Online appendix

C Empirical

C.1 Experimental design

Figure 5: The Noisy Information Generators

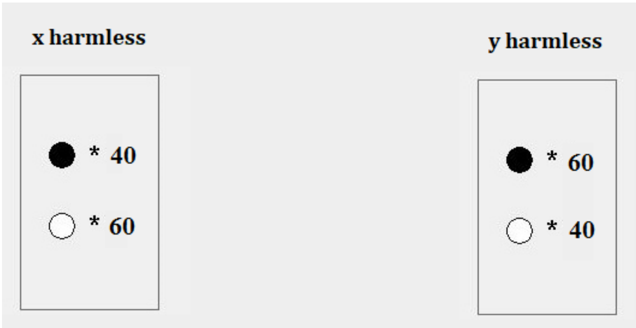

Figure 6: Screenshot of the Information Stage

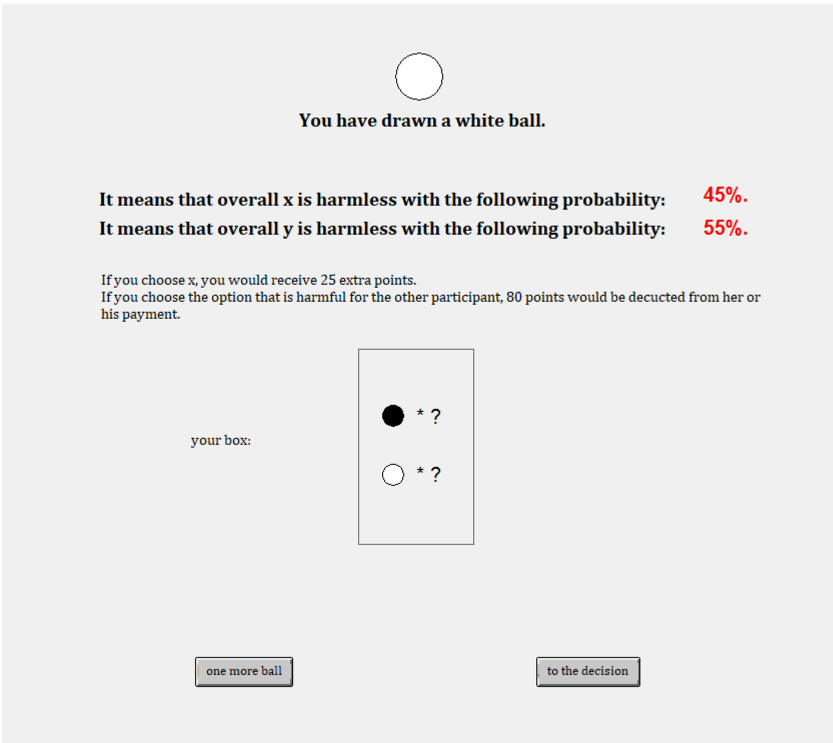

C.2 Summarizing statistics

Table 6: Basic information of the dictators

	N	Good state (%)	female (%)	student (%)	av. age (years)
Tradeoff	250	35	59	94	24
Control	246	36	61	93	24
p-value		.82	.62	.56	.49

This table summarizes the basic characteristics of the dictators in each treatment. We compare these characteristics between *Tradeoff* and *Control*. For the state, gender, and student status we report the p-values of the Chi-Square test. For the dictators' age, we report the p-value of the two-sided t test.

Table 7: Information acquisition behaviour

	median no. balls	av. belief at decision
Tradeoff	6	.33
Control	5	.34
p-value	.24	.30

This table presents the median number of information pieces drawn by the dictators and their average Bayesian posterior beliefs in the Good state. The p-values are of the two-sided Mann-Whitney-U test comparing between *Tradeoff* and *Control*.

Figure 7: Histogram of the number of information pieces drawn

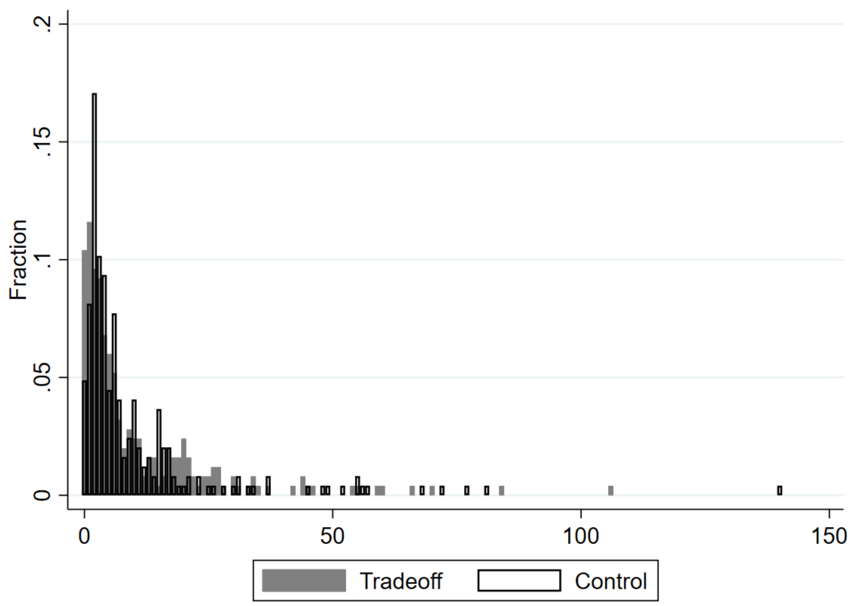

This figure presents the number of information pieces that the dictators drew in *Tradeoff* and *Control* respectively. The distribution is not significantly different between *Tradeoff* and *Control* (two-sided Mann-Whitney-U test, $p = .98$).

C.3 Robustness check of the Cox model estimate: a logistic regression

Using the data at the person-draw level, we estimate the following logistic model as a robustness check of the Cox model estimate and find a result similar to Finding 2 from Section 3.1.2.

$$logit\, h(X) = X_t \cdot b + Z \cdot a + (C + T \cdot c), \tag{18}$$

where $h(X)$ is the probability that the dictator stops acquiring information after that draw; X denotes the same covariates of interest as in the Cox model, i.e.,

$$X \cdot b = \beta_1 \text{Tradeoff} + \beta_2 \text{Info} + \beta_{12} \text{Tradeoff} \times \text{Info}. \tag{19}$$

The control valuables in Z include gender, cognitive ability, prosociality and belief accuracy, all measured in the same way as in the Cox model in Section 3.1.2. T is a vector of time dummies, which captures the time dependency of the probability to stop acquiring information.

