Heads in the Sand: Theory and Experiment on the Interdependence of Information Avoidance

Silvia Sonderegger Fabio Tufano Marcelo Woo*

*Job Market Paper

This version: November 10, 2025

Click here for the most recent version

Abstract

We study whether the deliberate avoidance of information that is freely available and instrumentally useful is interdependent across people. Agents derive direct utility from anticipating their future, state-dependent consumption, potentially creating incentives to avoid information that induces such anticipation. Building on Bénabou (2013), our model predicts that when remaining ignorant imposes negative externalities on others' prospects, individuals' decisions to acquire or avoid information may exhibit strategic complementarities (making ignorance contagious), strategic substitutabilities (making it self-limiting), or be unconditional. We test these predictions in a novel experiment where participants choose whether to learn a payoff-relevant state and where remaining ignorant worsens a future payoff common to all group members. Consistent with contagious ignorance, exogenously increasing expectations about the prevalence of information avoidance raises participants' own propensity to avoid information, on average. The contagion of ignorance appears mediated by a deterioration in anticipatory utility upon learning bad news —which become even worse news when ignorance is more prevalent. Individuals vary substantially in their best-responses to others' informational choices: around 40% of subjects condition their informational choices on those of others, split evenly between strategic complements and strategic substitutes. The mixture of such strategy types affects equilibrium share of information avoiders monotonically and nonlinearly. Our findings suggest that optimal incentive schemes aiming to promote the aggregate take-up of information should account for interdependencies in informational decisions.

Silvia Sonderegger: University of Nottingham, CeDEx, silvia.sonderegger@nottingham.ac.uk. Fabio Tufano: University of Leicester, CeDEx, fabio.tufano@leicester.ac.uk. Marcelo Woo: University of Nottingham, CeDEx, marcelo.woo@nottingham.ac.uk.

1 Introduction

Information avoidance —the deliberate choice to avoid readily available and instrumentally valuable information— is well documented in individual decision-making (Thornton, 2008; Ho et al., 2021).[1] For instance, people often choose not to learn whether they carry a potentially fatal disease, despite the availability of accurate tests (Oster et al., 2013) and investors have been found to behave like "ostriches", avoiding financial news (Galai & Sade, 2006; Karlsson et al., 2009; Fong & Hunter, 2022; Olafsson & Pagel, 2025), even though learning this information would enable them to make better decisions and achieve better outcomes. In many situations, however, individual outcomes are not independent; rather, they are interlinked with those of others, creating the possibility that avoidance of decision-relevant information by one agent imposes negative externalities on other agents' material payoffs. Examples include global warming, infectious diseases, or a company's success (or failure). Consider, for instance, a bank investing in assets that may be either of high quality or junk. Each employee understands that the bank's future depends not only on their own decisions but also on the choices of others —whether his peers stick to the current strategy (optimal if the assets are of high quality) or liquidate the assets (optimal if they are junk). Making the right decision, however, requires assessing the quality of the assets and potentially acknowledging that the bank may be on the brink of bankruptcy (and that all employees might be facing imminent unemployment). Importantly, if the assets turn out to be junk, the severity of the situation depends on the choices of others. If others continue as usual rather than taking action, the outcome will be much worse. Therefore, each employee's decision to bury his head in the sand or not impacts not only his own payoff but also everyone else's.

This paper studies information avoidance in a group setting where remaining ignorant imposes a negative payoff externality on other members. People choose to remain ignorant or acquire information (and act on it) in an environment where, as in the example above, their decision impacts both their own payoff and the payoffs of others. When more individuals remain ignorant, the outcome in the bad state of the world becomes worse for everyone. In this context, a key question arises that set it apart from information acquisition in individual decision-making. Are there interdependencies in information avoidance? In other words, do individuals base their decisions to acquire or avoid information on the choices of others in their group, and if so, how?

These questions are crucial to understand whether and how information avoidance self-sustains

[1] For a comprehensive review on information avoidance, see Golman et al. (2017). Hertwig & Engel (2016) define the related concept of *deliberate ignorance* as "the conscious individual or collective choice not to seek or use information", specifically "in situations where the marginal acquisition costs are negligible and the potential benefits potentially large, such that —from the perspective of the economics of information (Stigler, 1961)— acquiring information would seem to be rational (Martinelli, 2006)".

within organizations and groups. In a seminal theoretical paper, Bénabou (2013) describes what he refers to as contagious ignorance —a situation in which individuals become more likely to avoid information if others are doing the same. Intuitively, if everyone else chooses to bury their head in the sand, the severity of the bad state will be so overwhelming that individuals may prefer to stay unaware rather than confront such a potentially disastrous reality. History is filled with examples of collective information avoidance that appear consistent with the notion of contagious ignorance, from Easter Island to Lehman Brothers. However, the existing empirical literature lacks a systematic study of whether and how information avoidance self-reinforces in groups. This work aims to fill this gap.

This paper answers these questions by combining theory, experiment, and simulations. Our investigation begins with a theoretical model inspired by Bénabou (2013), which studies the strategic interdependencies that may arise in group settings. As well as contagious ignorance, we show that in our setup there is also an alternative possibility. Depending on the nature of their utility function, individuals may exhibit strategic substitutability in information avoidance: when others bury their heads in the sand, these individuals become less inclined to shun information. These findings indicate that the dynamics of information avoidance in groups may vary considerably. Furthermore, the theory also shows that some individuals may not base their choices on the actions of others at all, and may instead always avoid or always acquire information, regardless of what others do.

Informed by this theory, we design a novel lab experiment to assess the interdependence of information avoidance empirically.[2] Participants are matched in groups and informed that there are two possible states of the world: *Good* and *Bad*. In the *Bad* state, everyone in the group will be exposed to an aversive future event: at the end of the experiment, they will hear a series of distressing and unpredictable screams, a consumption event known to produce anticipatory anxiety.[3] In the *Good* state, no noises will be heard. At the beginning of the experiment, each participant chooses whether to discover the state immediately ("Now" option) or to wait until the future event unfolds ("Later" option). Importantly, the volume of the screams in the *Bad* state decreases in the number of members who acquire information about the state (that is, who choose Now), reflecting the decision value of information. In other words, each participant understands that by selecting Now, they can mitigate the intensity of the noises not only for themselves but also for everyone else in their group, should the state turn out to be *Bad*.

To assess the interdependence, we aim to characterize how one's demand for information

[2] The experiment was preregistered, AsPredicted #153871 and #175588.

[3] This is shown in Beaurenaut et al. (2020). We further discuss the rationales for the use of screams as a consumption event in Section 4.2.

depends on others' information decisions. For this, we employ two empirical strategies. First, using an information provision experiment, we exogenously vary subjects' expectations about the prevalence of information avoiders (people who choose Later) in their group by informing them about an analogous share observed in a pilot study. We randomize two treatments at the level of experimental groups: in the *Many Avoided* condition, subjects are told that 80% of subjects in a prior group avoided information, whereas in the *Few Avoided* condition, subjects are told only 20% did. The resulting variation in posterior beliefs allows us to test whether, on average, a higher expected prevalence of avoidance among peers increases the likelihood of own's avoidance. Second, using the strategy method (Selten, 1967; Fischbacher et al., 2001), we elicit each subject's best-response *schedule* —his conditional response to a range of shares of avoiders within the group. This allows us to identify individual heterogeneity in interdependence beyond the average effect.

We find evidence that, on average, information avoidance is interdependent, and in particular, contagious. Our treatment intervention strongly shifts participants' expectations about their peers' information decisions: beliefs about the share of avoiders increase by 22 percentage points on average, from 37% in *Few Avoided* to 59% in *Many Avoided*. The treatment also affects behavior: consistent with contagious ignorance, the share of subjects avoiding information increases by one-third, from 24% to 33%. 2SLS regressions further confirm that the average relationship between beliefs and behavior (among subjects who respond to the treatment) is positive: for every 10 percentage points of increase in beliefs about the prevalence of avoidance among others, the likelihood of avoidance increases by about 5 percentage points. Thus, on average, information avoidance exhibits strategic complementarity, in line with the prediction of Bénabou (2013).

We provide supportive evidence that a key mechanism driving the interdependence of information choices is a deterioration of anticipatory utility upon finding bad news in response to exogenous changes in expectations about the prevalence of avoidance. To show this, we develop and validate a novel and portable elicitation of anticipatory utilities, as expected by the subjects at the moment of the choice between acquiring or avoiding information.[4] Its key characteristic is its granularity: instead of simply eliciting the anticipatory utility upon acquiring information, we further condition on the events of finding out good news and bad news, separately. In line with the theoretical channel, subjects in the *Many Avoided* condition (relative to those in *Few Avoided*) report worse anticipatory utility when imagining the event of acquiring information and discovering *bad* news about the state, but we find no differences across treatments in the anticipatory utility upon discovering *good* news. This pattern lends support for the focal mechanism of this paper

[4]Note that it is the expectation at that time, and not the actual anticipatory utilities, what matter for the information decisions.

and, in contrast, is hard to square with more traditional channels such as conformity, herding, and social norms, which predict that belief about others' behavior —shifted by the treatment— affects the utility associated to acquiring information, without distinguishing between the events of discovering bad versus good news.

Importantly, beyond the average complementarity, individuals differ substantially in the type of interdependence they exhibit. Consistent with theoretical predictions, we find substantial heterogeneity in how subjects best-respond to varying shares of information avoiders in their group. The key finding is that we document strategic interdependencies in information decisions at the individual level. Among those influenced by others' choices (about 40% of the sample), half exhibit a pattern consistent with contagious ignorance (strategic complementarity), as in Bénabou (2013), while an equally sizeable share display the opposite tendency (strategic substitutability), imposing limits to the spread of ignorance. There is also a substantial share of participants whose information choice is independent of their teammates' decisions. The larger portion consists of those who choose to unconditionally acquire information (approximately 42%), while just under 9% of participants consistently choose to avoid information. Therefore, we document that, empirically, individuals exhibit substantial variation in their best-response types.

The diversity of such types of best-response schedules suggests that their composition within a group matters for the equilibrium prevalence of avoidance. Motivated by this intuition, we employ simulations to show how the interaction of strategy types determines the equilibrium prevalence of information avoidance in a group. Groups primarily composed of individuals who exhibit strategic complementarity in information decisions generate multiple equilibria, and risk getting trapped in a so-called *Mutually Assured Ignorance (MAI)* equilibrium (see Section 3), where everybody prefers to remain ignorant, at the cost of worsening the future. The inclusion of unconditional avoiders precipitates this outcome: with 30% unconditional avoiders, the MAI equilibrium arises with 75% probability, and with 50% it emerges with certainty. Likewise, unconditional information getters can drive the group toward a virtuous equilibrium where everybody acquires information (*Mutually Assured Awareness* equilibrium). In contrast, groups composed of individuals who exhibit strategic substitutability never reach a MAI equilibrium but are also unable to sustain an equilibrium characterized by universal information acquisition.

Furthermore, the relationship between group composition and the equilibrium prevalence of information avoidance is monotonic and nonlinear. For example, when the participation of Always Avoiders increases from 0% to 10% in a group conformed otherwise by Strategic Complements, the median equilibrium share of information avoiders increases from 0% to 22%. However, when the participation increases by a similar amount from 10% to 20%, the median jumps sharply from 22%

to 100%. Similar nonlinear dynamics are observed for other quantiles and for the average of the distribution of equilibria.

The remainder of the paper is organized as follows. In Section 2, we discuss related literature. In Section 3, we lay out our theoretical framework. In Section 4, we present the experimental design. In Section 5, we report the empirical findings. In Section 6, we use our data to inform the relationship between group compositions and equilibrium levels of avoidance. Finally, in Section 7, we conclude.

2 Related Literature

Our study relates mainly to two strands of literature: attitudes towards information and suboptimal decision-making in groups.

First, we contribute to the literature on attitudes towards information from both theoretical and empirical perspectives. On the theoretical side, standard economic theory recognizes that information is a valuable resource for decision-making and accordingly predicts that, if available at no cost, agents will typically acquire information about payoff-relevant states of the world to improve their subsequent decisions. However, a growing literature documents and acknowledges the possibility that economic agents may have intrinsic preferences for information beyond and above the instrumental, decision-value of information. Starting from Kreps & Porteus (1978), the theoretical literature has generated a variety of models in which agents may prefer to remain uninformed, or delay the acquisition of information, despite information being available at no or low cost. Golman et al. (2017) provides a review of models of information avoidance, defined as the deliberate choice to avoid freely available information when the agent is aware of its existence. Bénabou (2013), the model closest to our work, extends the analysis of information avoidance decisions to group settings. It shows theoretically that when willful blindness to information makes other agents worse off, it can become contagious. This rises the possibility of multiple equilibria, including a pervasive one in which all agents remain ignorant. We contribute to this literature by extending Bénabou (2013) to show that, under a broader class of preferences, strategic substitutability can also arise in that setting. In addition, we provide an empirical test of the prediction of interdependence.

On the empirical side, several studies both in the lab and in the field documented that individuals often prefer to remain ignorant despite information being available at no or low cost and being useful for subsequent decision-making (e.g., Dana et al., 2007; Karlsson et al., 2009; Oster et al., 2013; Ganguly & Tasoff, 2017; Pagel, 2018; Ho et al., 2021; Meissner & Pfeiffer, 2022; Masatlioglu et al., 2023;

Momsen & Ohndorf, 2023; Engelmann et al., 2024; Falk & Zimmermann, 2024). Across studies and contexts, rates of avoidance range from 5% (Ganguly & Tasoff, 2017) to more than 90% (Oster et al., 2013). The vast majority of empirical studies have focused on information avoidance in individual decision-making settings. Most closely related to our paper, Falk & Zimmermann (2024) studies intrinsic preferences for information about a potentially adverse future event in an individual setting. In their experiment, information confers no instrumental value: the decision-maker's future payoff is independent of whether information is acquired. We build on their experimental design to study information avoidance decisions in groups. Departing from them, we introduce a positive private return and externalities to information acquisition: becoming informed benefits the decision-maker and other members in her group. Therefore, we contribute to this literature in several ways. First, we document that, even in cases where information avoidance makes others worse off (negative payoff externalities), a substantial proportion of subjects still prefers to remain ignorant. Second, we provide, to the best of our knowledge, the first test of the interdependence of information decisions. Third, we document substantial heterogeneity of strategy types, and study its implications for the equilibrium levels of avoidance in groups. Fourth, we document how social preferences correlate with information decisions, which affect others' wellbeing.

Second, our study also relates to the literature on non-Bayesian information processing, overconfidence, and suboptimal decision-making in groups. A recent literature suggests that decisions in groups can become correlated through the social transmission of biases. Suvorov et al. (2024) shows that members of a group share noisy feedback about their own ability to other group members in a selective and asymmetric way: group members are more likely to share positive feedback about their ability to perform in an investment task than negative feedback. This selective information sharing leads to groups becoming overconfident about the group's ability to perform in investment decisions, and this overconfidence leads to suboptimal investment decisions by low ability groups. A closely related study by Oprea & Yuksel (2022) shows that when individuals can socially exchange ego-relevant beliefs, their initial biases amplify because subjects respond to their counterparts' beliefs in an asymmetric way. Such studies have focused on ex-post processing of information, after its content has been revealed to the decision-makers. Complementing these studies, we focus on ex-ante avoidance of information, and show that groups can also attain suboptimal (material) outcomes by deliberately avoiding learning from freely available signals, when signals reveal the state perfectly, and without any communication among group members.

