.17$).

Our theory predicts the possibility of such a non-negative impact on others; see Prediction 3. In the theory, this is driven by the agent changing her information acquisition relative to the baseline: She fishes for good news.

If she did not change her information choice, the sole effect of the selfish incentive on the other's welfare would be negative. It would only change her preferences over the two options. This might make her choose the incentivized option even when it implies more harm to the other in expectation than the other option. We formalized this "preference effect" in (6) as well as the residual effect related to information changes, the "information effect".

Empirical strategy. In light of this theory discussion, we construct a counterfactual with our data to learn about the empirical information effect. The counterfactual proxies the outcome that would occur had the dictators in *Tradeoff* not fished for good news but followed the information strategy they would have hypothetically chosen in *Control*.

We rely on the data observed in both treatments to proxy these hypothetical choices. First, we randomly draw a final posterior belief given their empirical distribution in *Control*. Then, given the empirical distribution of choices in *Tradeoff* at the drawn posterior, we draw a random choice. Table 2 illustrates the *Counterfactual* outcome constructed in this way.

Table 2: Constructing the *Counterfactual* scenario

Counterfactual	*Tradeoff*	*Control*
information as in		×
average decision given final posterior as in	×	
comparison with *Counterfactual* informs about	*information effect*	*preference effect*

When comparing the receivers' welfare in *Counterfactual* to *Control*, we learn about the "preference effect" by keeping the information fixed, that is, the distribution of final posterior beliefs. When comparing the receivers' welfare in *Counterfactual* to that in *Tradeoff*, we learn about the "information effect" by keeping fixed the average revealed preference between the options given each final posterior.

Finding 4 *The estimated preference effect is negative: Controlling for the dictators' stopped beliefs, having a self-benefiting option makes the dictators choose the harmful option more often.*

We compare the *Counterfactual* with the *Control* and find a negative effect. In the *Counterfactual*, the proportion of harmed receivers is higher than in the *Control* treatment (38% compared to 27%, Chi-Square $p = .00$).

Finding 5 *The estimated information effect is positive: Having a self-benefiting option changes how dictators acquire information in a way that they choose the harmful option less often, controlling for the dictators' decisions given their belief.*

We compare *Tradeoff* with the *Counterfactual* and find a positive effect. In *Tradeoff*, the proportion of harmed receivers is lower than in the *Counterfactual* (32% compared to 38%, Chi-Square $p = .046$). This suggests that had the dictators in *Tradeoff* acquired information the way the dictators in *Control* did, they would have inflicted *more harm* on the receivers. In other words, the changed information acquisition likely alleviates the effect of the preference bias towards x.

Aggregating the two effects, the proportion of harmed receivers does not significantly differ between *Tradeoff* and *Control* (Finding 3).

Information Acquisition in *Control*. In *Control*, where the dictators' own payments are unaffected by their decisions, they would (likely) not cause harm to the receivers if they acquired complete information.[29] However, the dictators in *Control* only acquire a limited amount of information, leaving room for fishing for good news to improve the receivers' welfare, in line with Finding 5. In *Control*, 27% of the dictators choose the option that harms the receiver. The dictators in *Control* only acquire a limited amount of information—the median number of pieces of information they draw is 5.

4 Literature

Motivated Beliefs. The paper contributes to the literature on motivated beliefs; see, e.g., Haisley and Weber (2010); Di Tella *et al.* (2015); Falk and Szech (2016); Zimmermann (2020); Gneezy *et al.* (2020); Möbius *et al.* (2022); Exley and Kessler (2024). Most related is the prior work involving information choices.

In the moral context, a rich body of work documents "strategic moral ignorance".[30] Meta-analytic results by Vu *et al.* (2023) reveal that within the experimental binary-choice paradigm of Dana *et al.* (2007) where subjects make a binary choice between no information and complete information, 40% of participants avoid information. In our *NoForce* treatments, subjects can also avoid information altogether: 15% of the subjects do so in the treatment group, while only 7% in the control. Our different design, where subjects can sample information piece-wise and slowly, thus seems to encourage some information acquisition.

[29] In our data, all the 7 *Control* dictators who acquire information until the Bayesian posterior beliefs displayed to them are rounded to certainty cause no harm to the dictator.

[30] See, e.g., Dana *et al.* (2007); Grossman (2014); Bartling *et al.* (2014); Van der Weele (2014); Serra-Garcia and Szech (2021) and the review by Gino *et al.* (2016).

Further, in our setting, the numbers show that only about half of the choices to avoid information completely are driven by selfish incentives, and thus "strategic" or "willful."[31]

In our *Force* treatments, in which participants are made to receive at least one piece of information, information choices differ with the initial information substantially: The fraction of participants avoiding additional information is 3% and 21% in the control and treatment groups respectively after one piece of good news (implied posterior $p = .46$), and 12% versus 7% after one piece of bad news (implied posterior $p = .26$). Prior work by Feiler (2014) has observed more avoidance at higher priors, yet with only minimal differences at the comparable initial information (18.25% at $p = .5$ vs 17.5% at $p = .2$). The comparison suggests that our design with noisy information makes choices more sensitive to the initial information.Our findings also suggest that the behavior we observe at low beliefs does not document "willful" ignorance, i.e., avoidance out of a selfish motive; instead, at low posteriors, the selfish motive causes strategic information seeking relative to the control.