When interpreting the results, this logistic model can be viewed as a hazard model in which the covariates proportionally affect the *odds* of stopping the information acquisition (Cox, 1975). Formally, consider such a hazard model,

$$\frac{h(t)}{1-h(t)} = \frac{h_0(t)}{1-h_0(t)} \cdot exp(X_t \cdot b + Z \cdot a).$$

Then,

$$\underbrace{log\left(\frac{h(t)}{1-h(t)}\right)}_{logit\, h(X)} = \underbrace{log\left(\frac{h_0(t)}{1-h_0(t)}\right)}_{C+T\cdot c} + X_t \cdot b + Z \cdot a. \tag{20}$$

Unlike in the framework of the Cox model, the coefficients here cannot be interpreted as hazard ratios. Instead, they should be interpreted as odds ratios. Our prediction that the hazard to stop acquiring information is lower in *Tradeoff* when bad news dominates suggests a negative β_1 in (19). And the prediction that the hazard is higher when good news dominates suggests a positive $\beta_1 + \beta_{12|good}$. Results reported in Table 8 support these predictions.

Table 8: The logistic model results (with p in the brackets)

Coef.	Covariate	(1)	(2)
$\hat{\beta}_1$	***Tradeoff***	-.26	-.22
		(.08)	(.126)
$\hat{\beta}_{12}$	**Tradeoff** $\times$		
	Good news dominance	.36	.35
		(.099)	(.08)
	Balanced	.-.54	-.59
		(.18)	(.14)
$\hat{\beta}_2$	Good news dominance	-.19	-.62
		(.30)	(.00)
	Balanced	-.66	-1.11
		(.02)	(.00)
Control		Yes	No
N (person-draws)		4,658	4,658
Pseudo R2		.07	.05

This table presents the estimated coefficients of the logistic model in (18), with standard errors clustered at the individual level. The dependent variable is the decision 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 dictator's odds to stop acquiring further information, given information histories dominated by bad news. And $\exp(\hat{\beta}_1 + \hat{\beta}_{12}|\text{good})$ reflects the treatment effect on the odds, given information histories dominated by good news. We control for belief accuracy, gender, the prosocial types (categorized by the SVO test), and the cognitive ability (measured by a Raven's matrices test). The time dependency of the odds is accounted for by including a dummy for each period.

C.4 Discussion: Self-selection

Below we explain the self-selection facing our empirical analysis and show evidence that self-selection can only *weaken* our finding of fishing for good news. That is, our results are lower bounds for the effects. We also provide intuition based on the model for why this is the case.

In our experiment, individuals repeatedly decide whether to continue or to stop acquiring information. When one compares between treatments the stopping decision at any point of the information process, dictators who have already stopped earlier are not in the sample. That is, the sample is dynamically self-selected. This could confound the empirical analysis if the self-selected samples differ between treatments.

In Section 3.1.2, we use the Cox proportional hazard model to estimate the effects across all information histories. When there are omitted variables in the Cox model, dynamic selection may confound the analysis. However, Bretagnolle and Huber-Carol (1988) show that this can only lead to *underestimating* the effects of interest.

Additionally, in our experiment it is intuitive why the self-selection weakens the observed "fishing for good news" behaviour: let's first consider the selection into the information process (extensive margin) in the *NoForce* treatment variation. Here, in *Tradeoff*, 25 out of 26 dictators who do not acquire information choose x, while in *Control*, 10 out of 12 dictators who do not acquire information choose y. Had these dictators received a further piece of information *supporting* their dictator decisions, i.e., good news in *Tradeoff* and bad news in *Control*, theory predicts that they would also stop directly and take the same decisions, an effect that would strengthen Finding 1. This prediction follows directly from the cutoff structure of optimal strategies, which are given by beliefs $p_l \leq p_h$.[41] This (theoretical) argument extends to the decisions about further pieces of information, suggesting that the potential self-selection could only weaken the results about the entire information histories (Finding 2).

To directly address potential self-selection on the extensive margin, i.e., the selection into the information acquisition process at the start, we implement a second treatment variation: *Force* and *NoForce*.[42] When we analyze the stopping decision after the first piece of information using the *Force* treatment data, the sample is not self-selected: all the dictators in *Force* receive the first piece of information and are randomly assigned between *Force-Tradeoff* and *Force-Control*.[43] Section A.3 shows that Finding 1 replicates in the Force treatments; in fact, the size of the effects is slightly larger than with the pooled sample.

[41] For example, if an agent prefers to stop at the prior and choose x, she would also stop and choose x at any higher belief $p_t > p_0 = p_h$. An agent who would stop at the prior and choose y, would also stop and choose y at any lower belief $p_t < p_0 = p_l$.

[42] Recall that while in '*NoForce*', the dictators are *not forced* to acquire any information, in *Force* the information stage starts for all dictators, and the dictators acquire at least one piece of information.

[43] To give a sense of the scale of the selection on the extensive margin: 15% and 7% dictators opt to acquire no information in *NoForce-Tradeoff* and *NoForce-Control*.

In summary, the potential issue of dynamic self-selection could only have weakened our result instead of having driven it. Further, in a sample clean of self-selection, we find similar results as in Finding 1.

C.5 Individuals with higher IQ have a higher tendency to fish for good news

Table 9: The Cox model results for about and below median IQ (with p in brackets)

Coef.	Covariate	Above Median (1)	Below Median (2)	All Dictators (3)
$\hat{\beta}_1$	***Tradeoff***	-.38	-.22	-.29
		(.01)	(.25)	(.02)
$\hat{\beta}_{12}$	**Tradeoff** $\times$			
	Good news dominance	.62	.24	.43
		(.02)	(.43)	(.03)
	Balanced	.31	-1.00	-.35
		(.55)	(.08)	(.35)
$\hat{\beta}_2$	Good news dominance	-.14	-.25	-.14
		(.51)	(.31)	(.38)
	Balanced	-1.01	-.23	-.52
		(.01)	(.47)	(.03)
Stratified		Yes	Yes	Yes
Violation of the PH assumption		No	No	No
Control variable: belief accuracy		Yes	Yes	Yes
Observations (individuals)		267	191	458
Chi2 p-value		.00	.00	.00

This table presents the Cox model results for the dictators with above and below median cognitive ability, measured by the number of correctly answered questions in a Raven's matrices test. Standard errors are clustered at the individual level. For comparison, we include the result for the whole sample in Column (3). In Column (1) and (2), the estimation is stratified by gender and prosociality, but not by IQ since we explicitly compare the dictators with above and below median IQ here. In Column (3), the estimation is stratified by gender, prosociality and IQ. The median number of correct answers to the Raven's test is four out of five in our experiment. In this table, the dictators above the median have given correct answers to four or five questions in the Raven's test, and the subjects below the median have correctly answered less than four questions in the Raven's test. The finding is that dictators with a higher cognitive ability have a higher tendency to "fish for good news". For a comprehensive description of the Cox model estimation, please see the description of Table 4.