3 Theoretical Framework

We now describe our theoretical framework and key results. All proofs are in the Appendix. Consider the following setup, inspired by (Bénabou, 2013). There are three periods, date 0, 1 and 2 and two states of the world, B (bad), occurring with probability $p \in (0, 1)$ and G (good), occurring with probability $1 - p$. Material payoffs are received at date 2 and generate utility u_i. At date 1, agents evaluate lotteries over date-2 outcomes according to the expected utility function $U_i = E[u_i]$. At date 0, agents evaluate lotteries over date-1 utilities U_i according to the expected utility function $E[v_i(U)]$, where $v_i(.)$ is a strictly increasing function capturing i's preferences over expected utility lotteries in the spirit of Kreps-Porteus (1978). We may interpret $v_i(E[u_i])$ as an anticipatory utility experienced in date-1 over date-2 payoffs. For example, if $v_i(.)$ is concave and the agent has interior priors $p_i \in (0, 1)$, then the loss in anticipatory utility upon learning that the state is Bad with certainty overweighs the gains in elation upon learning that the state is Good with certainty, creating an incentive to avoid learning the state. Denote with g the utility obtained at date 2 when the state of the world is good and with b when the state of the world is bad. b depends on whether an individual chooses to acquire or to avoid information about the state of the world at date 0. Intuitively, individuals who are aware that the state is bad can take corrective measures to reduce its negative consequences, whereas those who remain unaware cannot.[5] Thus, the utility at date 2 when the state is bad is equal to

$$b \equiv \bar{b} - \lambda_{-i} - \Delta \text{ if } i \text{ avoids information}$$

and

$$b + \Delta = \bar{b} - \lambda_{-i} > b \text{ if } i \text{ acquires information}$$

where $\bar{b} < g$ is the payoff when the state of the world is bad and all group members acquire information (and undertake corrective measures as a result), $\lambda_{-i} \in [0, \frac{n-1}{n}]$ is the share of group members other than i who avoid information, n is the number of individuals in the group and $\Delta \equiv \frac{1}{n}$. In words, Δ captures the instrumental value of acquiring information. It reflects how much an agent can lessen the adverse effects of a bad state by obtaining information and taking appropriate corrective measures. Note that, when the state of the world is bad, utility is decreasing in λ_{-i}, the share of i's fellow group members who avoid information. By avoiding information, agents not only damage themselves but also exert a negative externality on others.

[5]This is imposed exogenously for simplicity but could arise endogenously in a setup where taking corrective action is costly, so that under ignorance agents would choose not to exert any corrective measure.

In what follows, we assume that

$$(1-p)(g-\bar{b}) > \Delta. \quad (1)$$

This assumption is not necessary for our results but simplifies the exposition. In words, it ensures that the expected value of a lottery where i obtains $\bar{b}$ with probability p and g with probability $1-p$ exceeds the value of obtaining $\bar{b}+\Delta$ for sure.[6]

An individual's date 0 *net* expected utility from avoiding information is

$$\varphi_i \equiv v_i(pb+(1-p)g) - [pv_i(b+\Delta) + (1-p)v_i(g)] \quad (2)$$

where $v_i(pb+(1-p)g)$ is the date-0 expected anticipatory utility from avoiding and $pv_i(b+\Delta) + (1-p)v_i(g)$ is the expected anticipatory utility from acquiring information. If expression (2) is positive, individual i strictly prefers to avoid information at date 0, while if expression (2) is negative, individual i strictly prefers to acquire it.

3.1 Types of Best Response ("Strategy Types")

We now explore how the nature of v_i determines an individual's optimal information acquisition choice as a function of others' choices (i.e., her best reaction function). The first two results consider v_i that is everywhere convex or everywhere concave.

Lemma 1 *If v_i is everywhere convex, individual i is an "Always Getter". At date 0, she acquires information independently of the information acquisition choices of other agents in her group.*

If v_i is convex, the utility boost obtained from discovering that the state is good outweighs the utility drop when the state is bad. Consequently, the agent always prefers to acquire information. Consider now the case where v_i is concave, as in Bénabou (2013).

Lemma 2 (Bénabou, 2013) *If v_i is everywhere concave, individual i (i) is an Always Getter (ii) is an "Always Avoider" or (iii) is affected by the choices of others in her group in the following way: there exists an interior threshold value λ_i^* such that i strictly prefers to avoid information if $\lambda_{-i} > \lambda_i^*$ and strictly prefers to acquire information if $\lambda_{-i} < \lambda_i^*$ (Strategic Complementarity).*

In the Appendix, we characterize the necessary and sufficient conditions for each of (i), (ii) or (iii) to apply. Always Getters and Always Avoiders make the same information acquisition choice independently of the choices of others in their group. Alternatively, the agent may condition

[6]Formally, $p\bar{b}+(1-p)g > \bar{b}+\Delta$.

her decision on the choices of others in her group, exhibiting strategic complementarity. The intuition for the result is as follows. Without the instrumental value of information, the agent would always prefer ignorance, since the utility loss from discovering the bad state outweighs the gain from finding out that the state is good (this follows directly from concavity). In the presence of instrumental value, as in our setup, the agent may opt to acquire information. However, as more agents choose to avoid information, the outcome in the bad state deteriorates, increasing the utility loss associated with discovering the bad state. As a result, meeting the conditions for information acquisition becomes increasingly challenging. This pattern may give rise to strategic complementarity in the agent's information acquisition strategy. The agent acquires information only if a sufficient share of her peers also do so; otherwise, she avoids information.

Consider now a reference-dependent v_i (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992). The agent's utility is largely insensitive to expected payoffs except around some reference level. Denoting i's reference expected utility as $\widehat{U_i}$, we have (i) $v_i''(U) < 0$ for $U > \widehat{U_i}$ and (ii) $v_i''(U) > 0$ for $U < \widehat{U_i}$.

Lemma 3 *If v_i is reference dependent, under some conditions the following applies: there exists an interior threshold value λ_i^{**} such that i strictly prefers to avoid information if $\lambda_{-i} < \lambda_i^{**}$ and strictly prefers to acquire information if $\lambda_{-i} > \lambda_i^{**}$ (Strategic Substitutability).*

In the Appendix, we provide sufficient conditions for strategic substitutability. To illustrate, assume that if the state is known to be bad utility is always below the reference level while it is above it in the good state. First, suppose that many agents in i's group avoid information. The final payoff in the bad state is so low that, at date 1, the expected utility under ignorance falls below the reference level. Relative to remaining ignorant, the utility loss from discovering the bad state is minimal (since both situations generate utility that is below the reference level), while the gain from discovering the good state is large. As a result, the agent prefers to acquire information. Now, suppose that many agents acquire information. The final payoff in the bad state is not too low and the expected utility under ignorance is above the reference level. In this case, under mild conditions, the agent prefers to avoid information, since the utility gain from discovering the good state is minimal relative to the utility loss from discovering the bad state.[7]

The following proposition summarizes our results.

Proposition 1 *Depending on the nature of v_i, i's propensity to avoid information may be (i) independent of, (ii) increasing (strategic complementarity) or (iii) decreasing (strategic substitutability) in the share of group*

[7]For example, suppose $b(\lambda_{-i} = 0) = -10$, $b(1) = -90$, $g = 50$, $\Delta = 5$, $p = 0.5$, $\widehat{U}_i = 0$, and $v_i(U) = U^{1/3}$ if $U > 0$ and $= -(-U)^{1/3}$ if $U < 0$. Then $\varphi^i(0) \approx 2.58$ and $\varphi^i(1) \approx -2.87$, so the agent exhibits strategic substitutability.

members who avoid information.

In what follows, we will use the term "strategy type" to indicate whether an individual is an Always Getter, an Always Avoider, or alternatively exhibits preferences featuring strategic complementarity (in short, is a "Complement") or substitutability (is a "Substitute").

Naturally, the specific information avoidance equilibria that may be reached vary depending on the agents' strategy types. Bénabou (2013), for example, examines the risks that emerge when agents are Complements. In that case, groups may become trapped in a "Mutually Assured Ignorance" (MAI) equilibrium where no one seeks out information, and as a result, no corrective actions are taken to shield the group from adverse outcomes in the face of a bad state of the world —potentially with disastrous consequences.[8] By the same token, a "Mutually Assured Awareness" (MAA) equilibrium, where everyone acquires information, is also possible. The path taken depends critically on the agents' expectations about the actions of others in their group. Conversely, a group of Substitutes cannot fall into either a MAI or MAA equilibrium, as its agents are inherently inclined to resist such dynamics.[9]

We examine the empirical existence of strategy types by designing an experiment that randomly varies beliefs about others' information choices, as well as by using the strategy method to elicit full best response schedules; results are discussed in Section 5. We also assess how the mixture of types impacts aggregate equilibrium avoidance by employing simulations in Section 6.

4 Experimental Design and Procedures

4.1 Experimental Design

We investigate whether, in groups, information avoidance is interdependent among people. Specifically, we ask: do individual decisions to remain uninformed depend on the expectation that others will do so?

4.1.1 Decision environment

To study the interdependence of avoidance, we design an experiment in which individuals in a group choose whether to acquire or avoid information. After being randomly assigned to groups, participants are told that, at the end of the experimental session, a state of the world (either good or

[8] Bénabou (2013) uses the acronym MAD (Mutually Assured Delusion) equilibrium. However, in the context of information avoidance MAI provides a more accurate description of the equilibrium. The same holds for the acronym MAA (Mutually Assured Awareness).

[9] Consider, for instance, a group composed of Substitutes with a 0.5 threshold. It is easy to show that, in the unique equilibrium, exactly half the group acquire information and half avoid it.

bad) will be randomly drawn. In the bad state (or the “Screams" state), the group will experience a negative consumption event: listening to a series of unpleasant screams. In the good state (or the “Quiet” state), no such event happens.

Each member of the group can choose one of two options, at no monetary cost. If she chooses option “Now”, she learns the state of the world immediately, while if she chooses option “Later” she defers learning the state of the world until the consumption event takes place at the end of the experimental session.[10] The choice between “Now” and “Later” is the key decision of interest. We refer to this binary choice as the “information decision”.[11]

Our experiment departs from other studies in that information decisions generate a payoff externality: the acquisition of information creates a positive return not only for oneself (standard instrumental value of information), but also for other members of the group (externality). If the state is *Screams*, the volume of the screams that all group members will hear decreases in the total number of members who acquire information (option Now). More precisely, each individual knows that by choosing Now, she lowers the volume of the screams in the Screams state by 2 points *for all* group members.[12]

In our decision environment, the acquisition of information is costless under standard accounts. First, there are no monetary costs associated with choosing either Now or Later. Second, cognitive costs of processing information are negligible, as upon acquisition subjects learn only a binary state of the world.

4.1.2 Treatments

We aim to identify whether information decisions are interdependent. To this end, we designed our experimental treatment to induce exogenous variation in participants’ beliefs about the share of information avoiders (i.e., people who would choose Later) in their group. We do so through an information provision treatment (Haaland et al., 2023), where we randomize information about the share of information avoiders in the past. There are two treatment conditions, labeled the *Many Avoided* and *Few Avoided*, randomized at the group level. In *Many Avoided* (*Few Avoided*), we inform that 80% (20%) of subjects in a particular group of a past pilot similar to the current experiment chose to avoid information about the state. Participants are also reminded of the potential impact

[10]Participants are aware that if they choose Later, they will be informed about the outcome just before the period in which screams are potentially played.

[11]In each group, decisions are made simultaneously and without communication. As a result, there is no scope for social learning among group members.

[12]As described in Section 4.2, the contribution to the volume of screams is additively separable. Therefore, we examine whether strategic interdependence in information decisions can arise even in the absence of any in-built interdependencies in the technology.

on scream volume if their group members exhibit similar information avoidance behaviors as those observed in the pilot study.[13] To check whether the treatment manipulation works as intended, after the information provision treatment we elicit the participants' beliefs about the fraction of avoiders among their group mates.

To assess the (average) interdependence of information decisions, we compare the fraction of information avoiders (i.e., people who chose option Later) between *Many Avoided* and *Few Avoided*.

4.1.3 Individual best-response schedules and heterogeneity

Proposition 3.1 suggests that individuals can be heterogeneous in how they respond to others' information decisions. To delve into this heterogeneity, we use the strategy method (Selten, 1967) to examine how each participant responds to a range of choices of their group mates, in an incentive compatible way. We elicit the full schedule of each individual's best response to all the possible choices made by others in the group. We classify the best reaction schedules based on whether a participant chooses Later when the share of avoiders among other members is sufficiently large (Strategic Complement) or sufficiently small (Strategic Substitute). If the participant's choice is the same for all shares of information avoiders in their group, we classify them as Always Getter (if they always choose to acquire information) or Always Avoider (if they always choose to avoid information). The remaining cases (namely, subjects who switch multiple times) are classified in a residual category labeled "Others". We refer to these categories of best responses as "strategy types", paralleling the theoretical classification defined in Section 3.1. Figure 1 illustrates examples of best reaction functions for each strategy type.

4.1.4 Additional measurements

Baseline attitudes towards information To measure baseline preferences over information, we administer the "Information Preferences Scale" ("IPS"; Ho et al., 2021). The IPS captures an individual's willingness to obtain information that may be unpleasant but is instrumentally valuable, across multiple domains. Higher IPS scores indicate greater general tendency to seek information. Measuring these general preferences serves three purposes: (i) to verify balance between treatment groups on a trait potentially relevant for information decisions; (ii) to include as a control to improve the precision of our estimates by reducing residual variance; and (iii) to use as a benchmark in the

[13] Specifically, we inform each participant that, if their group mates behave like the participants of the pilot, the volume in the *Screams* state would be, depending on whether she herself avoids or acquires information, 89 or 91 (in the *Many Avoided* condition), and 60 or 62 (in the *Few Avoided* condition). The difference of 2 points within each condition arises from the focal participant's choice. Conditional on her choice, the difference in scream volume between conditions arises from the information choices of others.