The literature has provided other comparative statics consistent with our theory: For example, Spiekermann and Weiss (2016) have shown that individuals exhibit more substantial information avoidance with signals that generate good news less often (in our terminology). In our theory, the signals that maximize the likelihood of good news yield the highest expected utility to the agent.[32]

Another stream of this literature has analyzed the drivers of ignorance. The most prominent explanation is that self-image concerns drive it: Ignorance may provide an excuse that helps reduce the damage to one's self-image when acting selfishly (Grossman and van der Weele, 2017). Self-image concerns (about how one trades off selfish gain with others' welfare) are not concerned with behavior absent selfish motives. This implies that absent additional deviations from outcome-based social preferences, they would predict complete information acquisition in our control treatment. This way, self-image concerns do not explain well our finding of heightened information acquisition after bad news relative to the control and the non-negative welfare consequences (see Finding 1, 2 and 3). Notably, they do predict strategic information avoidance after good news. We prove this within a

[31] This result has been replicated by Exley and Kessler (2023) who pick up our idea of a control without selfish incentives within the standard experimental framework by Dana *et al.* (2007).

[32] This is a standard observation from the literature on Bayesian persuasion; see Kamenica and Gentzkow (2011).

variation of the main model in which the belief utility $u(a, p)$ arises endogenously in a self-signaling equilibrium (Bodner and Prelec, 2003). This variation can be found in a previous version of this paper, Chen and Heese (2020), and generalizes the prior binary-choice model of Grossman and van der Weele (2017).[33]

In a different context, Eil and Rao (2011) have studied how a self-related motive affects information demand.[34] They find that individuals' willingness to pay for information about their IQ (or beauty) is increasing in their current belief about their rank relative to others. In contrast, our fishing for good news findings imply the inverse result in the moral context. Demand for information is comparably high at currently low beliefs and comparably low at high beliefs.

Prior work in psychology by Ditto and Lopez (1992) in teamwork and medical contexts shares with our work that it also features dynamic information acquisition choices. However, the prior work has not analyzed the dynamics of information choices but measures of the aggregate information acquired (precisely, the time spent or the number of information pieces acquired). They document that individuals require less supportive information to reach a preferred conclusion than a non-preferred one. They suggest this is due to individuals' interpreting non-preferred information more skeptically (”motivated skepticism”; Kunda (1990)). In comparison, we facilitate Bayesian updating in the experiment and focus on how individuals strategically exploit the ability to sample information rather than the notion that information deemed more valid leads to a conclusion faster. Relatedly, to mimic real-life situations, the prior work attempts to create substantial time pressure; in contrast, we try to avoid an additional speed-accuracy tradeoff.

Dynamic Information Acquisition. In economics, there is little experimental work on dynamic information acquisition. This is the first paper in the context of moral decision-making. Prior work has studied collective information acquisition by committees (Chan *et al.*, 2018; Reshidi *et al.*, 2021), strategic experimentation (see, e.g., Hoelzemann and Klein, 2021, for a recent contribution), and dynamic search by individuals with a focus on analyzing the mental models employed (Gabaix *et al.*, 2006; Brown *et al.*, 2011; Caplin *et al.*, 2011).

Most related is Caplin *et al.* (2011) who emphasize the importance of the satisficing model of Simon (1955) and find that it describes dynamic information

[33] Our model allows for unrestricted information choices by leveraging techniques from Bayesian persuasion.

[34] Their primary analysis is about information updating.

acquisition behavior well across a range of settings. Similarly, our theory and empirical findings highlight the role of satisficing.

In neuroscience, a rich body of work inspects binary perceptual tasks using the "drift-diffusion model" to describe information processing in the brain (Swensson, 1972; Luce, 1991; Ratcliff and Smith, 2004; Ratcliff and McKoon, 2008). Our theory shows how egoistic incentives affect information acquisition strategies and subsequent choices within a standard drift-diffusion model.

Methodologically, we implement dynamic information acquisition with an experiment close to Wald's sequential sampling model (Wald, 1945). To our knowledge, such a sequential sampling design is quite novel. It makes the experiment particularly "clean" by giving us tight control over prior beliefs and the access and interpretation of information and allowing us to monitor information histories closely. Concurrent work by Reshidi *et al.* (2021) structures information choices similarly via sequential sampling but emphasizes a cost tradeoff, which we deliberately attempt to eliminate to focus on the tradeoff between the two motives.

Bayesian Self-persuasion. The paper relates to the literature on Bayesian persuasion (Kamenica and Gentzkow, 2011). First, our model has an interpretation as a model of "self-persuasion". The equilibrium characterization in Section 1.2 (Lemma 1) shows that equilibrium can be characterized by a Bayes-consistent distribution of stopped beliefs p_τ with support on two posteriors $p_l \leq p_0 \leq p_h$. This distribution maximizes

$$\mathrm{E}(V(p))$$

across all Bayes-consistent distributions of posteriors, for $V(p) = \max_{a \in \{x,y\}} U(a, p; r)$. This is as in standard *interpersonal* Bayesian persuasion, just that the sender and receiver are two selves of the same agent (compare to Kamenica and Gentzkow, 2011). A notable difference concerns the interpretation of the commitment assumption. In our information acquisition setting, commitment means the agent cannot strategically lie about the observed information to a future self. In many *intrapersonal* contexts, this seems more natural than the opposite assumption.