C.6 The posterior beliefs

In this section, we investigate the dictators' final Bayesian posterior beliefs in the likelihood of the remunerative option being harmless. We compare the belief distributions between *Tradeoff* and *Control*. Our analysis is motivated by the main empirical results (Finding 1 and Finding 2) and Theorem 1. We find that having drawn more good news than bad news, dictators in *Tradeoff* stop acquiring information at a lower belief in the Good state than those in *Control* (0.63 in *Tradeoff*, 0.66 in *Control*, Mann-Whitney-U test $p = 0.10$). In contrast, having drawn more bad news than good news, dictators in *Tradeoff* stop acquiring information at a similar belief in the Bad state than those in *Control* (0.84 in *Tradeoff*, 0.84 in *Control*, Mann-Whitney-U test $p = 0.13$). Figure 8 presents the histogram of the dictators' posterior beliefs.

This asymmetry in the posterior beliefs mirrors the asymmetry in the information acquisition strategy in Finding 2. It is also in line with Theorem 1.

When it comes to the overall posterior beliefs, Bayes' consistency states that the mean posterior beliefs must be equal to the prior belief in both treatments. In accordance with it, we find that the mean posterior beliefs are not significantly different between *Tradeoff* and *Control* (mean: 0.33 in *Tradeoff*, 0.34 in *Control*; student t-test, $p = 0.56$).[44]

[44]This finding is reconciled with the distributional differences by the observation that in *Tradeoff* slightly more dictators ended up with a posterior belief above the prior, although the difference is insignificant (*Tradeoff*: 30%; *Control*: 29%; Chi-Square, $p = .14$).

Figure 8: Histogram of the posterior beliefs

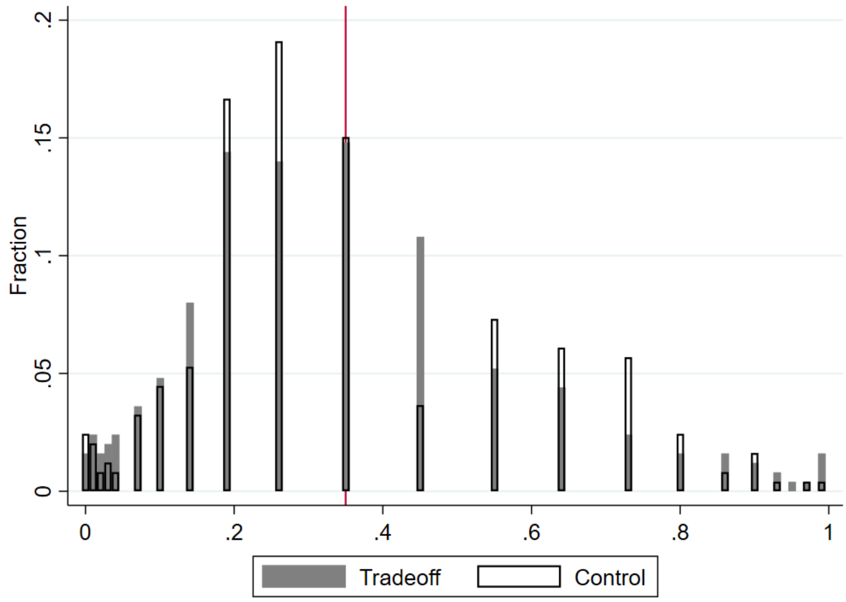

This figure shows the histogram of the dictators' Bayesian posterior beliefs about the likelihood of the Good State when they end the information stage in *Tradeoff* and *Control*. The red vertical line represents the prior belief.

C.7 Dictator Game Decision

Table 10: Dictator game decisions

	Choosing x (%)			Harm (%)
	Good	Bad	Overall	
Tradeoff	82	40	54	32
Control	55	16	30	27
p-value	.00	.00	.00	.17

The first three columns of this table present the proportions of dictators who choose x in *Good* and *Bad* state and in each treatment. Recall that in the *Good* state, x does not harm the receiver, while in the *Bad* state it does. The last column presents the percentage of dictators whose decision reduce the receivers' payoffs in the dictator game. The p-values are from the Chi square tests comparing between *Tradeoff* and *Control* respectively.

C.8 The supplementary stage

After the experiment, we elicited the dictators' posterior beliefs on the state and their SVO scores. We also asked them to answer a questionnaire consisting of questions on their sociodemographics (gender, age, etc.). Five items from Raven's progressive matrices intelligence test are also included.

Elicited beliefs. In the experiment, we display to the dictators the Bayesian posterior belief on the states (rounded to the second decimal) after each draw of information. After a dictator stops acquiring information, we elicit her belief of option x being harmless, given all the information acquired. The elicitation is incentivized by the randomized Quadratic Scoring Rule adapted from Drerup *et al.* (2017) and Schlag *et al.* (2013).

Figure 9: The belief difference

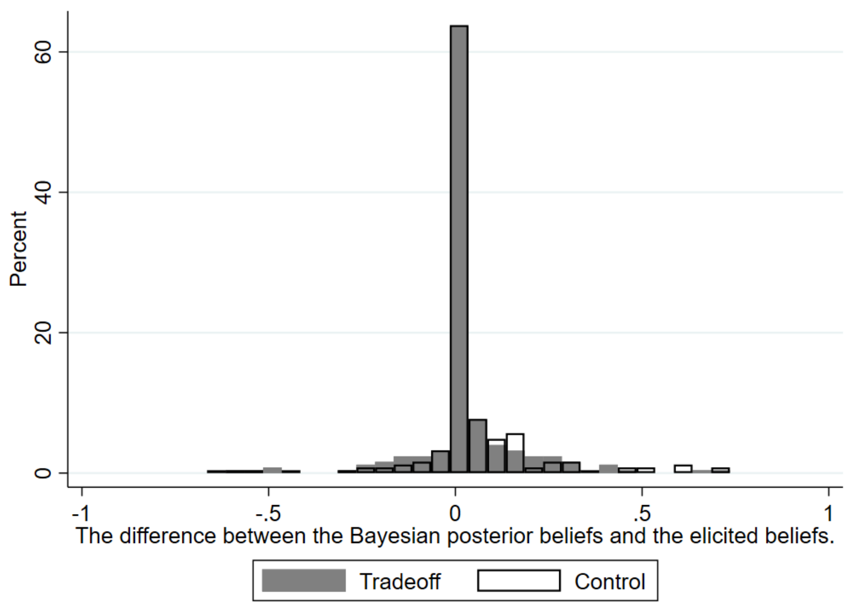

This figure shows the histogram of the difference between the rounded Bayesian posterior belief and the elicited belief that option x is harmless.