Figure 1: Examples of best-response functions, by strategy type

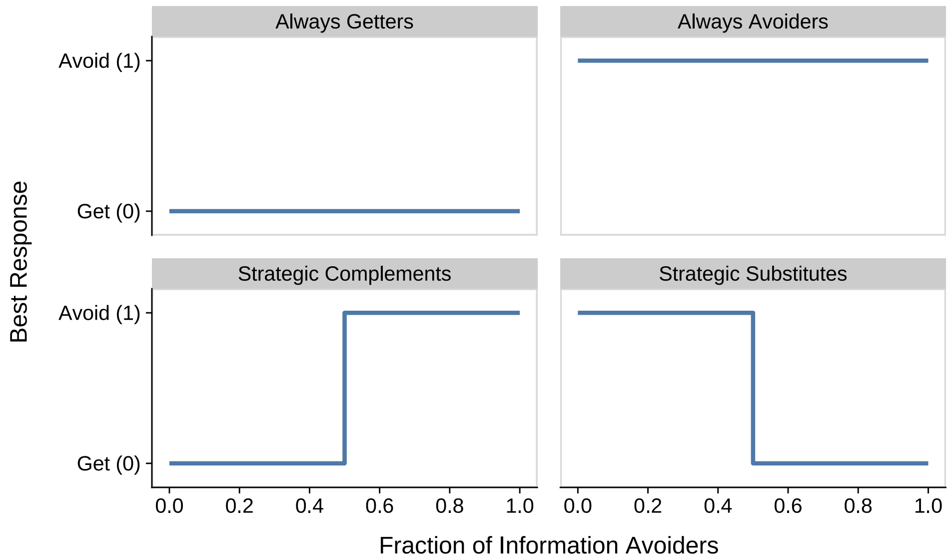

Notes: This figure illustrates examples of best-response functions, categorized by "strategy type". Strategy types classify best-response schedules by whether they are independent from others' choices ("Always Getters" and "Always Avoiders") or conditional ("Strategic Complements" and "Strategic Substitutes"). The x-axis represents the share of other group members avoiding information. The y-axis indicates the individual's choice to acquire or avoid information. Blue lines depict best-response schedules. Vertical lines indicate example switching thresholds for Strategic Complements and Strategic Substitutes (actual thresholds may vary across individuals).

validation of our novel elicitation of anticipatory utility (see below).

Anticipatory utility To shed light on mechanisms and determinants of information decisions, we introduce a novel elicitation designed to measure anticipatory utility over the future consumption event (screams or silence). A crucial feature of our elicitation is its granularity: anticipatory utilities are measured *conditional on each information set* —namely, upon acquiring information and discovering a positive signal ("good news" event), upon acquiring information and discovering a negative signal ("bad news"), and avoiding information and discovering no signal at all ("no news"). Specifically, we ask participants:

a) "How happy would you be *thinking about the outcome ahead* if you chose Now and then discovered that the lottery outcome is Screams?";

b) "How happy would you be *thinking about the outcome ahead* if you chose Now and then discovered that the lottery outcome is Quiet?"; and

c) "How happy would you be *thinking about the outcome ahead* if you chose Later?"

The elicitations above attempt to measure, respectively, the anticipatory utilities $v_i(b + \Delta)$, $v_i(g)$,

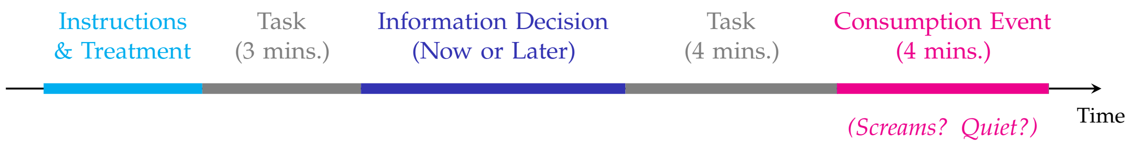

Figure 2: Timeline of experiment

and $v_i(pb + (1-p)g)$. To compute the net utility of avoiding information φ_i, we also elicit p_i, the participant's subjective belief that the state is Screams (i.e., the prior belief of finding out bad news). Together, these four questions measure the components of expression (2), which defines the net utility of avoiding versus acquiring information and theoretically determines the information decision. This granular elicitation of the components of φ_i allows us to examine how the treatment shifts the expected anticipatory utility of learning good versus bad news separately, which helps to distinguish among mechanisms. We call this set of questions the "Kreps-Porteus questions" (or "KP questions").

Social preferences Since information decisions generate payoff externalities on other members, social preferences may be a relevant predictor of the demand for information. To empirically investigate the relationship between social preferences and information decisions, we elicit individual measures by administering the Global Preferences Survey (GPS) (Falk et al. (2018)), which captures altruism and reciprocity, as well as trust, risk attitudes and time preferences.

4.1.5 Task and earnings

Following the design in Falk & Zimmermann (2024), we include an incentivized task for participants to complete in the first part of the experimental session, after the participants have taken the Now or Later decision and before the second part of the experiment where they may or may not experience the screams depending on the realization of the state. The experimental timeline is summarized in Figure 2.

4.2 Experimental Procedures

The experiment was conducted at the Centre for Decision Research and Experimental Economics (CeDEx) Lab, University of Nottingham.[14] We ran 17 experimental sessions (each constituting an independent group) across two waves of data collection (December 2023 and May–June 2024),

[14]The experiment was preregistered, AsPredicted #153871 and #175588. The experiment was approved by the University of Nottingham's School of Economics Research Ethics Committee. At the start of the session, participants were informed about the nature of the experiment and provided written consent to participate.

recruiting participants from the CeDEx subject pool. In total, the sample for analysis comprised 465 participants, with an average of 27.35 (± 3.66) subjects per session.[15] Each session lasted about one hour, and participants earned on average £12.52 (± 2.00), including a £5 show-up fee. The average score on the Information Preferences Scale (IPS) was 3.12 (± 0.41), indicating a general tendency to seek potentially undesirable information.[16] Full descriptive statistics of baseline characteristics are reported in Table B.1 in the Appendix.

At the start of each experimental session, participants were matched into groups, and each group was randomly assigned to a treatment condition.[17] Participants were then provided with detailed instructions on computer screens.[18] The instructions included a description of the state of the world lottery, the negative consumption event (screams), the timeline of the experiment, the information decisions with its implications, and the task. In order to ensure understanding of the decision environment, we included control questions about the timing of learning the outcome and about the nature of returns and externalities from information acquisition/avoidance. During the instructions phase, we also implemented our information provision treatment.

A computerized lottery determined the outcome (either *Screams* or *Quiet*) for each group. If the outcome was *Screams*, in the second part of the experiment the members of the group listened through headphones (worn throughout the session) to a series of unpleasant screams, drawn from the validated stimulus set in Beaurenaut et al. (2020).[19] To ensure hearing safety, the maximum possible volume level (100), was calibrated at 80dB in accordance with UK hearing safety regulations. The calibration was carried out separately for each headphone-computer pair, ensuring the hearing safety of every subject. To ensure exposure to the potential screams, participants were informed at the outset that headphone use was mandatory throughout the session and that attempts to remove them would be a basis for expulsion.[20] Compliance and correct headphone use were monitored by two invigilators per session. Participants were told they could leave the session at any time.

The use of screams instead of electric shocks (employed in experimental studies on information

[15] Originally, we recruited 493 participants in 18 sessions. However, due to technical issues, data (including the main outcomes of interest) was not recorded for 22 subjects in a single session of 28 participants and is excluded from the analysis; our results are robust to the inclusion of the 6 individuals from this session for which data is available. In a separate incident, one participant left early, about an hour after the session started, reducing observations for final-stage variables by one.

[16] A score of 3 means "Probably want to know".

[17] The random assignment of the treatment was blocked by week of year and time of session (morning and afternoon).

[18] Full instructions are available in Appendix Section D. The experiment was programmed in Lioness (Giamattei et al., 2020).

[19] Participants did not hear any sample screams before making their information decisions, but were informed that "[p]revious research has shown that people consider similar screams as *more aversive (disturbing) than* the sound of *nails sliding on a chalkboard*". Evidence of this statement can be found in Kumar et al. (2009).

[20] We are grateful to the people who generously helped as invigilators: Pierce Gately, Kieran Stockley, Tong Fang, Xue Wang, Matías Golman, Jesús Rodríguez, Yifan Li, Adrian Brown, Thomas Barnes, Lara Suraci, and Niramay Chugh.

avoidance in individual decision making) is based on three rationales. First, one of the motives that may lead to information avoidance are anticipatory emotions, such as anxiety about a future outcome, and studies show that the screams are able to sustain anxiety for longer periods than electric shocks (Beaurenaut et al. (2020)). This feature is especially desirable for experiments with group decision-making, which typically take longer compared to individual decision-making. Second, the provision of human scream sounds is more scalable than electric shocks, since the only required equipment are computers and headphones, as opposed to devices specialized for the provision of shocks.[21] Non-scalability imposes a natural limit to group sizes, which increase the likelihood of the emergence of complementarities in information decisions (Bénabou (2013)). Third, screams allow a more granular manipulation and more intuitive understanding of variations in severity (i.e., volume levels as opposed to electric shock intensity); this granularity in severity is important as group sizes become larger and the marginal return of acquisition of information (Δ) by any individual member becomes smaller. Fourth, the use of human screams alleviates ethical concerns compared to electric shocks.[22]

The experimental and model assumption is that screams are a bad. To assess whether this is the case, we elicit the participants' utility over two different levels of volume of screams. We ask: "Imagine that the Outcome is Screams, and that the volume of the screams will be level 100. How happy would you be thinking about the Outcome ahead?"; we ask a similar question for volume level 50. This allows us to identify subjects who, contrary to our assumption, derive higher utility from higher levels of volume of scream sounds ("volume lovers").

For groups whose lottery outcome was *Screams*, the volume of the screams was determined by the information decisions of its members. Specifically, the volume of the screams in group g was determined by the following function:

$$vol_g = 50 + (100 - 50)\left(\frac{1}{N_g}\sum_{i=1}^{N_g} a_{ig}\right), \tag{3}$$

where $vol_g \in [50, 100]$ is the volume of screams for group g conditional on the lottery outcome being Screams, N_g is the number of members in group g, and $a_{ig} = 1$ if subject i in group g

[21]The lower implementation requirement also facilitates replicability.

[22]Recently, experimental literature in Economics has started using the prospect of money losses as a negative future event ((Pagel, 2018; Engelmann et al., 2024). In particular, Engelmann et al. (2024) finds that money losses induce anticipatory anxiety and wishful thinking about the future event. The literature on intrinsic preferences for information has focused attention on future *consumption* events as distinct from *income* events, because knowing the latter introduces a planning advantage that motivates the acquisition of information about the future income event. To avoid introducing this extra planning advantage, in our experiment, we follow the latter tradition and focus on a consumption event. We argue that, theoretically, the relevant feature of the event is that the prospect of hearing human screams, much like the prospect of money losses, is perceived as a bad event by participants. In Section 5.1.1, we show evidence that, in our experimental sample, this is the case.

chooses Later and 0 otherwise. The acquisition of information about the state, $a_{ig} = 0$, thus partially mitigates the volume of the screams.

While the volume production function in Equation (3) appears related to a public good game (PGG), our decision environment differs in several ways that reverse the standard theoretical predictions. First, the nature of the cost is distinct: in a PGG, the cost of contributing is a known monetary amount, identical for all players and independent of others' actions. Here, there is no monetary cost of contributing: the cost is instead psychological (from potentially learning bad news about the future) and is heterogeneous across individuals. Moreover, the psychological cost plausibly varies with the choices of others (which influence how severe the bad state is). These differences lead to a divergence in the prediction for a purely material-payoff maximizing agent: such an agent would optimally free-ride in a standard PGG but would always prefer to "contribute" (i.e., acquire information) in our setting. Consequently, the central puzzle in each setting is inverted: whereas the literature on PGGs seeks to explain why people *do* contribute, in our setting the question is why people *do not* (i.e., why they avoid information) and how this reluctance to contribute is conditional on others' decisions.

Each group was composed by an entire experimental session, with group sizes of approximately $N_g = 30$. This implies that the return to each participant's information acquisition (in terms of volume reduction) was $1/30 * 50 = 1.67$ volume points.[23]

The information decision (the binary decision between options Now and Later) is our main outcome of interest. To obtain an alternative, continuous measure of preferences, we also ask participants, after they chose either Now or Later, to rate the strength of preference for their selected option on a Likert scale ranging from 0 ("I am indifferent") to 10 ("I have a very strong preference for the selected option"). From this measure, we construct the "Information Avoidance Scale" (IA scale) as follows. For those choosing Later, the IA scale equals the reported strength; for those choosing Now, it equals the negative of that strength. Thus, the IA scale ranges from -10 (very strong preference for Now) to +10 (very strong preference for Later), with 0 indicating indifference. The IA scale thus provides more variation than the binary outcome. We also ask participants to write down the reason for the selected option in an open text box.

To explore individual heterogeneity in best responses, we use the strategy method (Selten, 1967) to elicit each participant's complete schedule of choices between options Now and Later as a function of others' choices. This schedule captures their choice conditional on every possible share of information avoiders among others in the group, specifically for the shares: 0%, 20%, 40%, 60%,

[23]We described the return to information acquisition to participants as being equal to "around 2 out of 100 volume level points".

80%, and 100%.

We make the strategy method elicitation incentive-compatible by assigning a strictly positive probability that these choices are implemented. Specifically, before the elicitation, we inform participants about the following procedure. After the group members make their (unconditional) informational decisions (between Now and Later), we randomly select $N_g - 1$ subjects and calculate the proportion of information avoiders *among these* $N_g - 1$ *subjects*, based on their unconditional decisions. For the remaining participant, we implement her selected best response to this share.

To measure general preferences for information, we administer the "Information Preferences Scale", or "IPS" (Ho et al., 2021). This scale measures an individual's desire to obtain or avoid information that may be unpleasant but is instrumentally valuable, across several domains.[24] The scale ranges from 1 = "Definitely not want to know" to 4 = "Definitely want to know". Therefore, higher IPS scores indicate greater general willingness to obtain information.

To measure social and economic preferences, we administer the Global Preferences Survey (GPS) (Falk et al., 2018). The GPS measures risk attitudes, time preferences, altruism, reciprocity, and trust. Each dimension of preferences is elicited through the responses to two questions.[25] Following Falk et al. (2018), we standardize the answers, transforming them into z-scores. Higher z-scores indicate, respectively, higher risk seeking, patience, altruism, reciprocity, and trust. We also elicited demographics (e.g., age, gender), prior work experience, ratings about the perceived difficulty of instructions and quiz questions. Participants were given the opportunity to feedback on their experience in the experiment in open text boxes.

Following Falk & Zimmermann (2024), we include an incentivized task before the Now or Later decision and between such decision and the Risk Period. Participants are asked 90 general knowledge quiz questions, with earnings increasing in the number of correct answers. Each participant is paid according to her own performance, which is the only source of monetary earnings in our experiment, apart from the show-up fee.

5 Results

This section presents the analysis of our experimental data. We begin by establishing two necessary preconditions for our study. In Section 5.1.1, we verify that, consistent with our theoretical and

[24]The scale is constructed by averaging the answers to 18 questions that ask the subject's willingness to obtain information in hypothetical scenarios. The scenarios involve five domains: health, finance, social relations, ego-relevant characteristics, and occupation. For each question, the 4 possible answers are "Definitely don't want to know" (encoded as 1), "Probably don't want to know" (2), "Probably want to know" (3), and "Definitely want to know" (4).