Second, we provide a direct test of our self-persuasion model. Few prior papers have tested Bayesian persuasion models (Nguyen, 2017; Au and Li, 2018; Aristidou *et al.*, 2023; Fréchette *et al.*, 2022), and none a self-persuasion model. We use our experiment data to impute the empirical distribution of the information strategies—as given by the cutoffs p_h and p_l—and how they vary with the selfish incentive; see Figure 3. Appendix A.4 details how we constructed the empirical cutoff distributions. We find that among the responsive types, the prediction of Theorem 1 holds, that is, the distributions of the cutoff p_h and p_l are shifted downwards when there is an egoistic motive (one-sided Kolmogorov-Smirnov test, $p = 0.045$ for p_h and $p = 0.074$ for p_l).[35]

Figure 3: The CDF of the imputed persuasion cutoffs

(a) Upper cutoffs p_h

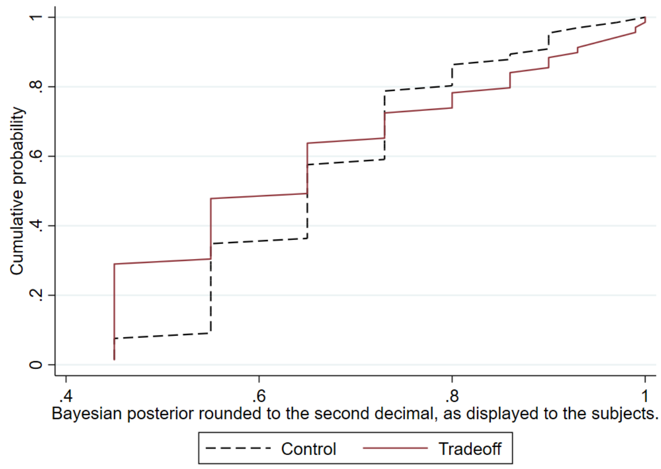

(b) Lower cutoffs p_l

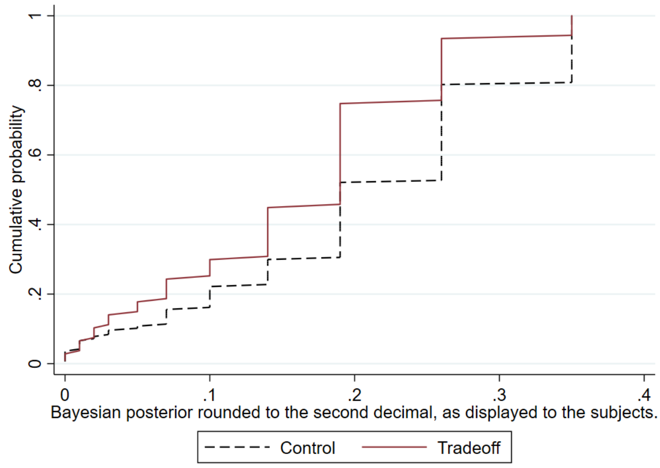

Third, the model variant with a self-signaling component mentioned when discussing the motivated beliefs literature is an example of Bayesian persuasion with signaling (see Heese and Liu, 2023). It features a sender with a private type and (psychological) preferences about the receiver's belief about her type,

[35]In our data, most dictators behave responsively (88%). Most unresponsive dictators are those in *Tradeoff* who choose the egoistic option x despite having received more bad news (77%).

complementing work on persuasion of a receiver with psychological preferences by Lipnowski and Mathevet (2018); see also Schweizer and Szech (2018).

5 Conclusion

We have analyzed decision-makers for whom information might reconcile a motive to feel moral and a competing selfish motive. We developed a novel experimental paradigm and a suitable theoretical framework to understand how the tradeoff between the two motives shapes information acquisition and welfare consequences. Both speak to real-life situations where single observations yield partial, inconclusive information and individuals face dynamic information acquisition decisions.

The main result is that the presence of a selfishly optimal action makes individuals strategically exploit their ability to sample information: They "fish for good news". In tendency, they seek and wait for positive information that aligns their selfishly optimal action with moral considerations and eagerly choose it once such positive information arrives.

Here, "strategic moral ignorance" (compare to Dana *et al.*, 2007) shows up as part of a more general strategy that includes a counterpart of "information seeking". Ample prior empirical work has provided insights on "strategic information avoidance" (see, e.g., Golman *et al.*, 2017), whereas "motivated information seeking" is a previously unexplored phenomenon and may deserve further attention.

Our welfare analysis highlights a novel and potentially counterintuitive aspect of prosocial behavior that arises under incomplete information. Namely, when a decision-maker has additional selfish incentives, this sometimes does not worsen and, in theory, may even improve outcomes for others affected. The reason is that these incentives make people change how they acquire information about the externalities of their choice on others before making it.

This welfare observation may speak to current debates: For example, minimizing personal incentives—such as when targeting perfect neutrality in a hiring decision—may not benefit and, in theory, even reduce the likelihood of selecting the most appropriate candidate.

Appendix

A Empirical

A.1 The difference in difference in Finding 1

To test the difference in the difference in Finding 1, we estimate the following logistic regression for the probability to continue acquiring information after the first piece of information and the first two pieces of information:

$$logit(continue) = b_1\text{Tradeoff} + b_2\text{good} + b_{12}\text{Tradeoff} \cdot \text{Good} + c, \tag{12}$$

where "Tradeoff" is a factor variable for the treatment; "Good" is a factor variable for whether the dictator has acquired more good news or bad news. Table 3 presents the regression estimates.

The interaction effect between *Tradeoff* and having acquired more good news, i.e., $b_{12|\text{good}}$, is significantly negative, showing that the treatment effect on the probability to continue acquiring information significantly differs after having received more good and after having received more bad news.