In Figure 9, we present the distribution of the difference between the elicited and the rounded Bayesian belief in *Control* and *Tradeoff* respectively. The deviation of the elicited beliefs from the Bayesian beliefs does not significantly differ between *Control* and *Tradeoff* (two-sided Mann-Whitney-U test, $p = .29$). In *Control*, the benchmark treatment, the dictators' elicited beliefs of x being the harmless option are on average 3.30 percentage points higher than the rounded

Bayesian beliefs (two-sided Wilcoxon signed rank test, $p = .00$). In *Tradeoff*, the dictators state beliefs that are higher than the rounded Bayesian beliefs by 1.84 percentage on average (two-sided Wilcoxon signed rank test, $p = .00$) .

Cognitive abilities On average, the subjects answered 3.6 out of 5 questions in Raven's matrices test correctly. There is no significant difference between *Control* and *Tradeoff* treatments (Chi-square $p = .12$). When asked a simple question on probability, in both treatments more than 90% of the subjects answer correctly (*Treatment*: 92%, *Control*: 94%; Chi-square test $p = .51$).[45]

SVO measure. We elicit the social value orientation (SVO) of the subjects as a measure of their altruism. The average SVO score of all the subjects is 20.49, with no significant difference between *Tradeoff* and *Control* treatments (two-sided Mann-Whitney-U test, $p = .84$). According to Murphy *et al.* (2011), 48% subjects are categorized as 'prosocials', 15% 'individualists' and 37% 'competitive type'. The categorization is similar between *Tradeoff* and *Control*, suggesting that the treatment variation in the dictator game has no influence on the SVO measure of altruism (Chi-Square, $p = 1.00$). Further, in *Tradeoff*, prosocial dictators choose the self-benefiting option X significantly less often (Chi-Square, $p = .029$), while in *Control* the three categories of dictators' decisions in the dictator game are similar (Chi-Square, $p = .573$). This result suggests that the SVO measures altruistic traits that are relevant in our experimental setup.

[45] We use the following question to elicit the subjects' understanding of probabilities:
Imagine the following 4 bags with 100 fruits in each. One fruit will be randomly taken out. For which bag, the probability of taking a banana is 40%?
A. A bag with 20 bananas.
B. A bag with 40 bananas.
C. A bag with 0 banana.
D. A bag with 100 bananas.
The correct answer is B.

D Theory and Further Discussion

D.1 Extensive margin: Complete information avoidance

In this section, we discuss two additional results of our model and the respective empirical evidence: the *complete* avoidance of noisy information and of information revealing the state at once. While our experiment focuses on *noisy* information, our preference model (1) can be used to analyze both the acquisition of noisy information that arrives sequentially, as well as information that reveals the state at once. These additional theoretical results are in line with the empirical findings in this paper and also with the empirical findings by Dana *et al.* (2007) and Feiler (2014), as explained further below.

Section D.1.1 and Section D.1.2 present the results, and Section D.1.3 the proofs.

D.1.1 Avoidance of noisy information

Our model predicts that with or without a remunerative option, some agent types move on to the decision without acquiring *any* noisy information.

This may be surprising in the scenario where no option is remunerative. In this scenario, indeed, for the types with $\frac{\partial u(a,q)}{\partial q} > 0$, it is optimal to acquire as much information as possible. In particular, these agents will not avoid information completely. However, other agent types have a threshold level of certainty $l(a) < 1$. Thus, they stop immediately if their threshold is met by the prior, $p_0 > l(x)$ or $p_0 < 1 - l(y)$.

When one of the options is remunerative, the incentives to acquire information are different. First, the the agent would decide not to acquire noisy information only if she would choose the remunerative option at the prior belief. Otherwise, she would "fish for information" that justifies this choice. Second, when considering avoiding information and choosing the remunerative option immediately or acquiring further information, the agent is aware that further information poses an undesirable risk since it might reverse her decision from the remunerative to the non-remunerative option. When this risk outweighs her utility gain from having more certain beliefs, the agent avoids noisy information completely.

Proposition 1 *For any $r = 0$ ($r > 0$) and for any prior $p_0 \in (0, 1)$, there is a set $S_r(p_0)$ of preference types u for which it is (strictly) optimal to avoid information completely.*

Empirical findings. In line with Proposition 1, in the experiment, we find that 15% of the dictators do not acquire any noisy information in the *NoForce-Tradeoff* treatment (Chi-Square $p = 0.00$).[46] Among those, 96% choose the remunerative option x (25/26). Here, theory suggests that these dictators avoid information because they are worried about bad news arriving and indicating that x harms the other.

We find that 7% of the dictators do not acquire any noisy information in the *NoForce–Control* treatment (Chi-Square $p = 0.00$). Among those, only 17% choose the remunerative option x (2/12). Here, theory suggests that the dictators who avoid noisy information completely are satisfied with 65% certainty that y is the harmless option.

D.1.2 Avoidance of Information Revealing the State at Once

Proposition 2 shows that when there is a remunerative option, then, for any prior belief, there are types of dictators who would avoid information that reveals the state all at once.

Proposition 2
1. *Take $r \geq 0$. For any prior $p_0 \in (0, 1)$, there is a set of preference types u that prefer no information over receiving a signal that perfectly reveals the state.*
2. *Take any prior beliefs $p_0' < p_0 \in (0, 1)$. Take a type where u is strictly concave in the second argument. If she prefers no information over receiving a signal that perfectly reveals the state when holding the prior belief p_0, she also does so when holding the prior belief p_0.*

Proposition 2 is consistent with the empirical finding of Dana *et al.* (2007), who, in a dictator experiment similar to ours, find that a significant fraction of dictators avoids information that reveals the *ex-ante* unknown state all at once. Feiler (2014) further documents that the fraction of dictators who avoid such perfectly revealing information increases with the dictator's prior belief that a self-benefiting

[46] Recall that in Force, it is not feasible to stop immediately.

option has no negative externality. The second part of Proposition 2 shows that the model also predicts this finding for a large class of preference types.

D.1.3 Proof of Proposition 1

Take $r = 0$. For any prior $p_0 \in (0,1)$, there is an open set of thresholds $l(x) \in (0,1)$ and $l(y) \in (0,1)$ such that $p_0 > l(x)$ or $p_0 > 1 - l(y)$. The claim for $r = 0$ follows then from Lemma 5. Take $r > 0$. It follows from the characterization of the belief cutoffs after Lemma 2 that $p_h \leq l(x)$. Hence, if $p_0 > l(x)$, the agent stops acquiring information immediately in the equilibrium given by p_h and p_l.

D.1.4 Proof of Proposition 2

The first item of Proposition 2 is a corollary of Proposition 1: It says that a preference type prefers to receive no information over all possible information acquisition strategies, including those given by the belief cutoffs $p_l = \epsilon$ and $p_h = 1 - \epsilon$, which yield information arbitrarily close to a signal that reveals the state perfectly, for $\epsilon \approx 0$ arbitrarily close to 0.