[25]For risk attitudes and time preference, one of the questions used by Falk et al. (2018) involve the staircase method. Since administering this could be time-consuming and cognitively burdening in the middle of our experiment, we did not include such questions.

experimental assumption, screams are perceived as a bad by most participants. In Section 5.1.2, we document the existence of information avoidance in our decision environment (a context where the positive instrumental value and externalities from acquisition of information introduces an additional incentive to *acquire* it). Further, we establish that the decision to avoid information in our environment is likely not random nor driven by indifference, as information avoiders report strong preferences for their choice.

Following these prerequisite validations, we turn to our primary results. In Section 5.2, we leverage experimental variation in beliefs about others' information decisions to test for strategic interdependence, assessing whether the likelihood of information avoidance increases (strategic complementarity) or decreases (strategic substitutability) on average. In Section 5.3, we move beyond the average to document and characterize the individual heterogeneity of the interdependence. Using the strategy method, we document the empirical distribution of strategy types and report heterogeneous treatment effects by type. In Section 5.4, we explore the mechanisms driving the average treatment effect. Specifically, we use the KP questions to decompose the effect by identifying which components of the expression for the net incentive to avoid information (2) are most impacted by the treatment. We also discuss the plausibility of alternative mechanisms. Finally, the next section is devoted to analyzing how the group composition, in terms of strategy types, shapes the distribution of equilibrium share of information avoiders.

5.1 Preliminary Checks

5.1.1 Are screams a bad?

Before presenting our primary findings, we first conduct a validation check of a foundational assumption underlying our study: that participants actually perceive the prospect of listening to screams as a bad. We provide evidence confirming that the assumption holds for our sample. First, we asked subjects how "(un)happy" they would be if screams were played at volume 50 and, separately, at volume 100. The difference in these ratings provides a proxy of each participant's disutility of listening to screams at higher volume. Comparing these ratings reveals that a large majority of subjects report lower utility levels for screams at higher volumes: specifically, 82% of subjects report that screams at volume 100 would be strictly worse than at volume 50.[26]. Second, using the KP questions, we asked subjects to rate their "(un)happiness" upon discovering that the state was Screams after choosing Now, and separately, if it was Quiet. The difference in these ratings provides a measure of each participant's utility for listening to screams relative to silence. If

[26]Among the remaining subjects, 3% are indifferent between volume levels, and 9% would prefer volume 100. We call the later type of subjects "volume-lovers".

the screams are perceived as a bad, we would expect lower ratings for the Screams state than for the Quiet state. As expected, this pattern holds for the majority of subjects: 81% report that discovering that the Screams state is strictly worse than discovering that the Quiet state.[27] The distribution of types of preference screams and volume is shown in Figure B.2 and Figure B.1 in the Appendix. The evidence hence supports the view that the vast majority of subjects perceive screams as a bad.

Furthermore, the fact that higher-volume screams are perceived as worse than lower-volume ones is compatible with the notion that participants perceive the instrumental value of information acquisition (which allows them to reduce the volume of screams) as non-trivial.

For our main analysis, we consider all subjects, including those with weak preferences for screams and volume (or both). We refer to such subjects respectively as "scream-lovers" and "volume-lovers".[28] Analyses that exclude these individuals are presented in the Appendix Section B.5. They confirm that our findings are robust to the exclusion of volume- and scream-lovers, and, in some cases, become even stronger.

5.1.2 The Prevalence of Information Avoidance

In order to study the interdependence of information avoidance, it must first be established that avoidance occurs.

We find a positive and statistically significant share of information avoiders. Panel A of Figure 3 shows the prevalence of avoidance by data collection wave and for the pooled sample. In the pooled data, the prevalence is 0.286, significantly greater than zero at 1% level. The proportion is stable across waves.

This result relates to prior findings in the literature. In an individual decision making context, Falk & Zimmermann (2024) find that 52% of subjects avoid information. The higher proportion in their study could be attributed to a series of factors. In addition to some differences in implementation, in their setup information provides no material benefit (i.e., it does not affect the severity of the future negative consumption event), eliminating instrumental and prosocial motives for information acquisition.

The choice to avoid information does not appear to be driven by indifference nor trembling. Panel B of Figure 3 plots the distribution of preference strength among avoiders (people who chose Later), elicited on a 0-10 scale (0="indifferent"; 10="very strong preference"). The average strength of preference is 6.6, well above indifference. Moreover, the cumulative distribution (orange line)

[27] Among the remaining subjects, 6% are indifferent between the two states and 14% would prefer discovering the Screams state. We call the later type of subjects "scream-lovers".

[28] There are 111 subjects who are either scream lovers or volume lovers or both. The distribution of preferences is reported in Appendix Section B.2.

Figure 3: Information avoidance: prevalence and preference strength

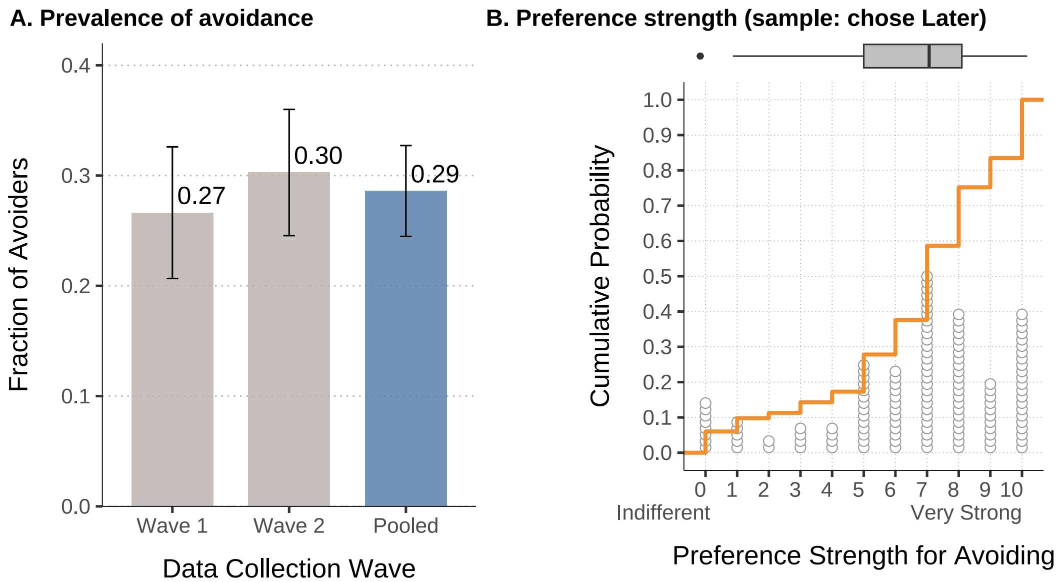

Notes: The figure displays the prevalence of and preference for information avoidance. Panel A shows the fraction of subjects who chose Later (avoided information), by data collection wave (brown bars) and in the full sample (blue bar). Whiskers indicate 95% confidence intervals. Pairwise comparisons using two-sample proportion tests yield p-values greater than 0.4. Panel B restricts attention to the subset of subjects who chose Later (avoided information) and shows the distribution of the reported strength of preference for this choice (0 = "I am indifferent" between Now and Later; 10 = "I have a very strong preference" for the chosen option). Gray dots represent the absolute frequency of each preference strength (one dot per subject). The orange line shows the empirical cumulative distribution function (CDF). The boxplot shows the median and the interquartile range, with whiskers extending to the most distant point within 1.5 times the interquartile range.

shows that only 6% of avoiders report indifference, while more than 80% report a strength of at least 5. Thus, evidence suggests that most information avoiders strongly prefer to remain uninformed, rather than choosing randomly or out of indifference.

5.2 The interdependence of information avoidance: experimental evidence

Having established the two prerequisites, we now turn to our main question: are information avoidance decisions interdependent across people? We examine this by leveraging exogenous variation in expectations about the prevalence of avoidance induced by the information provision experiment.

We first verify that treatment groups are balanced in all key covariates. Balance checks are reported in the Appendix Table B.3. Importantly, IPS scores are balanced between treatment groups, both for the overall scale and for each of the domain-specific subscale. Hence, any differences in information avoidance between treatments cannot be attributed to baseline differences in general or domain-specific attitudes toward information. The representation of some fields of study is somewhat imbalanced; to account for this, in subsequent analyses we report results both with and without field-of-study fixed effects in the regressions.[29]

Our treatment manipulation successfully shifts expectations about the share of avoiders among other members in the group. Panel A in Figure 4 reports the average belief by condition. As intended, the information-provision treatment substantially shifts beliefs upwards. On average, subjects in *Many Avoided* expect a 0.219 higher share of their peers to avoid information than those in *Few Avoided*, representing a 59.12% increase relative to the control group. The difference is statistically significant at 1% level. Appendix Table B.4 reports OLS estimates from a linear regression of beliefs on the treatment dummy and shows that the estimate of the average treatment effect is robust —both in magnitude and statistical significance— to the inclusion of different sets of controls (detailed in the Appendix). Furthermore, the shift in beliefs is not only in the mean: Panel B of Figure 4 plots the empirical CDFs of beliefs by treatment, showing that the distribution of beliefs in *Many Avoided* first-order stochastically dominates that in the *Few Avoided* condition. In sum, as intended, the experimental treatment is successful and effective in shifting beliefs about others' behavior.

We now turn to examine whether the treatment also generated differences in actual information decisions. Panel A of Figure 5 plots the proportion of information avoiders by treatment condition. Compared to *Few Avoided*, the share of avoiders in *Many Avoided* is 0.084 higher ($p = 0.045$), representing a 34.81% increase. Regression estimates from a linear probability model are displayed

[29] Compared to *Few Avoided*, in *Many Avoided* there is a lower share of Business and of Economics students.

Figure 4: Manipulation check

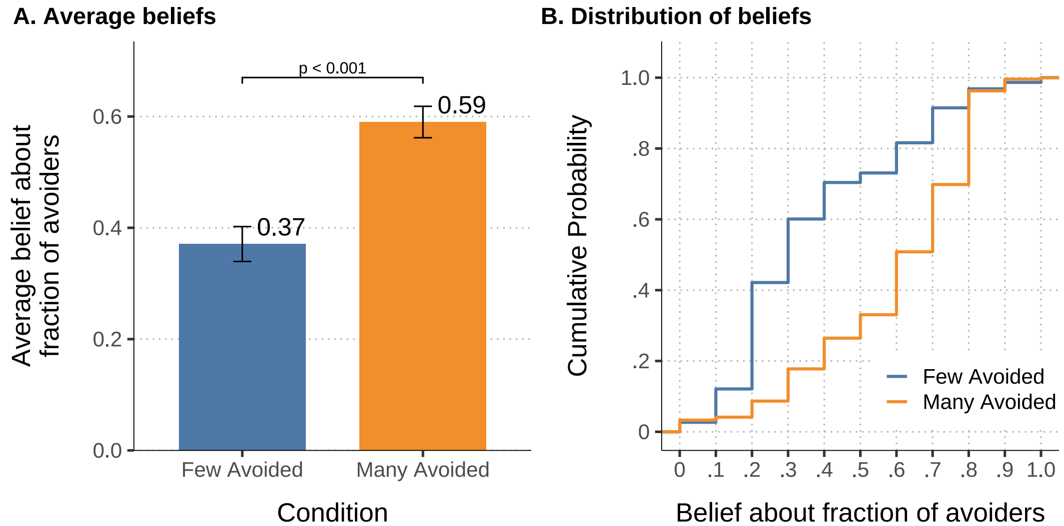

Notes: The figure shows subjects' beliefs about the fraction of information avoiders among other subjects in their own group, by treatment condition. Panel A shows the average belief by condition. Error bars show 95% confidence intervals; p-value from a two-sample t-test. Panel B shows the empirical CDF of beliefs by treatment condition.

in panel B of Figure 5, and reported in Table 1 in more detail. In the unconditional specification (blue circles, or column 1), the treatment effect is 0.084, significant at the 5% level. This effect is robust to the inclusion of individual characteristics and strata and field-of-study fixed effects (red crosses, column 2; $\hat{\beta} = 0.109$, significant at 5%), and increases to 0.123 when all controls are included, significant at the 1% level (gray squares, column 4).

The estimate of the treatment effect on information avoidance is robust to a range of alternative specifications. Excluding Scream- and Volume-Lovers from the sample makes the effects even stronger (Appendix Tables B.5 and B.6). Appendix Section B.6 shows that the estimate remains robust and significant across alternative specifications, regardless of whether the set of controls is selected by the researcher (Simonsohn et al., 2020) or via the data-driven post-double selection procedure (Belloni et al., 2014). Finally, we obtain qualitatively similar results when using as the outcome variable the IA scale (the continuous preference strength for information avoidance) (Appendix Section B.7).

To quantify the relationship between beliefs and behavior (the interdependence), we move beyond the reduced-form results and estimate a 2SLS model instrumenting beliefs with the treatment assignment (Appendix Table B.9).[30] The estimated coefficient ranges from $\hat{\beta}^{2SLS} \approx 0.4$

[30] The exogeneity assumption is satisfied by design. The monotonicity assumption —requiring that for every subject potential beliefs are weakly higher in *Many Avoided* relative to *Few Avoided*— is likely satisfied, as suggested by the first-order stochastic dominance of the belief distribution in *Many Avoided* relative to *Few Avoided* shown in Figure 4. The exclusion restriction requires that the effect of the information provision treatment operates only through posterior beliefs about others' information decisions. While this assumption is not directly testable, we discuss mechanisms in Section 5.4.