Table 3: Logistic regression estimates (with p in the brackets)

Coef.	(1) After the 1st piece of info	(2) After the first 2 pieces of info
b_1	.45	1.17
	(.27)	(.008)
$b_{2\|\text{good}}$	1.77	1.64
	(.006)	(.001)
$\mathbf{b_{12\|\text{good}}}$	**-2.46**	**-1.46**
	(.001)	**(.049)**
c	1.82	.32
	(.00)	(.26)
N	458	409
Chi2 p	.00	.00

A.2 Finding 2: The Cox model results

Table 4: The Cox proportional hazards model results (with p in brackets)

Coef.	Covariate	(1)	(2)
$\hat{\beta}_1$	***Tradeoff***	-.29	-.21
		(.02)	(.09)
$\hat{\beta}_{12}$	**Tradeoff** ×		
	Good news dominance	.43	.28
		(.03)	(.14)
	Balanced	-.35	-.53
		(.35)	(.15)
$\hat{\beta}_2$	Good news dominance	-.14	-.09
		(.38)	(.59)
	Balanced	-.52	-.51
		(.03)	(.03)
Stratified by: gender, IQ, prosociality		Yes	No
Violation of the proportional hazards assumption		**No**	**Yes**
Control variable: belief accuracy		Yes	Yes
Observations (individuals)		458	458
Chi2 p-value		.00	.00

This table presents the estimated *coefficients* of the Cox model in (9), with standard errors clustered at the individual level. In the brackets, we report the p value of the corresponding coefficient estimate. The dependent variable is the hazard to stop acquiring information, and the key coefficients of interests are $\hat{\beta}_1$ and $\hat{\beta}_{12}$. $\exp(\hat{\beta}_1)$ reflects the treatment effect on the dictators' hazards to stop acquiring further information, given information histories dominated by bad news; and $\exp(\hat{\beta}_1 + \hat{\beta}_{12}|\text{Good news dominance})$ reflects the treatment effect on the hazards, given information histories dominated by good news (see the derivation in Equation (11)). The violation of the proportional hazard assumption of the Cox model (PH) is tested using Schoenfeld residuals. Without stratification, the PH assumption is violated, as shown in column (2), implying that the baseline hazard might differ for subgroups of the sample. Hence, we follow the literature and use stratification to allow the baseline hazard to vary according to the control variables, i.e., gender, the prosocial types (categorized by the SVO test), and the cognitive ability (categorized by the score in a 5-element Raven's matrices test). With the stratification, PH is no longer violated. We also control for the belief accuracy, measured by the Brier score of the beliefs after each draw (see Footnote 28). The reported likelihood Chi-square statistic is calculated by comparing the deviance ($-2\times$ log-likelihood) of each model specification against the model with all covariates dropped. We use the Breslow method to handle ties.

A.3 Robustness: *Force* treatments

Difference in difference. Below, we estimate the logistic regression in (12) in the *Force* treatment, where there is no self-selection into the sample. We find that consistent with Finding 1, having received more good news and being randomly assigned to *Tradeoff* has a significantly negative interaction effect on the dictators' tendency to continue acquiring information.

Table 5: Logistic regression estimates (with p in the brackets)

Coef.	(1)
b_1	.48
	(.53)
$b_{2\|good}$	1.52
	(.17)
$\mathbf{b_{12\|good}}$	**-2.74**
	(.039)
c	2.03
	(.00)
N	161
Chi2 p	.05

First Piece of Information. Further, we find: After having received a piece of good news, significantly fewer dictators in *Force-Tradeoff* continue acquiring information; while having received a first piece of bad news, the proportion of dictators who continue acquiring information is similar between *Force-Tradeoff* and *Force-Control*. This result aligns with the model; see Prediction 1 and the sentence thereafter.

A.4 Bayesian Self-Persuasion: Imputed Cutoff Distributions

We infer from the data the distribution of the optimal strategies, characterized by upper and lower belief cutoffs, and we compare them between treatments.

Recall that in the model, we analyze the information acquisition behavior of an agent who does not avoid information completely. Theorem 1 considers "responsive" types, i.e., those who choose x if they stop at a posterior weakly above the prior or y if they stop at a posterior weakly below the prior. In our data, considering the subjects who do not avoid information completely, we find that the large majority of subjects behaves responsively (405 out of 458; *Control*: 225 out of 234; *Tradeoff*: 180 out of 224).[36]

Figure 3 shows the empirical cumulative distribution functions of the lower belief cutoff p_l (3b) and the upper belief cutoff p_h (3a). Both CDFs reflect the

[36] In the *Control* treatment, 5 dictators choose y after having received more good news, 4 dictators choose x after having received more bad news. In the *Tradeoff* treatment, 4 dictators choose y after having received more good news, 37 subjects choose x after having received more bad news.

Figure 4: Proportion of Dictators Who Continue Acquiring Information After the First Piece in *Force*

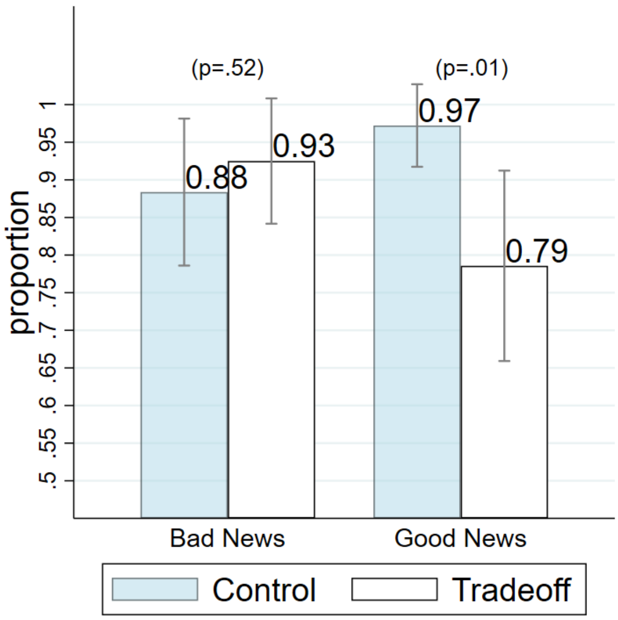

This figure presents the proportion of dictators in *Force* who continue acquiring information after the first piece of information (N=161). In the parentheses, we report the p-value from the Chi-square test.

dictators who acquire some information. The CDF of the upper belief cutoff reflects the stopped beliefs of the dictators who stop weakly above the prior and choose x. The CDF of the lower belief cutoff reflects the stopped beliefs of the dictators who stop information acquisition at posterior beliefs weakly below the prior and choose y.