Now, we prove the second item of Proposition 2. Fix a prior $p_0 \in (0,1)$. Take any prior belief $p_0' < p_0 \in (0,1)$. Consider an agent type with u strictly concave in the second argument. Suppose that, given the prior p_0', she prefers to receive no information over receiving a signal that perfectly reveals the state. We show that the agent also prefers to avoid information when the prior is p_0. There are two cases. In the first case, $x = \arg\max_{a \in \{x,y\}} U(a, 0; r)$. This implies that the agent strictly prefers x over y at any belief $p \in [0,1]$. The strict concavity of u implies

$$u(x, p_0) > p_0 u(x, 1) + (1 - p_0) u(x, 0), \tag{21}$$

which is equivalent to

$$V(p_0) > p_0 V(1) + (1 - p_0) V(0), \tag{22}$$

which shows that the agent strictly prefers to avoid information at the prior belief p_0. In the second case, $y \in \arg\max_{a \in \{x,y\}} U(a, 0; r)$. Since we assumed that the agent avoids information given the prior belief p_0',

$$V(p_0') \quad > \quad (1 - p_0') V(0) + p_0' V(1) \tag{23}$$

Now, we use that the agent prefers x at p_h given Lemma 3, and that $V(0) = \max_{a\in\{x,y\}} U(a, 0; r) = U(y, 0; r) = u(y, 1) = 0$. Thus, (23) implies

$$r + u(x, p_0') \quad > \quad rp_0'. \tag{24}$$

Rearranging,

$$r > -\frac{u(x, p_0')}{1 - p_0'} \tag{25}$$

It follows from the concavity of u that $\frac{-u(x,p)}{1-p} = \frac{u(x,1)-u(x,p)}{1-p}$ is strictly decreasing in p. Thus,

$$r > -\frac{u(x, p_0)}{1 - p_0}, \tag{26}$$

or equivalently,

$$r + u(x, p_0) \quad < \quad rp_0. \tag{27}$$

Thus, the type also prefers to avoid information when the prior is p_0. This finishes the proof of the second item.

D.2 Equilibrium selection

In the experimental setup, we are careful to limit the information cost of the dictators to a minimum. However, one may think that still there can be minimal cost and that these drive equilibrium selection. In the following, we show that the equilibrium of Lemma 2 in which the agent stops whenever she is indifferent between stopping and continuing, is the unique equilibrium that is stable with respect to the introduction of minimal cost.

Formally, we consider a variation of the model in Section 1.1 with one modification. For every instant in time where the agent observes the information process $(Z_t)_{t\geq 0}$, she pays a bounded and positive flow cost $c(-)$ which may depend on her posterior belief, that is, if the agent observes the process until time $t \geq 0$, the total cost she pays is

$$\int_0^c c(p_s)ds. \tag{28}$$

The flow cost allow for diverse economic interpretations. For example, constant flow cost c are linear in time spent and may be interpreted as time cost.[47]

The variation of the model with flow cost appears in Morris and Strack (2019). We will use the results in Morris and Strack (2019) to analyze stability of equilibria with respect to the introduction of minimal cost. First, we summarize the relevant results of their work: Proposition 1 shows that for any Bayes-consistent distribution G, there is a stopping strategy so that the distribution F_τ of stopped posterior beliefs is equal to G. For any distribution G of posterior beliefs, the minimal cost of generating that distribution are denoted

$$C(G) = \inf_{\tau:F_\tau=G}\mathrm{E}(\int_0^c c(p_\tau)ds), \tag{29}$$

Given that the sum of material and belief utility $V(p)$ of the agent only depends on the posterior belief, in equilibrium, the agent chooses a cost-minimal stopping strategy. Proposition 2 directly implies a characterization of the flow cost associated with any cost-minimal stopping continuation strategy τ, for any given current belief p_t, only in terms of the distribution G of the stopped posterior beliefs p_τ that it generates:

$$\mathrm{E}(\int_0^c c(p_\tau)ds|(Z_s)_{s\leq t}) = \mathrm{E}_G(\phi_c(p_\tau, p_t)) \tag{30}$$

for $\phi_c(q, p_t) = \int_{p_t}^q \int_{p_t}^x \frac{c(y)}{2[y(1-y)]^2} dydx$ (compare to equation (1) in their paper). These results together allow to write the agent's continuation utility as a function of G,

$$\mathrm{E}_G(W(p_\tau, p_t)) \tag{31}$$

for $W(p, p_t) = V(p) - \phi_c(p, p_t)$ for $V(-) = \arg\max_{a\in\{x,y\}} U(a, p)$ and U as defined in (1). Now, we will use this characterization of the continuation utility to study the stability of equilibria in the dynamic information acquisition game of Section 1.1.

Stability. We want to select equilibria of the game without cost for which there is a close-by equilibrium in a close-by environment with minimal cost. We say that a Nash equilibrium τ of the game without cost is *stable* if there is sequence $(c_n)_{n\in\mathbb{N}}$ of cost functions with $\lim_{n\to\infty} c_n = 0$ and a sequence of equilibria $(\tau_n)_{n\in\mathbb{N}}$ given $(c_n)_{n\in\mathbb{N}}$ so that $\lim_{n\to\infty} \tau_n = \tau$.[48]

[47] In our experiment, we impose only a minimum time lag of 0.3 second between draws of information. On average, the dictators in the experiment spend 57 seconds acquiring information.

[48] Our notion of convergence is that the distribution of the stopped posterior beliefs converges

The following theorem shows that the equilibrium of Lemma 2 is the unique stable equilibrium.

Theorem 3 *Take any sequence $(c_n)_{n\in\mathbb{N}}$ of cost functions with $\lim_{n\to\infty} c_n = 0$. Suppose there is a converging sequence of equilibria $(\tau_n)_{n\in\mathbb{N}}$ given $(c_n)_{n\in\mathbb{N}}$. Then, τ_n converges to the equilibrium in Lemma 2, in which the agent stops if and only if $p_t \geq p_h$ or $p_t \leq p_l$.*

Proof. Take the equilibrium sequence $(\tau_n)_{n\in\mathbb{N}}$. First, we show that the agent does not stop at a belief $p_t \in (p_l, p_h)$ when n is sufficiently large. Given (31), the utility from stopping at $p_t \in (p_l, p_h)$ is $W(p_t, p_t) = V(p_t)$. Taking the analogous version of Claim 2 for the setting with cost, there is a continuation strategy that gives utility $\overline{W}(p_t, p_t)$, where $\overline{W}(-, p_t)$ denotes the concave envelope of $W(-, p_t)$, analogous to the definition of $\overline{V}$. Since cost converge to 0 as $n \to \infty$, $\overline{W}(p_t, p_t)$ converges to $\overline{V}(p_t)$ as $n \to \infty$ and $\overline{V}(p_t) > V(p_t)$ since $p_t \in (p_l, p_h)$ by the definition of p_l and p_h (see Section 1.2). This implies that when n is sufficiently large, it is optimal to continue at p_t.