Figure 5: Information avoidance: treatment effect and predictors

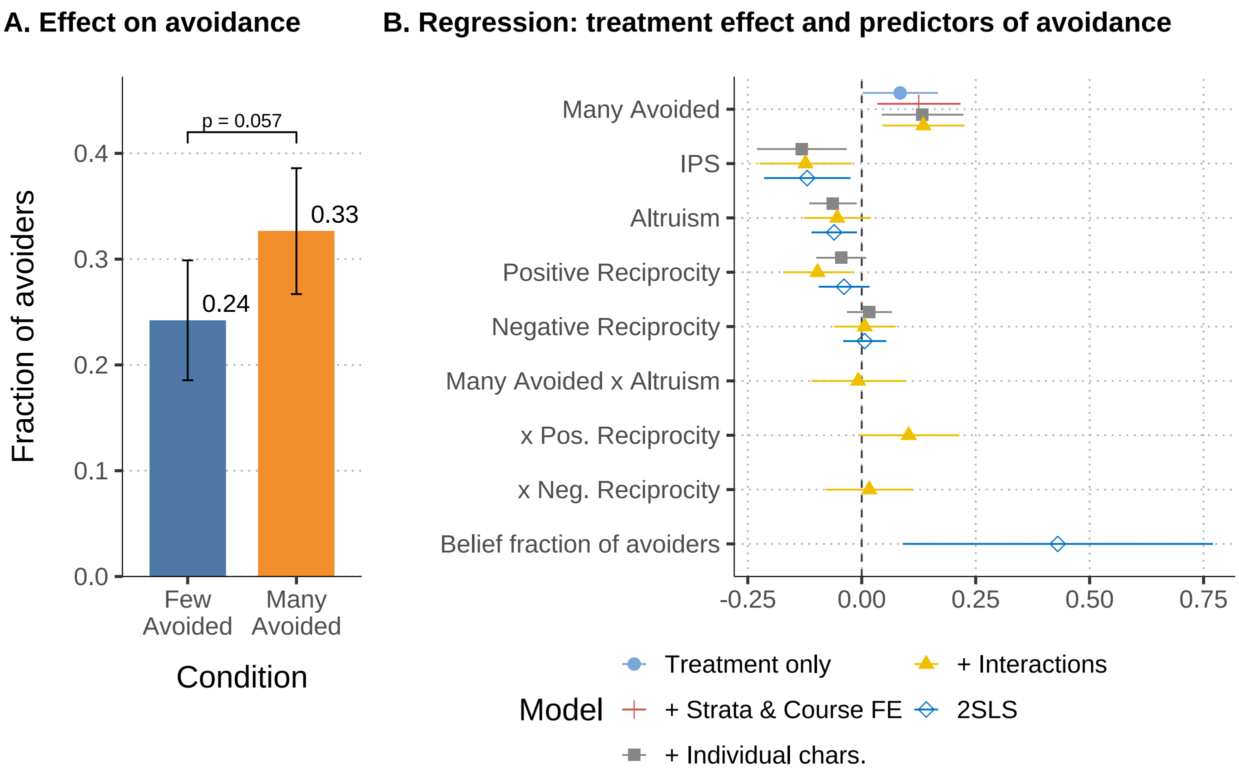

Notes: The figure shows the average treatment effect on information avoidance. In the *Few Avoided* condition, subjects were informed that 20% of participants avoided information in the past, and that, if current participants behaved similarly, screams volume would be around 60; in the *Many Avoided* condition, the analogous fraction was 80% (suggesting screams volume around 90). Panel A displays the fraction of avoiders (who chose Later) by treatment condition. Error bars indicate 95% confidence intervals; p-value derived from a two-sample proportion test. Panel B shows coefficients from linear probability models where the dependent variable is the "Avoided information" dummy (chose Later = 1, Now = 0). The first model (blue dots) includes only the treatment dummy. The second model (red plus signs) adds strata and field of study fixed effects. The third model (gray squares) further controls for individual characteristics: the Information Preferences Scale (Ho et al., 2021; with higher values indicating a stronger overall preference for *knowing* potentially undesirable information), gender, age, social preferences (altruism, positive and negative reciprocity), and economic preferences (risk aversion, patience, and trust; omitted from the figure for conciseness). The fourth model (yellow triangles) adds interactions between treatment and social preferences. The last model (blue hollow diamonds) instruments beliefs with the treatment assignment and shows 2SLS estimates. Error bars represent 95% confidence intervals computed based on robust standard errors.

to 0.6, implying that a 10 percentage point increase in expected avoidance by others raises an individual's probability of avoiding information by about 5 percentage points. This supports the conclusion that, on average, information avoidance is a strategic complement in our sample.

The regressions in Table 1 also show that social preferences are correlated with information decisions and may moderate the interdependence. Column (4) includes social preferences in the regression, and column (5) their interactions with the treatment dummy. The estimated coefficients associated to social preferences cannot be interpreted causally, and are interpreted as descriptive evidence. For reciprocity, we observe a significant and positive interaction with the treatment: more reciprocal individuals respond more strongly to the treatment, as they are on average less likely to avoid in the *Few Avoided*, and more likely in the *Many Avoided* condition. Intuitively, avoiding information can be interpreted as a way of taking revenge by harming others, so more reciprocal individuals are less inclined to avoid information when only a minority avoids information (*Few Avoided*) than when the majority does (*Many Avoided*). However, as Appendix Table B.5 and Table B.6 show, when Scream- and Volume-Lovers are excluded, the estimated coefficients shrink and become statistically insignificant, although the signs remain unchanged. For altruism, our data indicate that more altruistic individuals are less likely to avoid information ($\hat{\beta} = -0.068$), consistent with the idea that prosocial motives encourage information acquisition to benefit others.

Lastly, the Information Preferences Scale significantly predicts behavior in our experiment: individuals with a stronger tendency to seek potentially undesirable information (higher IPS scores) are less likely to avoid information in our experiment.

In summary, we provide experimental evidence that, in groups, deliberate avoidance of free and useful information can be interdependent, and in particular, contagious.

	Dependent variable: Avoided information				
	(1)	(2)	(3)	(4)	(5)
Many Avoided	0.084**	0.109**	0.114**	0.123**	0.124**
	(0.042)	(0.049)	(0.049)	(0.049)	(0.049)
Information Preferences Scale			-0.125**	-0.130**	-0.118**
			(0.051)	(0.052)	(0.054)
Social preferences					
Altruism				-0.065**	-0.052
				(0.028)	(0.041)
× Many Avoided					-0.006
					(0.057)
Positive reciprocity				-0.046	-0.101**
				(0.029)	(0.041)
× Many Avoided					0.109*
					(0.060)
Negative reciprocity				0.013	-0.005
				(0.026)	(0.036)
× Many Avoided					0.030
					(0.051)
Constant	0.242***	0.126	0.326	0.280	0.273
	(0.029)	(0.225)	(0.301)	(0.311)	(0.321)
Strata FEs		✓	✓	✓	✓
Field of study FEs		✓	✓	✓	✓
Individual characteristics			✓	✓	✓
Observations	465	465	460	460	460

* p < 0.1, ** p < 0.05, *** p < 0.01. Robust standard errors in parentheses.

Notes: The table reports estimates from a linear probability model regressing information avoidance on the treatment dummy. The dependent variable is a binary indicator equal to one if the subject avoid information (chose Later) and zero otherwise. The treatment dummy equals one if the subject was in the *Many Avoided* condition and zero in the *Few Avoided* condition. The regressions include strata fixed effects, demographic controls, field of study, and measures of economic and information preferences (see notes in Table B.4). The number of observations drops in column (3) to (5) since five participants did not provide information about their age.

Table 1: Treatment effect on information avoidance

5.3 The heterogeneity of interdependence

Section 5.2 showed that exogenously increasing expectations about the share of avoiders in a group raises the likelihood of information avoidance, on average. A key question, however, is whether individual responses to the information decisions of others are heterogeneous. As Proposition 1 showed, theory allows for a variety of best-response types. To empirically identify this heterogeneity, we employ the strategy method to elicit a *schedule* of best responses to different shares of avoiders in the group. As detailed in Section 4.1, we categorize these empirical schedules based on whether they are increasing in the share of avoiders (Strategic Complements), decreasing (Strategic Substitutes), or constant (either Always Getters or Always Avoiders). This allows addressing the empirical question: do we observe the theoretical heterogeneity in practice? How large is this heterogeneity?

The data reveal substantial heterogeneity in how individuals respond to the information avoidance decisions of others. Panel A of Figure 6 presents the empirical distribution of individual strategy types in the whole sample (red bars). The most prevalent type is the Always Getter (42.4 %), who seeks information regardless of the share of avoiding peers. At the opposite extreme, the least common type is the Always Avoider (8.8%), who avoids learning the state irrespective of the share of avoiders.

Beyond these unconditional types, we document the existence of conditional —or interdependent— types, whose behavior depends on that of others. Among these types (approximately 40% of our sample), we identify two patterns. 20.6% of participants are Strategic Complements: they avoid information only if a sufficiently high share of their group members also avoids it. The proportion of Strategic Complements is statistically different from zero, with a 99% confidence interval of (0.171, 0.246). A similar share (19.4%) are Strategic Substitutes, who avoid information only if a sufficiently *low* share of their group members avoids it. The remaining minority (8.8%) are subjects who switch their information decisions multiple times and, hence, whose best responses oscillate. In sum, we find substantial heterogeneity in strategy types, with a significant share of individuals exhibiting dependence on the choices of others.

We now assess the treatment effect heterogeneity by strategy type, by examining both the extensive margin (whether the treatment affects the distribution of types) and the intensive margin (how the treatment changes the probability of avoidance conditional on the type).

On the extensive margin, we find, reassuringly, that the treatment has no statistically significant effect on the distribution of types. Panel A of Figure 6 shows the distribution split by treatment condition (light and dark gray bars), which appear remarkably similar. A regression analysis confirms this graphical intuition. Appendix Table B.10 reports estimates from regressions where a dummy for each strategy type is regressed on the treatment indicator. In these regressions, the

Figure 6: Heterogeneity of strategy types

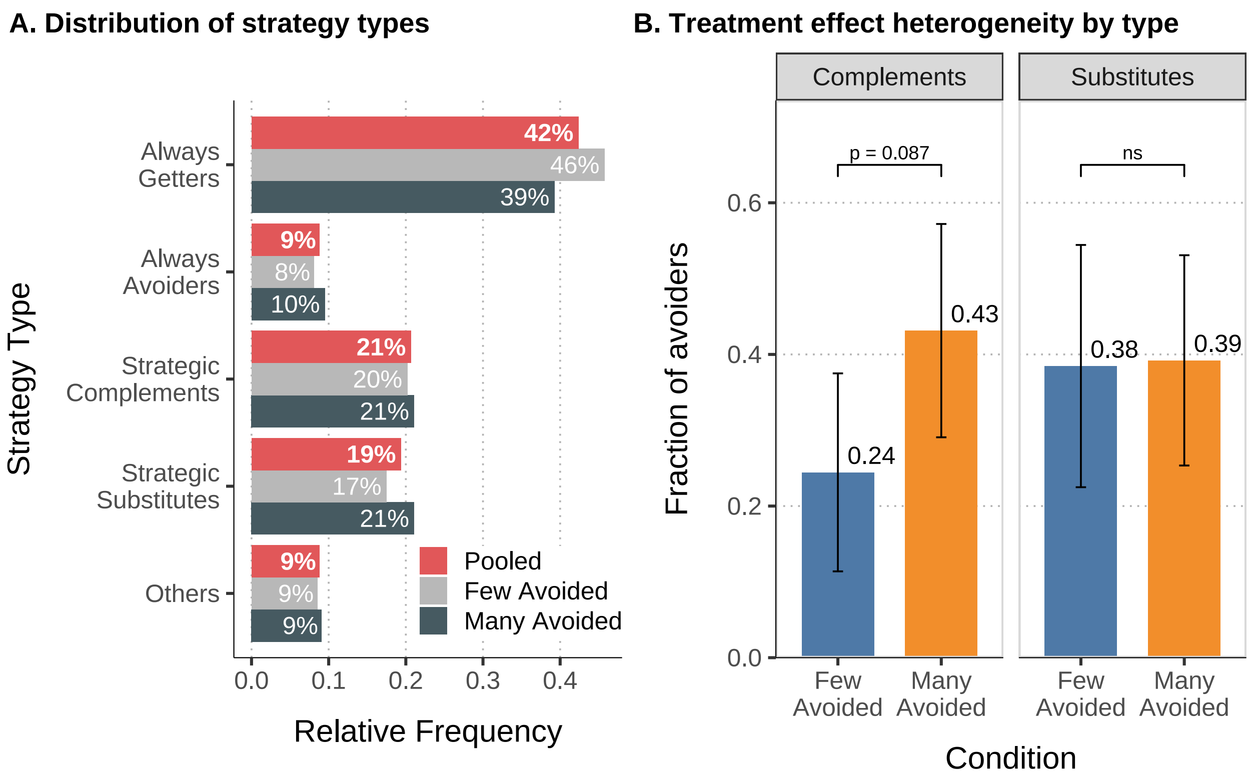

Notes: The figure illustrates the heterogeneity of strategy types. Panel A shows the distribution of types in the pooled sample (red bars), and separately for each treatment condition (gray bars), capturing the extensive margin (i.e., whether the treatment affects the distribution of types). Panel B shows the intensive margin: the treatment effect on the fraction of information avoiders conditional on the strategy type (Complements and Substitutes). The y-axis represents the fraction of subjects who avoided information (chose Later). Error bars represent 95% confidence intervals; p-values derived from a two-sample proportion tests.

treatment coefficient captures the change in the proportion of each type between *Many Avoided* and *Few Avoided*. No statistically significant differences emerge for any type.

The distributional invariance across treatments is consistent with theory. Best-response *schedules* themselves should not change unless some (perceived) parameter of the environment —beyond the expected behavior of others— changes. Since the treatment was designed to change only the expectation about others' behavior, but no other feature of the environment, there is no theoretical reason to expect the distribution of types to shift across treatments. Accordingly, we find that, empirically, the distribution of strategy types remains similar between treatments.

We now turn to the intensive margin. Panel B of Figure 6 compares the share of avoiders (who chose Later) between treatment conditions conditional on the strategy type. Table 2 reports a linear regression of the avoidance dummy on the treatment for each type. In these regressions, the treatment coefficient captures a conditional average treatment effect (CATE): the difference in the likelihood of avoidance between the *Many Avoided* and *Few Avoided* conditional on the strategy

	Dependent variable: Avoided information			
	(1) Always Getters	(2) Always Avoiders	(3) Complements	(4) Substitutes
Many Avoided	0.014	0.056	0.187*	0.008
	(0.033)	(0.057)	(0.096)	(0.106)
Control group mean	0.049**	0.944***	0.244***	0.385***
	(0.022)	(0.057)	(0.066)	(0.080)
Observations	197	41	96	90

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Notes: The table presents estimates from linear probability models where the dependent variable is an indicator equal to 1 if the subject avoided information (chose Later). Each column reports a separate regression estimated on the subsample of a distinct strategy type. The independent variable is the treatment dummy (equal to 1 for the *Many Avoided* condition and 0 for the *Few Avoided* condition.

Table 2: Treatment effect on information avoidance, by strategy type (intensive margin)

type. We find that the *Many Avoided* treatment significantly increases the likelihood of information avoidance for Strategic Complements. This is in line with theoretical intuitions: since the treatment successfully raised beliefs about others' avoidance —the argument of best-response schedules— and Complements' schedules are increasing, they should respond by avoiding more often. By contrast, however, we found no significant response among Strategic Substitutes.

The difference in responsiveness between Strategic Complements and Substitutes can be explained by the interaction between switching threshold distributions and the treatment-induced shift in beliefs, shown in Figure 7. Average beliefs are shown in pink lines and average switching threshold in black lines. While the treatment leaves the average switching thresholds unchanged for both types (around 0.5 for Complements and around 0.4 for Substitutes), it substantially increases average beliefs (from 0.35 to 0.62 for Complements and from 0.42 to 0.64 for Substitutes). For Strategic Complements, the belief shift crosses the threshold, triggering increased avoidance. For Strategic Substitutes, however, the new beliefs remain on the same side of the threshold, resulting in no significant change in average avoidance rates.