Figure 3 show that the belief cutoffs are systematically lower in *Tradeoff*, as predicted by the model in Theorem 1 (one-sided Kolmogorov-Smirnov test, $p = .045$ for p_h and $p = 0.074$ for p_l).

B Theory

B.1 Preliminaries for the proofs

First, we establish two claims that we will use to prove both Lemma 1 and Lemma 2. For this, recall the definition of the cutoff beliefs p_l and p_h following the statement of Lemma 2.

Claim 1 *Let $p_t \in [p_l, p_h]$. For any continuation strategy τ,*

$$\mathrm{E}(V(p_\tau)|(Z_s)_{s\leq t}) \leq \bar{V}(p_t) \tag{13}$$

Proof. We have

$$\mathrm{E}(V(p_\tau)|(Z_s)_{s\leq t}) \quad \leq \quad \mathrm{E}(\bar{V}(p_\tau)|(Z_s)_{s\leq t}) \leq \bar{V}(\mathrm{E}(p_\tau)|(Z_s)_{s\leq t}) = \bar{V}(p_t),$$

where we used that $V \leq \bar{V}$ for the first inequality and Jensen's inequality for the second inequality. For the final equality, we use that $\mathrm{E}(p_\tau|(Z_s)_{s\leq t}) = p_t$ by Doob's optional stopping theorem.[37] ■

Now, consider the candidate equilibrium strategy τ^* where the agent continues to observe the information process as long as $p_l < p_t < p_h$, and stops whenever $p_t \leq p_l$ or $p_t \geq p_h$.

Claim 2 *Let $p_t \in [p_l, p_h]$. The strategy τ^* satisfies*

$$\mathrm{E}(V(p_{\tau^*})|(Z_s)_{s\leq t}) = \bar{V}(p_t) \tag{14}$$

Proof. We consider two cases: if $V(p_0) = \bar{V}(p_0)$, by definition, $p_h = p_l = p_0$ and the agent immediately stops at $t = 0$, i.e., $\Pr(p_{\tau^*} = p_0) = 1$, which directly yields the result in this case. If $V(p_0) < \bar{V}(p_0)$, then, $\bar{V}$ is linear on all open intervals $I' \subseteq [\epsilon, 1-\epsilon]$ satisfying $p_0 \in I'$ and $V(p) < \bar{V}(p)$ for all $p \in I'$, by its minimality. Now, (p_l, p_h) is the largest such interval, which implies that V and $\bar{V}$ must coincide at p_l and p_h,[38]

$$V(p_h) = \bar{V}(p_h), \quad \text{and} \quad V(p_l) = \bar{V}(p_l). \tag{15}$$

Finally, for any $p_t \in [p_l, p_h]$,

$$\begin{aligned}
\mathrm{E}(V(p_{\tau^*})|(Z_s)_{s\leq t}) &= \Pr(p_{\tau^*} = p_h|(Z_s)_{s\leq t})V(p_h) + \Pr(p_{\tau^*} = p_l|(Z_s)_{s\leq t})V(p_l) \\
&= \Pr(p_{\tau^*} = p_h|(Z_s)_{s\leq t})\bar{V}(p_h) + \Pr(p_{\tau^*} = p_l|(Z_s)_{s\leq t})\bar{V}(p_l) \\
&= \bar{V}(p_t),
\end{aligned}$$

where we used (15) for the equality on the second line. For the equality on the third, we used the earlier observation that $\bar{V}$ is linear on (p_l, p_h) together with Bayes' law. ■

[37] See e.g., Revuz and Yor (2013).

[38] One checks that this is also true if $(p_l, p_h) = (\epsilon, 1-\epsilon)$ by the minimality of $\bar{V}$.

B.2 Proof of Lemma 1

Let τ^* be the candidate equilibrium strategy where the agent continues to observe the information process as long as $p_l < p_t < p_h$, and stops whenever $p_t \leq p_l$ or $p_t \geq p_h$. Claim 1 and Claim 2 together imply that at any point of time, following τ^* is weakly optimal, hence τ^* is a Nash equilibrium. This proves Lemma 1.

B.3 Proof of Lemma 2

Take any Nash equilibrium τ^{**} in which the agent stops observing the information process whenever she is indifferent between stopping and continuing. It follows from Claim 1 and Claim 2 that, when $p_t \in (p_l, p_h)$, it is strictly optimal for the agent to continue acquiring information: stopping yields $V(p_t)$, which is strictly smaller than $\bar{V}(p_t)$, and there is a continuation strategy which yields $\bar{V}(p_t)$ by Claim 2. When $p_t \in \{p_l, p_h\}$, if the agent would stop acquiring information, her payoff would be $V(p_t) = \bar{V}(p_t)$, given (15). Thus, it follows from Claim 1 that it is weakly optimal to stop acquiring information, so the agent stops under τ^{**}. Finally, we conclude that τ^{**} is identical to τ^* (see the proof of Lemma 1 for the definition of the equilibrium τ^*.)