Second, we compare the utility of τ_n with the utility from the strategy τ^* in which the agent stops if and only if $p_t \geq p_h$ or $p_l \leq p_t$. Given (31), the utility from τ_n converges to

$$\lim_{n\to\infty} \mathrm{E}(V(p_{\tau_n})) - \mathrm{E}(\phi_c(p_{\tau_n}, p_0)). \tag{32}$$

with

$$\lim_{n\to\infty} \mathrm{E}(V(p_{\tau_n})) \quad \leq \quad \lim_{n\to\infty} \mathrm{E}(\bar{V}(p_{\tau_n})) \leq \lim_{n\to\infty} \bar{V}(\mathrm{E}(p_{\tau_n})) = \bar{V}(p_0), \tag{33}$$

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.[49] Given Claim 2 and (31), the strategy τ^* yields utility

$$\bar{V}(p_0) - \mathrm{E}(\phi_c(p_{\tau^*}, p_0)). \tag{34}$$

Our first observation in this proof implies that the distribution of stopped beliefs given $\tau = \lim_{n\to\infty} \tau_n$ has support in $(p_l, p_h)^c$. Suppose that it would have positive mass in $[0, p_l) \cup (p_h, 1]$. This would imply that $\mathrm{E}(\phi_c(p_{\tau_n}, p_0)) > \mathrm{E}(\phi_c(p_{\tau^*}, p_0)$ when

in distribution.

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

n is sufficiently large. Together with (32) - (34), this implies that the utility from τ_n is strictly lower than that of τ when n is sufficiently large. This contradicts the assumption that τ_n is a Nash equilibrium. We conclude that the support of the distribution F_{τ_n} of stopped posterior beliefs converges to the set $\{p_l, p_h\}$ as $n \to \infty$. ■

E Original instructions in German

In this online appendix, we include the original instructions that we used in the experiment. The original instructions are paper-based and in German language. They are similar in all four treatment: *NoForce Tradeoff*, *Force-Tradeoff*, *Force-Control* and *NoForce-Control*. We include the instructions in *NoForce-Tradeoff* in full and point out the deviation from them in the three other treatments respectively. We include the English translation of these instructions in Online Appendix F.

E.1 Treatment: NoForce-Tradeoff

Allgemeine Erklärungen

Wir begrüßen Sie zu dieser Studie! Im Rahmen dieser Studie können Sie eine nicht unerhebliche Summe Geld verdienen. Lesen Sie die folgenden Erklärungen daher bitte gründlich durch! Wenn Sie Fragen haben, strecken Sie bitte Ihre Hand aus der Kabine – wir kommen dann zu Ihrem Platz.

Während der Studie ist es nicht erlaubt, mit den anderen Studienteilnehmern zu sprechen, Mobiltelefone zu benutzen oder andere Programme auf dem Computer zu starten. Die Nichtbeachtung dieser Regeln führt zum Ausschluss aus der Studie und von allen Zahlungen. Ihr Einkommen aus dieser Studie bekommen Sie am Ende der Studie bar ausbezahlt. Während der Studie sprechen wir nicht von Euro, sondern von Punkten. Ihre gesamte Auszahlung wird also zunächst in Punkten berechnet und dann am Ende in Euro umgerechnet, wobei gilt:

1 Punkt = 5 Cent.

Teilnehmerzuordnung: Durch eine Zufallsentscheidung hat Ihnen der Computer aus allen Studienteilnehmern in diesem Raum einen anderen Teilnehmer zugeordnet. Im folgenden bezeichnen wir den Ihnen zugeordneten Studienteilnehmer als 'den anderen Teilnehmer'.

Benutzen Sie diesen Teilnehmerbogen gerne als Referenz während der Bearbeitung am Computer. Im Vorlauf zur Bearbeitung der Studie werden wir Sie bitten, einige Kontrollfragen zu bearbeiten.

Wahrscheinlichkeiten

In diesem Abschnitt möchten wir Sie ein wenig mit mathematischen Wahrscheinlichkeiten vertraut machen.

Stellen Sie sich folgende Situation vor: In einem Raum sitzen 10 Teilnehmer. Jeder Teilnehmer hat eine Box. Die Teilnehmer können keine der Boxen sehen, wissen aber folgendes:

- In jeder Box befinden sich 10 Bälle.
- 5 der 10 Teilnehmer haben eine Box mit 6 weißen Bällen und 4 schwarzen Bälle (Situation A).
- 5 der 10 Teilnehmer haben eine Box mit 4 weiße Bällen und 6 schwarze Bällen (Situation B).

Stellen Sie sich nun folgendes vor: Sie sind einer der Teilnehmer. Die Wahrscheinlichkeit, dass Sie sich in Situation A befinden, ist also 50 %. Ein Computer zieht einen Ball aus Ihrer Box und legt ihn nach dem Ziehen wieder zurück. Wenn Sie die Farbe des gezogenen Balles erfahren, hift dies, besser einzuschätzen in welcher Situation Sie sich befinden.

Frage: Was ist die Wahrscheinlichkeit, dass Sie sich in Situation A befinden, gegeben dass der gezogene Ball weiß ist?

Wahrscheinlichkeiten, welche zusätzliche Information - wie die Farbe des gezogenen Balles - berücksichtigen, werden auf statistisch korrekte Weise nach einem mathematischen Gesetz, dem Satz von Bayes' berechnet. Die Berechnung ist kompliziert, und benötigt mehrere Rechenschritte. **Daher zeigen wir Ihnen die statistisch korrekten Wahrscheinlichkeiten an, wann immer sie entscheidungsrelevant sind.**

Beispiel (Antwort zur Frage):

$$
\begin{aligned}
& \text{Wahrscheinlichkeit von Situation A, gegeben, dass der gezogene Ball weiß ist} \\
= \;& \text{W'keit von Situation A} \cdot \left[\frac{\text{W'keit, dass der gezogene Ball in Situation A weiß ist}}{\text{W'keit, dass der gezogene Ball weiß ist}}\right] \\
= \;& 50\% \cdot \frac{60\%}{50\%} \\
= \;& 60\%.
\end{aligned}
$$

Die Wahrscheinlichkeit von Situation A, gegeben, dass der gezogene Ball schwarz ist, wird auf ähnliche Weise berechnet. **Bitten bearbeiten Sie nun die erste Kontrollfrage am Computer.**

Ihre Entscheidungen

Sie erhalten nun 100 Punkte auf ihr Punktekonto. Als Nächstes fällen Sie eine Entscheidung, die Ihre eigene Auszahlungshöhe und die Auszahlungshöhe des anderen Teilnehmers beeinflusst, welcher ebenfalls 100 Punkte auf sein Punktekonto erhalten hat. Die Entscheidungen des anderen Teilnehmers haben jedoch keine Auswirkung auf Ihre Auszahlungshöhe.