Finally, columns (1) and (2) of Table 2 show that, as expected, the treatment does not affect the avoidance rate for Always Getters and Always Avoiders, who, by definition, follow strategies that are irresponsive to the behavior of others. Taken together, the pattern of intensive margins of treatment effects by types suggest that the positive average treatment effect documented in Section 5.2 is primarily driven by the responsive behavior of Strategic Complements.

Our analysis reveals considerable heterogeneity underlying the average treatment effect. The significant increase in information avoidance is not uniform, but is instead almost entirely driven

Figure 7: Switching thresholds and beliefs by strategy type and treatment condition

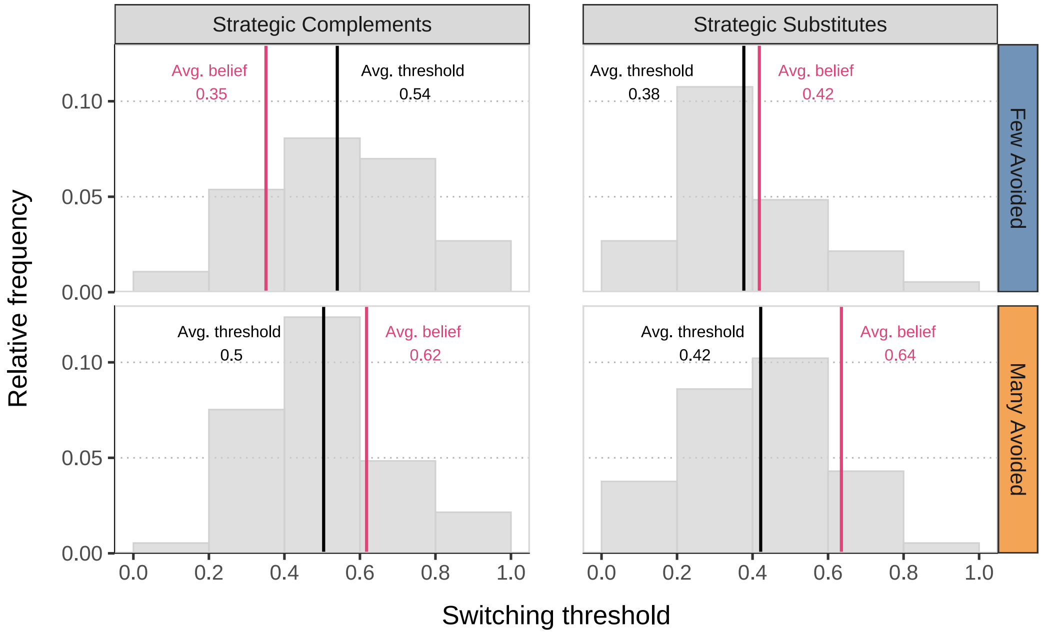

Notes: The figure shows the switching thresholds —the fraction of other group members avoiding information at which a subject changes her own decision— along with beliefs about the fraction of others avoiding information, by strategy type and treatment condition. The distribution of individual switching thresholds is represented by gray bars. The average threshold is indicated by a black line, and the average belief about the share of avoiders is shown by a pink line.

by Strategic Complements, while the behavior of Always Getters, Always Avoiders, and Strategic Substitutes is unchanged. This pattern of responses is explained by the interaction between switching thresholds and beliefs.[31] Our findings suggest that understanding *which* individuals respond and *how* they respond (in terms of direction) is important for predicting aggregate acquisition of information. We explore this point through simulations that vary the composition of the group in terms of strategy types and trace the equilibrium prevalence of information avoidance in Section 6.

5.4 Mechanisms

Section 5.2 showed a positive average treatment effect. In this subsection we explore underlying mechanisms: what makes subjects in the *Many Avoided* condition avoid discovering the state more often than subjects in the *Few Avoided* condition?

The strategic complementarity in information avoidance could be driven by a variety of mechanisms. First, we focus on the "Mutually Assured Ignorance" (MAI) channel, formalized by Bénabou (2013) and central in our model: the idea that when a subject expects others to

[31] Although Strategic Substitutes were unresponsive in this specific intervention, their behavior can potentially change under larger changes in beliefs. The data from the strategy method, which maps the best-response schedule for a wider range of beliefs, allow predicting decisions under such counterfactual beliefs.

remain ignorant about the state, her own payoffs in the bad state become worse, making the event of discovering bad news about the state more aversive. In Section 5.4.1, we provide evidence supporting this mechanism by using our KP measures of expected anticipatory utilities conditional on finding out good, bad, and no news. We then discuss the plausibility of alternative mechanisms in Section 5.4.2.

5.4.1 The Mutually Assured Ignorance mechanism

Anticipatory utility and information avoidance To shed light on determinants of information decisions and on mediators of the treatment effect, we use the KP questions to measure the expected anticipatory utility associated to the decisions of acquiring and avoiding information. First, we show that our measures of anticipatory utility significantly predict actual information decisions in our experiment, above and beyond the predicting capacity of the IPS, an established measure of information preferences. Next, we provide suggestive evidence that the treatment operates mainly through the utilities as captured by the KP measures. Finally, and crucially, we show that the treatment increases the net incentives of avoiding information by lowering the expected anticipatory utility of finding out *bad* news, but has no effect on the expected anticipatory utility of finding out *good* nor *no* news, in line with the MAI mechanism in Bénabou (2013) and with our model.

Recall that the KP questions elicit subjects' estimation of their own anticipatory utility *conditional on each possible information set* —that is, upon finding out good news ($\hat{v}_i(g)$), bad news ($\hat{v}_i(b+\Delta)$), and no news ($\hat{v}_i(\hat{p}_i b + (1-\hat{p}_i)g)$)—. In addition, the KP questions elicit the subjective belief that the state is bad ($\hat{p}_i$). All the KP questions are elicited before subjects make their information decisions, so that we interpret them as *expected* anticipatory utilities, where expectations are taken before the subject potentially learns any information.

Based on the answers to the KP questions, for each subject we infer the expected anticipatory utility of acquiring information, by taking the expectation over states

$$U_{0,i}^{Get} \equiv \hat{p}_i \cdot \hat{v}_i(b+\Delta) + (1-\hat{p}_i) \cdot \hat{v}_i(g) \tag{4}$$

based on the subjective belief about the state, $\hat{p}_i$. On the other hand, the expected anticipatory utility of avoiding information is simply

$$U_{0,i}^{Avoid} \equiv \hat{v}_i(\hat{p}_i b + (1-\hat{p}_i)g) \tag{5}$$

We then estimate the net expected anticipatory utility of avoiding information φ_i, defined in

Equation (2), as

$$\hat{\varphi}_i \equiv U_{0,i}^{Avoid} - U_{0,i}^{Get}, \tag{6}$$

The "hat" notation in $\hat{\varphi}_i$ emphasizes that this is our empirical estimate of the theoretical construct φ_i.

The model predicts that an individual i should avoid information whenever the net utility $\varphi_i > 0$, acquire it if $\varphi_i < 0$, and be indifferent if $\varphi_i = 0$. To assess the predictive power of our empirical anticipatory utility measures, we regress the avoidance dummy a_{ig} on the net expected anticipatory utility, $\hat{\varphi}_i$, or alternatively, a dummy variable indicating $\hat{\varphi}_i > 0$.[32] To benchmark our measures against an established alternative, we also include the IPS in the regressions. All specifications control for field of study and strata fixed effects, demographic characteristics, and economic preferences.

Panel A in Table 3 reports the results. Column (1) shows that, without the IPS nor the KP measures, the controls explain only 6% of the variation in choices, as indicated by the adjusted R^2. Column (2) adds the Information Preferences Scale (IPS). The IPS is a significant negative predictor of avoidance, consistent with the interpretation that individuals with a general willingness to *know* potentially undesirable information are *less* likely to avoid information in our experiment. Adding the IPS to the model improves explanatory power only modestly (R^2 = 7%). Columns (3) and (4) replace the IPS with our measures of anticipatory utility. Both versions of anticipatory utility, continuous and binary, are highly significant (1% level) and positively associated with avoidance: a higher (positive) net anticipatory utility of avoidance is associated with a higher probability of avoiding information. Notably, models including our measures explain approximately 20% of the variation, nearly three times than that with the IPS. This is possibly because the KP questions elicit aspects that are specific to the decision environment. Columns (5) to (6) include both the IPS and KP measures. Perhaps surprisingly, for both measures, the coefficients remain remarkably stable in sign and magnitude. The KP measures also explain variance above and beyond the IPS. An interpretation is that the information decisions are driven by both a general disposition (captured by the IPS) and environment-specific aspects (captured by the KP measures). This pattern of findings supports both the predictive and discriminant validity of our KP measures, suggesting that the latter are not merely proxies of general preferences for information, but instead elicit aspects of preferences for information that are specific to our environment and predictive of decisions. Overall, our net utility measure is a more powerful predictor of information avoidance, and its predictive

[32]The dummy is equal to 0 if $\hat{\varphi}_i < 0$, and encoded as a missing value for border cases where $\hat{\varphi}_i = 0$ (12 individuals in our sample).

power is independent of that of the IPS.

Next, we assess whether anticipatory utilities mediate the treatment effect. Existing methodologies of causal mediation typically rely on the often strong assumption that the mediator is independent of unobservables conditional on the treatment. Since this assumption may not hold, the evidence reported in the remainder of this paragraph may be interpreted as tentative. Panel B presents regressions that include both the treatment dummy and the IPS or KP measures. For reference, column (1) reproduces the treatment effect estimate in the baseline specification from in Table 1. Column (2) adds to this regression the IPS as a benchmark, with the treatment coefficient remaining virtually unchanged. Columns (3) and (4) replace the IPS with our measures of net expected anticipatory utility. When these utility measures are included, the treatment effect shrinks substantially in magnitude and becomes statistically insignificant (for the continuous utility specification) or less significant (binary specification). Moreover, comparing the coefficients of the IPS and KP measures (columns 2-6) between Panel A (which exclude the treatment dummy) and Panel B shows that the inclusion of the treatment dummy leaves those coefficients essentially unchanged. This pattern of findings is consistent with a mediating role for anticipatory utility: the treatment appears to operate primarily by influencing this utility. This theoretically intuitive: the decision should depend only on the associated utility, so once we control for the utility, variations in the treatment dummy should become insignificant.

	Dependent variable: Avoided information					
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Predictive power						
Information Preferences Scale (IPS)		-0.129**			-0.136***	-0.148***
		(0.052)			(0.047)	(0.049)
Net utility of avoidance (continuous)			0.023***		0.023***	
			(0.003)		(0.003)	
Net utility of avoidance (dichotomous)				0.372***		0.379***
				(0.046)		(0.045)
Controls and fixed effects	✓	✓	✓	✓	✓	✓
Observations	460	460	460	448	460	448
Adjusted R^2	0.059	0.071	0.185	0.211	0.199	0.226
Panel B: Treatment and utility						
Many Avoided	0.122**	0.123**	0.061	0.081*	0.062	0.080*
	(0.049)	(0.049)	(0.047)	(0.045)	(0.046)	(0.045)
Information Preferences Scale (IPS)		-0.130**			-0.136***	-0.147***
		(0.052)			(0.047)	(0.048)
Net utility of avoidance (continuous)			0.022***		0.022***	
			(0.003)		(0.003)	
Net utility of avoidance (dichotomous)				0.365***		0.372***
				(0.047)		(0.046)
Controls and fixed effects	✓	✓	✓	✓	✓	✓
Observations	460	460	460	448	460	448
Adjusted R^2	0.070	0.082	0.187	0.214	0.200	0.230

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Notes: The table presents OLS estimates analyzing the predictive power and potential mediating role of the KP measures on actual information decisions. The dependent variable is an indicator equal to 1 if the subject chose to avoid information about the state (option Later) and 0 otherwise. Panel A assesses the predictive power of a number of measures. "Information Preferences Scale (IPS)" (Ho et al., 2021) measures overall preference to know potentially undesirable information; higher values indicate stronger preference to know. "Net utility of avoidance (continuous)" is the empirical measure of the net expected utility of avoiding information (relative to acquiring it), $\hat{\varphi}^i$, derived from the KP questions. "Net utility of avoidance (dichotomous)" is an indicator equal to 1 if the empirical net utility of is higher than zero, $\hat{\varphi}^i > 0$. In Panel A, coefficients are standardized for comparability across measures. Panel B focuses on how the treatment dummy ("Many Avoided") changes upon inclusion of the predictive measures; coefficients are not standardized for comparability with Table 1. All regressions control for demographic characteristics (age and gender), economic and social preferences (patience, risk-seeking, trust, altruism, and positive and negative reciprocity), and fixed effects by strata and field of study of the participant.

Table 3: Anticipatory utility and information avoidance

To further illustrate this point, in Figure 8 we report a probit model where the latent variable is the net utility $\hat{\varphi}_i$. The model is estimated separately for the treated (*Many Avoided*) and control (*Few Avoided*) subsamples. The probability of avoiding information increases with the empirical net utility $\hat{\varphi}_i$. Crucially, this relationship is statistically indistinguishable between treatment conditions. We repeat the same exercise across interdependent strategy types (Complements and Substitutes, where there is variation in choices) and find the same pattern. Again, this is in line with the theoretical notion that for predicting information avoidance, once we condition on the net utility, it shouldn't matter whether the subject is a Strategic Complement or Substitute.

Overall, these findings suggest that the KP measures, designed to measure anticipatory utility, proxy relevant determinants of information decisions and act as (or at least, are correlated with) a key mediator of the treatment effect.

Figure 8: Association between anticipatory utility and information avoidance (Probit)

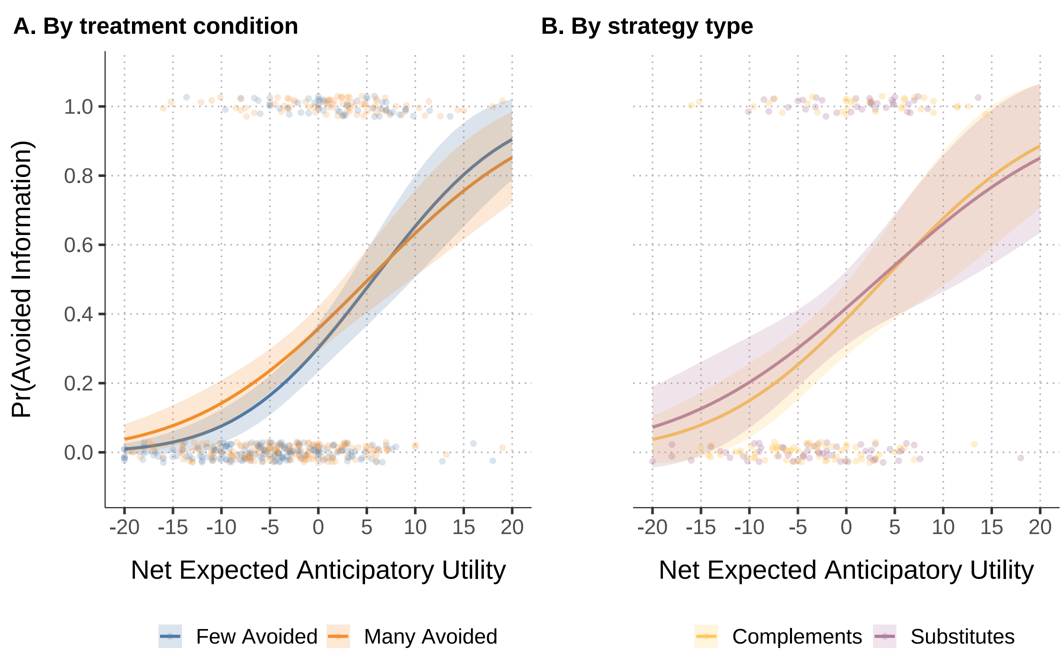

Notes: The figure illustrates the relationship between the measured net expected anticipatory utility of avoiding information and the propensity to avoid information. The net utility is calculated as the difference between the utility of avoiding information and the expected utility of acquiring information, elicited from the KP questions. Each point represents an individual participant's data, jittered vertically. The solid line shows the predicted probability of information avoidance from a probit regression model, estimated separately for each treatment condition (panel A) and for each strategy type (panel B). The shaded bands indicate the 95% confidence intervals.