B.4 Proof of Lemma 3

Take the strategy τ' where the agent never stops observing the information process (unless $p_t \leq \epsilon$ or $p_t \geq 1-\epsilon$, and she has to stop). Given $\epsilon \approx 0$, she acquires almost complete information about the state. Note that her expected utility when doing so is $\mathrm{E}(\mathrm{V}(\mathrm{p}_{\tau'})) \approx (1-p_0)V(0) + p_0 V(1) \geq (1-p_0)u(y,1) + p_0(u(x,1)+r)$ since she can almost always choose y in the state when y is harmless and x in the state when x is harmless. Given that $u(x,1) = 0$ and $u(y,1) = 0 < r$, we have $(1-p_0)V(0) + p_0 V(1) \geq (1-p_0)u(y,1) + p_0 r > u(y,1)$. It follows that the equilibrium strategy τ^* given by the cutoff beliefs p_l and p_h must yield a payoff strictly larger than $u(y,1)$ as well when ϵ is sufficiently small, that is $\mathrm{E}(V(p_{\tau^*})) > u(y,1)$.

First, this implies that the agent does not choose y at p_h when ϵ is sufficiently small: suppose she does so; then she will also choose y at $p_l < p_h$ since at p_l she is more certain that y is harmless, ceteris paribus. However, when she always chooses y, her payoff is weakly smaller than $u(y,1)$ since $U(y,p;r) = u(y,1-p) \leq u(y,1)$

for all p.

Second, this implies that $V(p_h) > u(y,1)$ when ϵ is sufficiently small: suppose that $V(p_h) \leq u(y,1)$. Then, also $V(p_l) = \max_{a\in\{x,y\}} U(a,p_l;r) \leq u(y,1)$ since $U(y,p;r) = u(y,1-p) \leq u(y,1)$ for all p and $U(x,p_l;r) \leq U(x,p_h;r) \leq u(y,1)$. However, $V(p_h) \leq u(y,1)$ and $V(p_l) \leq u(y,1)$ together imply $\mathrm{E}(V(p_{\tau^*})) \leq u(y,1)$, which contradicts with the observation $\mathrm{E}(V(p_{\tau^*})) > u(y,1)$ when ϵ is small enough. Given that we assumed that the agent weakly prefers y at p_l, we have $V(p_l) = u(y,1-p_l) \leq u(y,1)$. We conclude that $V(p_h) > V(p_l)$ since $V(p_h) > u(y,1)$.

B.5 Equilibrium when there is no egoistic motive ($r = 0$)

The following result formally describes the equilibrium when $r = 0$. It shows that the agent acquires information until she reaches her threshold level of certainty $l(a)$, unless $l(a) > 1-\epsilon$.[39] Here, it may be the case that the threshold level of certainty is already reached at the prior for one of the options, so that she stops directly.

Lemma 5 *Let $r = 0$. If $\max_{a\in\{x,y\}} l(a) \leq 0.5$ or $p_0 \in (1-l(y), l(x))^c$, then $p_l = p_0 = p_h$. If $p_0 \in [1-l(y), l(x)]$, then $p_l = \max\{\epsilon, 1-l(y)\}$ and $p_h = \min\{l(x), 1-\epsilon\}$.*

B.6 Proof of Theorem 1

Take any "responsive type" u, meaning that it is strictly optimal for the type to choose the option y at the belief p_l and the option x at the belief p_h when $r > 0$ and when $r = 0$. First, recall from Lemma 4 that $p_l(r) = \epsilon$. Hence $p_l(r) \leq p_l(0)$, which shows the right inequality of Theorem 1.

To show the left inequality of Theorem 1, first we note that it follows from Lemma 5 that $p_h(0) \in \{1-\epsilon, l(x), p_l(0)\}$. When $p_h(0) = p_l(0)$, the agent is not responsive, so the precondition of the theorem is not fulfilled. It remains to establish that $p_h(r) \leq \min\{1-\epsilon, l(x)\}$. Clearly $p_h(r) \leq 1-\epsilon$ since the agent has to stop at $1-\epsilon$ necessarily. Finally, we show that $p_h(r) \leq l(x)$. Given the definition of $l(x)$ in (5), we know that either $l(x) = 1$ or $\frac{\partial u(x,q)}{\partial q} = 0$ for all $q > l(x)$.[40] If $l(x) = 1$, clearly $p_h(r) \leq l(x)$. For the second case, observe that the derivative of the objective function with respect to $p_h(r)$, which is the left side of (4), is strictly

[39] Recall that, for technical reasons, we restrict the agent's strategies, imposing that the agent has to stop at $p_t = \epsilon$ and $p_t = 1-\epsilon$ for $\epsilon \approx 0$ arbitrarily small.

[40] In particular, $l(x) < 1$ implies the continuous differentiability of $u(x,q)$ for $q > l(x)$.

negative for any $p > l(x)$ when $r > 0$. This follows since $\frac{\partial u(x,q)}{\partial q} = u(x,q) = 0$ for all $p > l(x)$. Hence, $p_h(r) \leq l(x)$. This finishes the proof of the claim that $p_h(r) \leq \min\{1-\epsilon, l(x)\}$, and thereby the proof of Theorem 1.