Ihre Entscheidung besteht daraus, zwischen zwei Optionen, X und Y, zu wählen.

1. Eine dieser Optionen ist für den anderen Studienteilnehmer 'schädlich' und führt dazu, dass er 80 Punkte weniger ausgezahlt bekommt.

2. Die andere der Optionen hat keine Auswirkung auf den anderen Studienteilnehmer, diese Option ist 'sicher'.

3. Für je 7 von 20 Entscheidern ist die Option X sicher und Option Y schädlich (35 % Wahrscheinlichkeit), und für je 13 von 20 Entscheidern ist die Option Y sicher und Option X schädlich (65 % Wahrscheinlichkeit). Es wurde vom

Computer bereits zufällig ausgewählt, welche Option in Ihrem Fall sicher und welche schädlich ist.

4. Unabhängig davon, ob Option X oder Option Y für den anderen Teilnehmer sicher ist, bekommen Sie selber 25 Punkte mehr ausgezahlt, wenn Sie sich für Option X entscheiden.

Ihre Information

Bevor Sie sich entscheiden, können Sie zusätzliche Information darüber erhalten, welche der Optionen der Computer für Sie als sicher ausgewählt hat. Wenn Sie die Entscheidung ohne zusätzliche Information treffen möchten, drücken Sie bitte direkt auf ‘Entscheidung ohne zusätzliche Information’. Wenn Sie die Entscheidung mit zusätzlicher Information treffen möchten, drücken Sie bitte auf ‘Entscheidung mit zusätzlicher Information’

In einer Box befinden sich 100 weiße oder schwarze Bälle. Falls Option X sicher ist, befinden sich in der Box 60 weiße Bälle, und 40 schwarze Bälle. Falls Option Y sicher ist, befinden sich in der Box 40 weiße Bälle, und 60 schwarze Bälle.

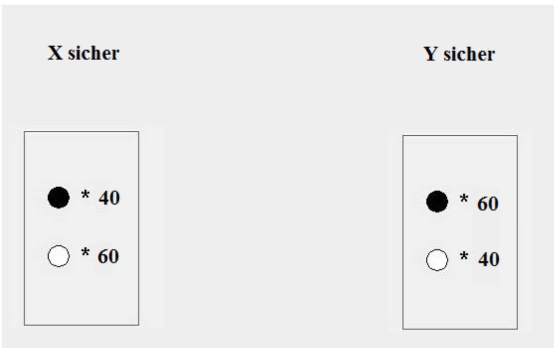

Figure 10: Box mit 40 schwarzen und 60 weißen Bällen (X sicher); Box mit 60 schwarzen und 40 weißen Bällen (Y sicher)

Sie werden die Box nicht gezeigt bekommen, aber sie können den Computer einen Ball zufällig aus der Box ziehen lassen. Dafür klicken Sie auf die Schaltfläche ‘Ein weiterer Ball’. Nach Ihrem Klick wird der Ball, den der Computer aus der Box gezogen hat, auf dem Bildschirm eingeblendet. Danach legt der Computer den Ball in die Box zurück, sodass sich wieder 100 Bälle in der Box befinden.

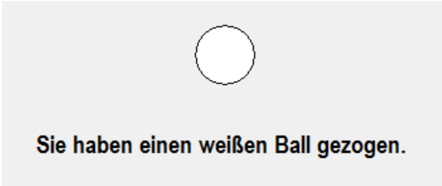

Figure 11: Beispiel: Der vom Computer gezogene Ball ist weiß.

Sie können danach einen weiteren Ball ziehen lassen. Dafür klicken Sie wiederum auf die Schaltfläche 'Ein weiterer Ball'. Sie können unbegrenzt viele Bälle ziehen lassen. Wenn Sie keine Bälle mehr ziehen lassen möchten, klicken Sie auf die Schaltfläche 'Zur Entscheidung'.

Erinnern Sie sich, dass mehr weiße Bälle in der Box sind, wenn Option X sicher ist (60 weiße Bälle), als wenn Option Y sicher ist (40 weiße Bälle)? Wenn Sie also sehen, dass ein weißer Ball gezogen wurde, ist dies ein Hinweis darauf, dass Option X sicher ist. Nach jedem Ball wird die statistisch korrekte Wahrscheinlichkeit, dass Option X sicher ist, eingeblendet. Dabei werden alle Bälle, die Sie bereits gezogen haben, berücksichtigt.[50]

Die betroffene Person weiß nicht, ob oder wieviel Sie sich informiert haben. Sie erfahren am Ende der Studie nicht, ob X oder Y die schädliche Aktion ist. Es ist kostenlos für Sie, zusätzliche Information vor Ihrer Entscheidung zwischen X und Y zu erhalten.

E.2 Treatment: NoForce-Control

...

[50] Für die genaue Berechnung der Wahrscheinlichkeit, dass Option X sicher ist, wird ein mathematisches Gesetz, der Satz von Bayes, benutzt. Die exakte Berechnungsformel, wenn z.B. ein einziger weißer Ball gezogen wurde, ist folgende:

$$\begin{aligned} & \text{Wahrscheinlichkeit, dass Option X sicher ist, gegeben, dass der gezogene Ball weiß ist} \\ = \; & \text{W'keit, dass Option X sicher ist} \\ & \cdot \frac{\text{W'keit, dass der gezogene Ball weiß ist, wenn Option X sicher ist}}{\text{W'keit, dass der gezogene Ball weiß ist}}. \end{aligned}$$

Dies ist die einzige objektiv richtige Berechnungsweise.

Ihre Entscheidungen

...

Ihre Entscheidung besteht daraus, zwischen zwei Optionen, X und Y, zu wählen.

...

4. Ihre eigene Auszahlungshöhe ist unabhängig davon, ob Sie sich für X oder Y entscheiden. Ihre Auszahlungshöhe ist auch unabhängig davon, ob Option X oder Option Y für den anderen Teilnehmer sicher ist.

...

E.3 Treatment: Force-Tradeoff

...

Ihre Information

Bevor Sie sich entscheiden, können Sie zusätzliche Information darüber erhalten, welche der Optionen der Computer für Sie als sicher ausgewählt hat.

...

E.4 Treatment: Force-Control

...

Ihre Entscheidungen

...

Ihre Entscheidung besteht daraus, zwischen zwei Optionen, X und Y, zu wählen.

...

4. Ihre eigene Auszahlungshöhe ist unabhängig davon, ob Sie sich für X oder Y entscheiden. Ihre Auszahlungshöhe ist auch unabhängig davon, ob Option X oder Option Y für den anderen Teilnehmer sicher ist.

...

Ihre Information

Bevor Sie sich entscheiden, können Sie zusätzliche Information darüber erhalten, welche der Optionen der Computer für Sie als sicher ausgewählt hat.

...