Prevalence of avoidance among others, anticipatory utility, and contagion Having established the validity of our KP measures, we now look deeper at specific components of the net utility of avoiding information. This allows us to provide a decomposition of the treatment effect on such utility, useful to assess the plausibility of the MAI mechanism.

In separate regressions, we estimate the effect of the treatment on the following potential mediators: a) the net utility of avoiding information, $\hat{\varphi}_i$, and b) each of its components separately; both excluding the set of controls and fixed effects ("no control" models) or including them ("control" models). The treatment coefficient of each regression, capturing the effect on each mediator, is shown in Figure 9. These coefficients have a causal interpretation, and are standardized for comparability across measures. Starting from the top of the figure, the first pair of coefficients shows that the *Many Avoided* treatment has a positive and highly significant effect on the net utility of avoiding information —which, as shown above, is predictive of actual information choices. How is this effect decomposed? The second pair of coefficients shows the effect of the treatment on the expected utility of finding out *bad* news about the state upon acquiring information. The anticipatory utility of discovering bad news significantly deteriorates in the *Many Avoided* condition relative to the *Few Avoided* condition. This is consistent with the model, as the payoff in the bad state b (severity of the screams) worsens as more group members avoid information. In contrast, the third pair of coefficients shows that the treatment has no effect in the expected utility of finding out *good* news. This again aligns with the theoretical intuition: since the payoff of the good state (silence, g) is unaffected by the number of avoiders in the group, there should be no treatment effect on the utility of discovering good news. Finally, the treatment also increases the utility of avoiding information and learning *no* news.

Unexpectedly, the treatment also increases the subjective belief that the state is bad. In principle, this suggests the possibility that changes in beliefs about the state in response to the treatment could be part of the mechanism, either as a competing channel or by interacting with and amplifying the MAI mechanism, as follows: updating beliefs about the information decisions of others may deteriorate the utility of learning bad news and also induce the subject to believe that the bad state is more likely, further reducing incentives to acquire information. The plausibility of this alternative channel can be informed by the empirical literature on the relationship between prior beliefs about the state and information decisions. In the setting closest to ours, Falk & Zimmermann (2024) find that, in an individual decision-making setting, experimentally manipulating prior beliefs about the bad state (p_i) does not lead to significant changes in information avoidance.

Taken together, the treatment effects on each utility component combine to increase the implied *net* anticipatory utility of avoiding information, suggesting that one of the main sources of stronger incentives to avoid information when others do so arises from the expected deterioration of anticipatory utility in the event of discovering bad news, a result consistent with the MAI mechanism of Bénabou (2013) and of our model.

Figure 9: Mechanisms: treatment effects on anticipatory utility

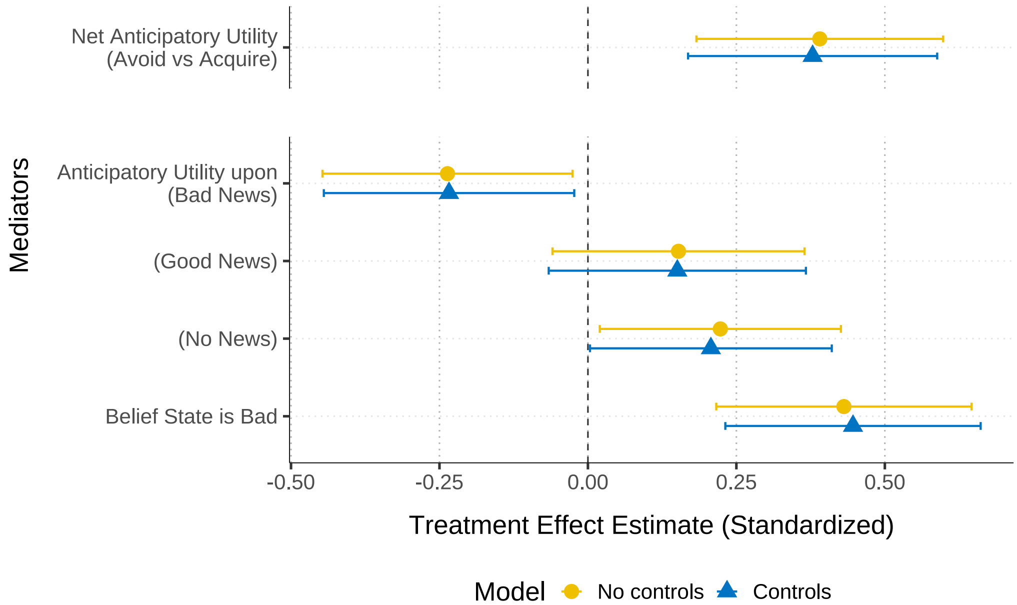

Notes: The figure reports the treatment effects on net anticipatory utility and each of its components from separate regressions. Yellow circles show estimates from models without further controls; blue triangles correspond to models that include controls for demographics, preferences, and fixed effects (see notes in Figure 5). The top row reports effects on the net anticipatory utility (of avoiding information relative to acquiring it); lower rows report effects on its components: the anticipatory utilities upon learning bad, good, or no news, and the subjective belief that the state is bad. Coefficients are standardized; whiskers indicate 95% confidence intervals (robust standard errors).

5.4.2 Other Mechanisms

While the evidence is consistent with the MAI channel, the interdependence of information avoidance could, in principle, arise from other behavioral mechanisms. We discuss some of the main possibilities below.

Conformity and herding. A first alternative possibility is that the positive correlation between beliefs about others' behavior and own behavior is driven by conformist preferences. We say that an individual is conformist if he derives direct utility from matching his actions to those of others. Alternatively, the positive correlation may be driven by information herding. Under this mechanism, subjects are unsure about which is the most convenient course of action (i.e., to avoid or acquire information) in our particular experimental setting; if they believe that the signal about others' behavior in the past is more informative about the most convenient decision than their own priors, they may rationally follow the crowd and match their actions to that of others. While each mechanism is grounded in distinct psychological motivations, they all formally predict the same empirical pattern: an individual's likelihood of avoidance increases with the perceived fraction of avoiders in their group. Given the similar empirical predictions, we label these mechanisms jointly as “conformity” for brevity. To assess them, we conduct four tests. Specifically, if the observed treatment effect is purely driven by conformity, then we would expect four observable implications. First, since the treatment effectively changes the beliefs about others' information decisions, the utility of acquiring information and discovering the *good* state should differ across conditions. We focus on the event of acquiring information and discovering good news because material payoffs g are invariant in the share of avoiders, and hence, constant across conditions (unlike the payoffs in the bad state), allowing a cleaner comparison. As Figure 9 shows, we don't find support for this implication of conformity. Second, under a pure conformity channel the belief about the share of avoiders should predict own avoidance *independently* of the subjective belief about the state. In contrast, in the groupthink channel, what others do is relevant to the extent that the subject believes that the state is bad. To test these predictions, we regress the dummy a_{ig} on both beliefs and their interaction, with results reported in Appendix Table B.11. We find that, contrary to the conformity channel, the interaction of both beliefs is highly significant, suggesting that the belief about others' behavior predicts own avoidance only if the subject believes that the probability of the bad state is high. Moreover, once the interaction is included, the coefficient on the belief about others' information decisions becomes insignificant. This pattern lends support to the Bénabou's (2013) mechanism rather than to conformity. Third, conformity, if effective, would be manifested in Strategic Complements only. Assuming that subjects derive utility from imitating behavior of people from *both* past and current sessions, we should observe a shift in the switching threshold of Strategic Complements caused by the treatment, given that the strategy method elicits choices conditional on the choices of group members in the current, but not past, session. As Figure 7 showed, we do not find evidence of different switching thresholds across conditions. Fourth, under a pure conformity channel the effect of the treatment should not interact with reciprocity. However, we find mild evidence of the interaction, as shown in Table 1. Taken together, the four tests do not provide significant support for either the conformity or the information herding channel, while in some instances reinforce the evidence in favor of the groupthink mechanism.

Social learning. A second alternative is social learning. Under this mechanism, subjects do not know the parameters of their preference over screams, and use the signal about behavior in the past to update their beliefs about their own preference parameters. As a result, this may generate differences in information avoidance across treatment conditions. Although we do not have a direct test for this channel, we expect this mechanism to be negligible.

Non-linear preferences over volume. Finally, a third alternative mechanism relates to differences in the perceived benefit of a reduction of the 2 volume points across conditions. If subjects perceive that a 2-point reduction is less valuable in the *Many Avoided* condition relative to the *Few Avoided* condition, they would have higher net incentives to avoid information. While this mechanism is possible in principle, casual experience suggests that an increase of 2 volume points in a computer (out of 100 points) generates only a subtly noticeable difference in perceived loudness. This suggests that the effect through this mechanism would be negligible, unless subjects exaggerate their expectations of changes in loudness both substantially and asymmetrically between conditions.

6 Composition of Strategy Types and Equilibrium Information Avoidance

Section 5.3 showed that individuals differ substantially in their strategy types — their type of best-response schedule to others' information decisions. This heterogeneity suggests that the composition of strategy types is an important determinant of the *equilibrium* prevalence of information avoidance in a group. We refer to the equilibrium prevalence of information avoidance as an "information equilibrium". Formally, an information equilibrium is a share π^* of information avoiders such that a proportion π^* of group members best-responds to it by avoiding information, while the remaining proportion best-responds by acquiring it. In this section, we explore how different strategy type compositions map into the set of information equilibria.

To derive the set of information equilibria of a group g, we first construct the group-level, or "aggregate", best response function by aggregating the individual best-response functions of its members, as follows. Each individual best-response schedule is a function $f_{gi} : [0, 1] \rightarrow \{0, 1\}$ that maps each possible share of avoiders to a decision to avoid (1) or acquire (0) information. To derive the aggregate best response schedule F_{gi}, we simply take the average of f_{gi} at each π, over group members. For example, consider a simple group composed by one Always Getter (with $f_{g1}(\pi) = 0$ for all π) and one Always Avoider (with $f_{g2}(\pi) = 1$ for all π). Then the aggregate best response function is $F_g(\pi) = 0.5$ for all π. In other words, for any proportion of avoiders π, half of the group members best-responds by avoiding information. Since, in practice, f_{gi} can be elicited only for a coarse set of values of π, we can accordingly obtain only a coarse representation of the aggregate schedule F_g. To estimate its full shape, we impute the missing values by linear interpolation. This approach is simple, provides good local approximations around the observed points of F_g (under the mild assumption that F_g is continuous), and ensures that the estimated schedule varies smoothly with π.

Once the full shape of F_g is estimated, identifying the information equilibria is straightforward: it is the set of fixed points of F_g, namely, the set of values π^* such that $F_g(\pi^*) = \pi^*$. Graphically, these are the points where F_g crosses the identity line.

To build intuition, we begin by considering "homogeneous" groups, namely, groups composed entirely by a single strategy type. Figure 10 plots the best-response function F_g of groups composed by the subsample defined by each strategy type. Blue dots indicate the observed values of F_g, while blue lines the segments estimated by linear interpolation. The set of equilibria, shown as orange diamonds, are given by the intersection(s) of F_g and the identity line (dotted). Homogeneous groups of unconditional types (Always Getters or Always Avoiders) display a flat aggregate schedule, since the best-responses of all its members are unaffected by others' decisions. Those groups have a unique information equilibrium at either extreme, where everybody is informed (Always Getters) or uninformed (Always Avoiders). In contrast, homogeneous groups of conditional types generate monotonic F_g: negatively sloped for Strategic Substitutes, and positively for Strategic Complements.[33] Substitutes feature a unique interior equilibrium. Complements, however, exhibit multiple equilibria —including fully informed (MAA), fully uninformed (MAI), and an interior equilbrium—, consistent with the theoretical prediction of Bénabou (2013). This possibility of multiple of equilibria is important, as it captures the risk of collective ignorance, even when full awareness is also sustainable. The plot therefore illustrates the theoretical insights discussed in Section 3.1: Complements can sustain both the MAA and the MAI equilibrium, while Substitutes can reach none of those.

The diversity of shapes in aggregate schedules F_g suggests that information equilibria may vary greatly with group composition, motivating a systematic analysis of how variations in the mixture of types influences equilibrium outcomes. To study this, we employ simulations to characterize the distribution of equilibria under different compositions of strategy types: for each composition, we randomly sample groups, compute their equilibrium sets, and trace how the resulting distribution of equilibria shifts as composition changes.

We detail the algorithm as follows. Fix a group size N. Fix a sequence of group compositions $\tau \equiv \{\tau_t\}_{t=1,2,...,T}$, where each group composition τ_t is a vector $\tau_t \equiv (p_{AG}, p_{AA}, p_{SC}, p_{SS})_t$ whose elements are probabilities (respectively, for Always Getters, Always Avoiders, Complements, and Substitutes) that add up to 1. Start with $t = 1$, and the first composition τ_1. For each of $b = 1, ..., B$ bootstrap iterations, draw a random group of size N through bootstrap sampling with replacement, where the sampling probabilities of each strategy type are given by τ_1. This generates B simulated groups of size N with composition τ_1. For each group $b = 1, ..., B$, construct the aggregate best

[33] The precise piece-wise slopes will depend on the distribution of individual switching thresholds.