B.7 Parametric example: Concave belief utility

We give a parametric example, with parameters l and α; further, the curvature of the belief utility varies depending on an elasticity parameter $\rho \in (-\infty, 1]$:

$$u(a,q) = \begin{cases} 0 \quad \text{if} \quad q > l, \\ -\alpha(1 - \frac{q}{l}^{1-\rho}) \quad \text{if} \quad q \leq l. \end{cases}$$

A calculation shows that

$$p_h(r) = \min(l, lz). \tag{16}$$

for $z = \left[\frac{1-\frac{r}{\alpha}}{\rho}\right]^{\frac{1}{1-\rho}}$ (which is a positive number for $r < \alpha$). Suppose that $1 - l < p_0 < l$ so that the agent acquires some information without egoistic motive. Given Theorem 2 and the discussion thereafter, the welfare effect can only be strictly positive if $r < \alpha$, and $z \geq 1$ or

$$\begin{aligned} 1 - lz &< \frac{1-l}{1-l+p_0} \\ \Leftrightarrow (1+p_0)(1-lz) + l^2 - 1 &< 0. \end{aligned} \tag{17}$$

Note that the function on the left is continuous in l and the inequality (17) is not satisfied for $l = 1$. This implies that there is a cutoff $l^* < 1$ so that (17) is not satisfied if $0 < l^* \leq l$. The function z is increasing in α, which implies that (17) is more easily satisfied with higher α.

References

Akerlof, G. A. and Kranton, R. E. (2000). Economics and identity. *The Quarterly Journal of Economics*, **115** (3), 715–753.

Aristidou, A., Coricelli, G., Vostroknutov, A. *et al.* (2023). *Incentives or persuasion? an experimental investigation.* Working paper.

Au, P. H. and Li, K. K. (2018). *Bayesian persuasion and reciprocity: theory and experiment*. Working paper.

Bartling, B., Engl, F. and Weber, R. A. (2014). Does willful ignorance deflect punishment?–an experimental study. *European Economic Review*, **70**, 512–524.

Battigalli, P. and Dufwenberg, M. (2007). Guilt in games. *American Economic Review*, **97** (2), 170–176.

Bénabou, R. and Tirole, J. (2006). Incentives and prosocial behavior. *American Economic Review*, **96** (5), 1652–1678.

Blossfeld, H.-P., Rohwer, G. and Schneider, T. (2019). *Event History Analysis with Stata*. Routledge.

Bock, O., Baetge, I. and Nicklisch, A. (2014). Hroot: Hamburg registration and organization online tool. *European Economic Review*, **71**, 117–120.

Bodner, R. and Prelec, D. (2003). Self-signaling and diagnostic utility in everyday decision making. *The Psychology of Economic Decisions*, **1**, 105–26.

Bretagnolle, J. and Huber-Carol, C. (1988). Effects of omitting covariates in cox's model for survival data. *Scandinavian Journal of Statistics*, pp. 125–138.

Brier, G. W. (1950). Verification of forecasts expressed in terms of probability. *Monthly weather review*, **78** (1), 1–3.

Brown, M., Flinn, C. J. and Schotter, A. (2011). Real-time search in the laboratory and the market. *American Economic Review*, **101** (2), 948–974.

Camerer, C. F. (2011). *Behavioral game theory: Experiments in strategic interaction*. Princeton university press.

Caplin, A. and Dean, M. (2013). Rational inattention, entropy, and choice: The posterior-based approach.

—, — and Martin, D. (2011). Search and satisficing. *American Economic Review*, **101** (7), 2899–2922.

Card, D., Chetty, R. and Weber, A. (2007). The spike at benefit exhaustion: Leaving the unemployment system or starting a new job? *American Economic Review*, **97** (2), 113–118.

Chan, J., Lizzeri, A., Suen, W. and Yariv, L. (2018). Deliberating collective decisions. *The Review of Economic Studies*, **85** (2), 929–963.

Chen, S. and Heese, C. (2020). *Motivated Information Acquisition in Social Decisions*. Crc discussion paper.

Cleves, M., Gould, W., Gutierrez, R. and Marchenko, Y. (2010). *An Introduction to Survival Analysis Using Stata*. College Station, TX, Stata Press.

Cox, D. R. (1972). Regression models and life-tables. *Journal of the Royal Statistical Society: Series B (Methodological)*, **34** (2), 187–202.

— (1975). Partial likelihood. *Biometrika*, **62** (2), 269–276.

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.

Di Tella, R., Perez-Truglia, R., Babino, A. and Sigman, M. (2015). Conveniently upset: Avoiding altruism by distorting beliefs about others' altruism. *American Economic Review*, **105** (11), 3416–42.

Ditto, P. H. and Lopez, D. F. (1992). Motivated skepticism: Use of differential decision criteria for preferred and non-preferred conclusions. *Journal of Personality and Social Psychology*, **63** (4), 568.

Drerup, T., Enke, B. and Von Gaudecker, H.-M. (2017). The precision of subjective data and the explanatory power of economic models. *Journal of Econometrics*, **200** (2), 378–389.

Eil, D. and Rao, J. M. (2011). The good news-bad news effect: Asymmetric processing of objective information about yourself. *American Economic Journal: Microeconomics*, **3** (2), 114–38.

Exley, C. L. and Kessler, J. B. (2023). Information avoidance and image concerns. *The Economic Journal*, **133** (656), 3153–3168.

— and — (2024). Motivated errors. *American Economic Review*, **114** (4), 961–987.

Falk, A. and Fischbacher, U. (2006). A theory of reciprocity. *Games and economic behavior*, **54** (2), 293–315.

— and Szech, N. (2016). *Pleasures of Skill and Moral Conduct*. Working paper.