F Instructions English translation

F.1 Treatment: NoForce-Tradeoff

General Explanations

Welcome to the study! In this study, you can earn a good amount of money. Please carefully read the following explanations! Shall you had questions, please stick your hand out of the cubicle—we will come to your seat.

During the study, it is not allowed to talk with other participant, to use mobile phones, nor to start other programs on the computer. The violation of these rules will lead to an exclusion form the study and any payment. You will receive your payment of the study at the end of the study in cash. During the study, we do not talk about Euro. Instead we will talk about points. Your total payment will be calculated in points and translated into Euro at the following rate:

1 Point = 5 Cent.

Participant Pairing: The computer has paired you with another participant who is randomly selected from all the participants in the room. In the following, we refer to the participant whom you are paired with as 'the other participant'.

Please feel free to refer back to this Instruction when you are working on the computer. Before the study starts, we will ask you to answer a couple of contrl questions.

Probabilities

In this section, we want to familiarize you with mathematical probabilities.

Imagine the following situation: in a room there are 10 participants. Every participant has a box. The participants cannot see the boxes but know the following:

- In each box there are 10 balls.
- 5 of the 10 participants have a box with 10 white balls and 4 black balls (Situation A).
- 5 of the 10 participants have a box with 4 white balls and 6 black balls (Situation B).

Now imagine the following: you are one of the participants. So the probability that you are in Situation A is 50%. A computer draws a ball out of your box and places it back into the box after the draw. When you find out the color of the drawn ball, it helps to better assess what situation you are in.

Question: What is the probability that you are in Situation A, given that the drawn ball is white?

Probabilities, which take additional information into account—such as the color of the drawn ball—are calculated in a statistically correct manner according to a mathematical law, Bayes' theorem. The calculation is complicated and requires several calculation steps. **We therefore show you the statistically correct probabilities whenever they are relevant for your decision.**

Example (Answer to the Question):

$$
\begin{aligned}
& \text{Probability of Situation A, given that the drawn ball is white} \\
= \; & \text{Probability of Situation A} \cdot \left[\frac{\text{Prob that the drawn ball in situation A is white}}{\text{Probability that the drawn ball is white}}\right] \\
= \; & 50\% \cdot \frac{60\%}{50\%} \\
= \; & 60\%.
\end{aligned}
$$

The probability of Situation A, given that the drawn ball is black, is calculated in a similar way. **Now, please process to the first control question on the computer.**

Your Decisions

You obtain now 100 points to your points account.

Next you make a decision that affects your own payout amount and the payout amount of the other participant, who also received 100 points on the points account. However, the decisions of the other participant have no impact on your payout amount.

Your decision is to choose between two options, X and Y.

1. One of these options is 'harmful' to the other study participant and leads to 80 points less being paid out to the participant.

2. The other of the options has no effect on the other study participant, this option is 'safe'.

3. For every 7 out of 20 decision-makers, option X is safe and option Y is harmful (35 % probability), and for every 13 out of 20 decision-makers, option Y is safe and option X is harmful (65 % probability). The computer has already chosen at random which option is safe and which is harmful in your case.

4. Regardless of whether option X or option Y is safe for the other participant, you will receive 25 points more yourself if you choose option X.

Your Information

Before you make a decision, you can get additional information about which of the options the computer has selected to be safe for you. If you want to make the decision without additional information, please click directly on ‘Decision without additional information’. If you want to make the decision with additional information, please click on ‘Decision with additional information ’.

There are 100 white or black balls in a box. If option X is safe, there are 60 white balls and 40 black balls in the box. If option Y is safe, there are 40 white balls and 60 black balls in the box.

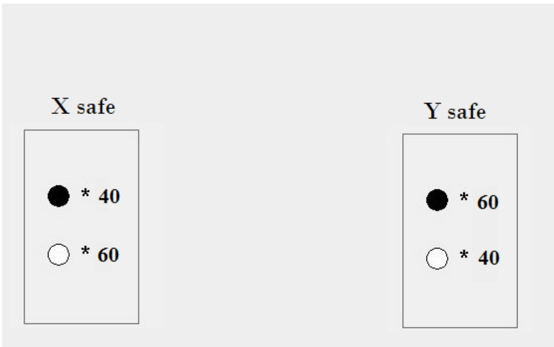

Figure 12: Box with 40 black and 60 white balls (X safe); Box with 60 black and 40 white balls (Y safe)

You won’t be shown the box, but you can have the computer draw a ball out of the box at random. To do this, click the button ’Another Ball’. After you click, the ball that the computer drew out of the box will appear on the screen. Then the computer puts the ball back in the box so that there are 100 balls in the box again.

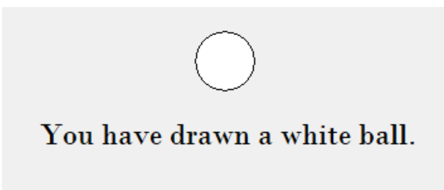

Figure 13: Example: The ball drawn by the computer is white.

You can then let another ball be drawn. To do this, click the ’Another Ball’ button again. You can have an unlimited number of balls drawn. If you do not want to have any more balls drawn, click the button 'To the decision' .

Do you remember that there are more white balls in the box when option X is safe (60 white balls) than when option Y is safe (40 white balls)? So if you see that a white ball has been drawn, it is an indication that option X is safe. After each ball, the statistically correct probability that option X is safe is displayed. All balls that you have already drawn are taken into account.[51]

The person affected by your choice does not know whether or how much you have obtained information. You won't find out at the end of the study whether X or Y is the harmful action. It is free of charge for you to obtain additional information before making your decision between X and Y.

F.2 Treatment: NoForce-Control

Your Decisions

...

Your decision is to choose between two options, X and Y.

...

4. Your own payout amount is independent of whether you choose X or Y. Your payout amount is also independent of whether option X or option Y is safe for the other participant.

[51] A mathematical law, Bayes' theorem, is used to accurately calculate the probability that option X is safe. The exact calculation formula, if e.g. a single white ball has been drawn, is the following:

$$\begin{aligned} & \text{Prob. that option X is safe conditional on the ball drawn being white} \\ = \; & \text{Prob. that option X is safe} \\ & \cdot \frac{\text{Prob. that a white ball is drawn conditional on option X being safe}}{\text{Prob. that a white ball is drawn}}. \end{aligned}$$

This is the only objectively correct calculation method.

F.3 Treatment: Force-Tradeoff

...

Your Information

Before you make a decision, you can get additional information about which of the options the computer has selected to be safe for you.

...

F.4 Treatment: Force-Control

Your Decisions

...

Your decision is to choose between two options, X and Y.

...

4. Your own payout amount is independent of whether you choose X or Y. Your payout amount is also independent of whether option X or option Y is safe for the other participant.

...

Your Information

Before you make a decision, you can get additional information about which of the options the computer has selected to be safe for you.

...