Fehr, E. and Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *The quarterly journal of economics*, **114** (3), 817–868.

Feiler, L. (2014). Testing models of information avoidance with binary choice dictator games. *Journal of Economic Psychology*, **45**, 253–267.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental economics*, **10** (2), 171–178.

Fréchette, G. R., Lizzeri, A. and Perego, J. (2022). Rules and commitment in communication: An experimental analysis. *Econometrica*, **90** (5), 2283–2318.

Gabaix, X., Laibson, D., Moloche, G. and Weinberg, S. (2006). Costly information acquisition: Experimental analysis of a boundedly rational model. *American Economic Review*, **96** (4), 1043–1068.

Geanakoplos, J., Pearce, D. and Stacchetti, E. (1989). Psychological games and sequential rationality. *Games and Economic Behavior*, **1** (1), 60–79.

Gino, F., Norton, M. I. and Weber, R. A. (2016). Motivated bayesians: Feeling moral while acting egoistically. *Journal of Economic Perspectives*, **30** (3), 189–212.

Gneezy, U., Saccardo, S., Serra-Garcia, M. and van Veldhuizen, R. (2020). Bribing the self. *Games and Economic Behavior*, **120**, 311–324.

Golman, R., Hagmann, D. and Loewenstein, G. (2017). Information avoidance. *Journal of Economic Literature*, **55** (1), 96–135.

Grossman, Z. (2014). Strategic ignorance and the robustness of social preferences. *Management Science*, **60** (11), 2659–2665.

— and van der Weele, J. (2017). Self-image and willful ignorance in social decisions. *Journal of the European Economic Association*, **15** (1), 173–217.

Haisley, E. C. and Weber, R. A. (2010). Self-serving interpretations of ambiguity in other-regarding behavior. *Games and Economic Behavior*, **68** (2), 614–625.

Heese, C. and Liu, S. (2023). *Image-Building Persuasion*. Working paper.

Hoelzemann, J. and Klein, N. (2021). Bandits in the lab. *Quantitative Economics*, **12** (3), 1021–1051.

Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, **101** (6), 2590–2615.

Kőszegi, B. (2006). Ego utility, overconfidence, and task choice. *Journal of the European Economic Association*, **4** (4), 673–707.

Kunda, Z. (1990). The case for motivated reasoning. *Psychological bulletin*, **108** (3), 480.

Lipnowski, E. and Mathevet, L. (2018). Disclosure to a psychological audience. *American Economic Journal: Microeconomics*, **10** (4), 67–93.

Loewenstein, G. (1987). Anticipation and the valuation of delayed consumption. *The Economic Journal*, **97** (387), 666–684.

Luce, R. D. (1991). *Response times: Their role in inferring elementary mental organization*. Oxford University Press.

Michelacci, C. and Ruffo, H. (2015). Optimal life cycle unemployment insurance. *American Economic Review*, **105** (2), 816–59.

Möbius, M. M., Niederle, M., Niehaus, P. and Rosenblat, T. S. (2022). Managing self-confidence: Theory and experimental evidence. *Management Science*.

Morris, S. and Strack, P. (2019). The wald problem and the equivalence of sequential sampling and static information costs.

Murphy, R. O., Ackermann, K. A. and Handgraaf, M. (2011). Measuring social value orientation. *Judgment and Decision Making*, **6** (2), 771–781.

Nguyen, Q. (2017). *Bayesian persuasion: evidence from the laboratory.* Working paper.

Rabin, M. (1995). *Moral Preferences, Moral Constraints, and Self-Serving Biases.* Working paper.

Ratcliff, R. and McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural computation*, **20** (4), 873–922.

— and Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological review*, **111** (2), 333.

Raven, J. C. *et al.* (1998). *Raven's Progressive Matrices and Vocabulary Scales.* Oxford Pyschologists Press.

Reshidi, P., Lizzeri, A., Yariv, L., Chan, J. H. and Suen, W. (2021). *Individual and collective information acquisition: An experimental study.* Working paper.

Revuz, D. and Yor, M. (2013). *Continuous Martingales and Brownian Motion*, vol. 293. Springer Science & Business Media.

Schlag, K. H. *et al.* (2013). Eliciting probabilities, means, medians, variances and covariances without assuming risk neutrality. *Theoretical Economics Letters*, **3** (1), 38–42.

Schweizer, N. and Szech, N. (2018). Optimal revelation of life-changing information. *Management Science*, **64** (11), 5250–5262.

Serra-Garcia, M. and Szech, N. (2021). The (in) elasticity of moral ignorance. *Management Science*.

Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, **69** (1), 99–118.

Spiekermann, K. and Weiss, A. (2016). Objective and subjective compliance: A norm-based explanation of 'moral wiggle room'. *Games and Economic Behavior*, **96**, 170–183.

Swensson, R. G. (1972). The elusive tradeoff: Speed vs accuracy in visual discrimination tasks. *Perception & Psychophysics*, **12** (1), 16–32.

Van der Weele, J. J. (2014). Inconvenient truths: Determinants of strategic ignorance in moral dilemmas.

Vu, L., Soraperra, I., Leib, M., van der Weele, J. and Shalvi, S. (2023). Ignorance by choice: A meta-analytic review of the underlying motives of willful ignorance and its consequences. *Psychological Bulletin*, **149** (9-10), 611.

Wald, A. (1945). Sequential tests of statistical hypotheses. *The annals of mathematical statistics*, **16** (2), 117–186.

Zimmermann, F. (2020). The dynamics of motivated beliefs. *American Economic Review*, **110** (2), 337–61.