Figure 10: Aggregate best-response functions and information equilibria (homogeneous groups)

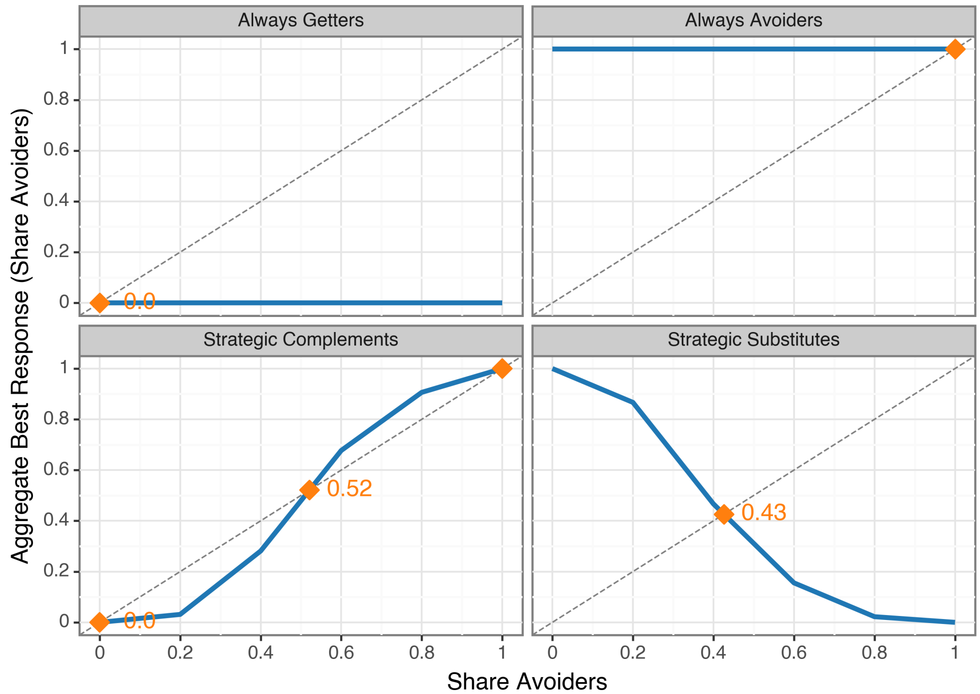

Notes: The figure shows the aggregate (group-level) best response schedules of "homogeneous" groups composed by a single strategy type. Each group is formed by subsetting the experimental sample on the type indicated by the caption. Blue dots represent values of the schedule at elicited values of π. Blue lines show the linear interpolation between blue points. The dotted, diagonal line shows the identity line. Orange diamonds indicate the fixed points of the schedule (i.e., the information equilibria).

response F_b (including the linear interpolation) from the group members' individual schedules and find the set of fixed points of F_b, Π_b, classifying their stability. Then, obtain composition-level statistics (computed across bootstrapped groups b): the mean and 95% confidence interval of F_b at each π, and the distribution of information equilibria π^*. Let $t := t + 1$ and iterate for group composition τ_t until $t = T$.

In the simulations, we iterate over a grid of all possible compositions where the proportion of types take on discrete values from 0 to 1 in steps of 0.1. Given the constraint that the probabilities add up to 1, the grid includes 286 different group compositions. For each composition, we bootstrap $B = 100$ groups, each with $N = 13$ members.[34] While studying large N is possible and would result in more precise estimates, as N grows it becomes more unrealistic for a principal to be able

[34]If N is a "round" number, such as 10, it is possible that, for some draws, continuous portions of F_b lie exactly on top of the identity line, leading to an infinite number of equilibrium points. This introduces technical details about how to count and adjust for these when calculating the relative frequency of equilibria, as well as apparent non-monotonicities in features of the distribution of equilibria as the group composition τ_t varies. To avoid these issues, we focus the discussion on cases constructed to prevent the emergence of infinite number of equilibria.

to control the group composition. We therefore focus on finite group sizes, where composition remains as a policy-relevant lever. The simulation thus involves 367,600 sampled individuals across 28,600 groups, across 286 group compositions.

Among the large number of possible group compositions, we focus on "restricted" sequences that begin with a group composed entirely of a conditional type and progressively add unconditional types.[35] Formally, these satisfy $p_C + p_U = 1$, where $C \in \{SC, SS\}$ denotes a conditional type and $U \in \{AG, AA\}$ an unconditional type, and where p_U progressively grows from 0 to 1.

In Appendix Figure C.1, we show how the shape of the aggregate best-response function F_g evolves as the group composition of types changes. Intuitively, starting with homogeneous groups of conditional types —whose schedule is either positively or negatively sloped—, a gradual increase in the participation of unconditional types progressively flattens the schedule and pulls it towards the extremes.

Figure 11 shows the resulting distribution of information equilibria by group composition of strategy types. Each panel starts with a conditional type (Strategic Complements in Panels A and C; Strategic Substitutes in Panels B and D) and gradually increases the participation of an unconditional type (Always Avoiders in Panels A and B; Always Getters in Panels C and D). Groups consisting entirely of Strategic Complements exhibit a wide range of equilibria. All groups are characterized by a full avoidance and a full acquisition equilibrium (both of which are typically stable). In addition, there is an interior equilibrium (typically unstable). As Complements are progressively mixed with Always Avoiders (Panel A), the distribution shifts to the right, so that full-avoidance remains as an equilibrium, while the full-acquisition equilibrium gradually disappears. For Strategic Substitutes (Panels B and D), the picture is very different. Homogeneous groups of Substitutes ($p_{SS} = 1$) feature a single, typically unstable interior equilibrium. As expected, introducing Always Avoiders shifts this equilibrium rightward (toward more avoiders in equilibrium), while introducing Always Getters shifts it leftward (toward less avoiders), so that when the participation of unconditional types is sufficiently high, only corner equilibria remain (either full avoidance or full awareness).

Since, in practice, unstable equilibria are unlikely to persist and to be observed, in what follows we restrict attention to the set of stable equilibria.

Policy-makers may be interested in targeting specific features of the distribution of equilibria. For example, they may be interested in the expected value of or in the highest possible prevalence of avoidance in a group (which one may think as a "worst-case" scenario). While the histograms in Figure 11 show the distribution of equilibria, they make it difficult to track how summary features (such as the mean, median, and other quantiles) vary with group composition. To illustrate

[35] We analyze higher-dimensional compositions in Appendix Section C.

Figure 11: Distribution of information equilibria by group composition

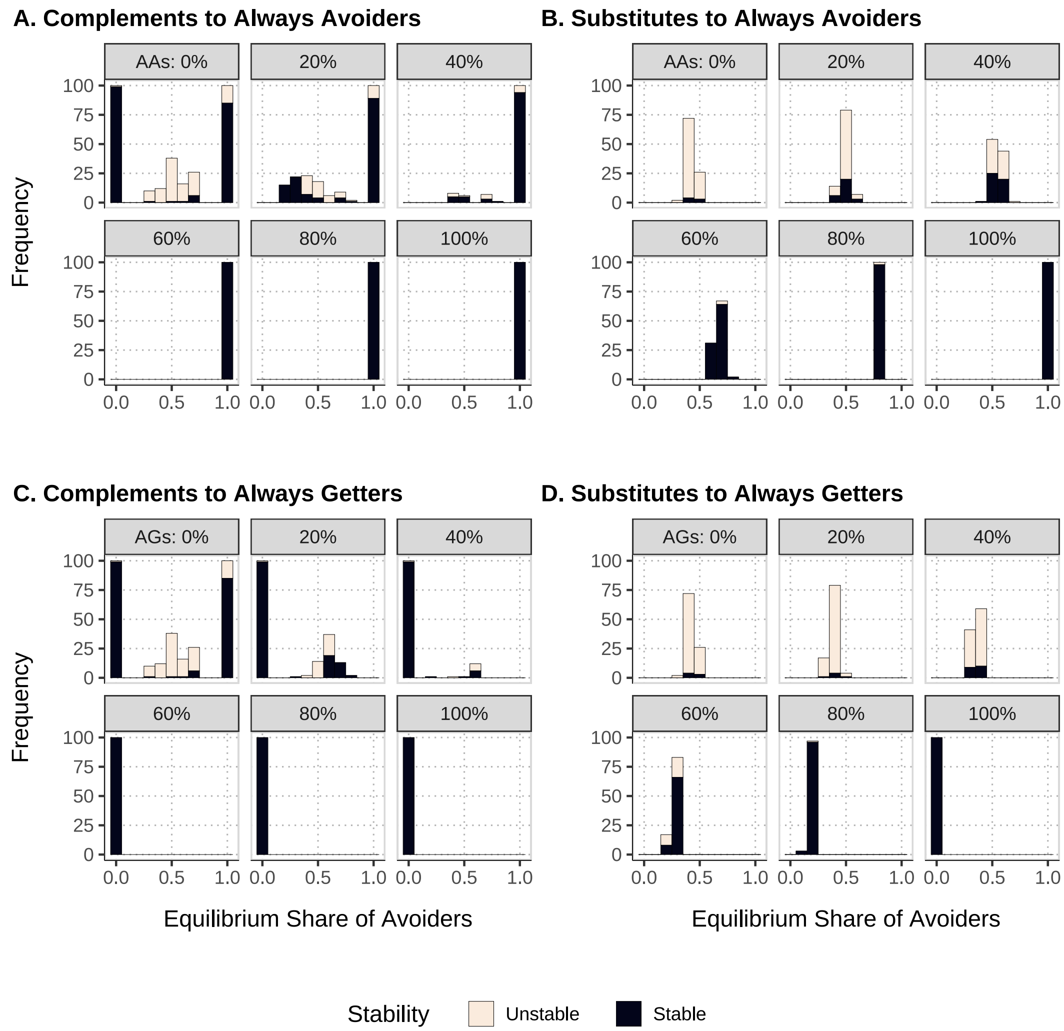

Notes: The figure shows how the distribution of information equilibria (i.e., the equilibrium prevalence of avoidance within each bootstrapped group) varies with group composition. Within each panel, each facet corresponds to a given share of unconditional types —Always Avoiders ("AA", panels A and B) or Always Getters ("AG", panels C and D)— and a remainder share of a conditional type (Strategic Complements or Substitutes). Bars report the absolute frequency (across bootstrapped groups) of the equilibrium share of avoiders indicated on the x-axis. Stable equilibria are shown in dark bars, and unstable equilibria in light bars.

these relationships more clearly, Figure 12 plots how such features change as the proportion of unconditional types increases. The summary features of the distribution of equilibria respond monotonically but nonlinearly to changes in the composition. To illustrate, Panel A shows that when the participation of Always Avoiders (AA) increases from 0% to 10% in a group conformed otherwise by Strategic Complements, the median equilibrium share of information avoiders (gray line) increases from 0% to 22%. However, when the participation increases by a similar amount from 10% to 20%, the median jumps sharply from 22% to 100%. Similar nonlinear dynamics are observed for other quantiles and for the average (orange diamonds) of the distribution of equilibria.

Overall, the figure shows that features of the distribution of information equilibria can exhibit strong, nonlinear responses to changes in strategy type composition. This implies that there are regions in the composition space (spanning the x-axis) where the effect of a fixed amount of change in group composition generates the strongest or weakest responses in the feature of interest.

Figure 12: Summary features of distribution of information equilibria, by group composition

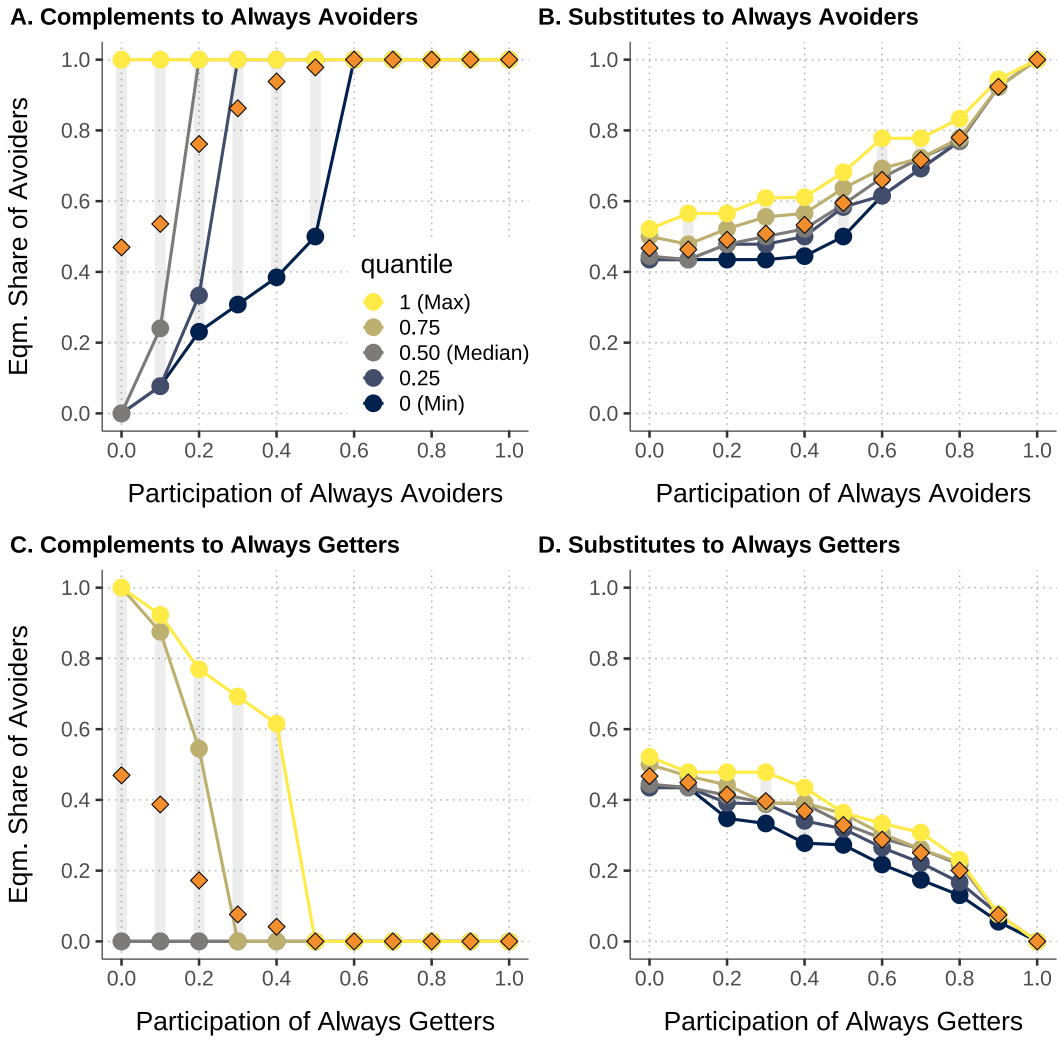

Notes: The figure shows how summary features of the distribution of information equilibria (i.e., the equilibrium prevalence of information avoidance within a group) respond to changes in the group composition of strategy types. Each panel focuses on mixtures of two strategy types: Panels A and B start with groups of Strategic Complements and gradually increase the participation of Always Avoiders (Panel A) and of Always Getters (Panel B); Panels C and D start with groups of Strategic Substitutes and gradually increase the participation of Always Avoiders (Panel C), and of Always Getters (Panel D). The x-axis represents the participation of the unconditional type, while the y-axis reports the values of the feature of interest. Gray vertical bars indicate the support of the distribution of information equilibria (over bootstrapped groups) of a given composition; orange diamonds denote the average of the distribution; and dots indicate quantiles of the distribution (maximum in yellow, median in dark gray, and minimum in dark blue).

7 Conclusion

We live in an era of abundance of information in a variety of domains —politics, health, financial, environmental. Despite the wide availability of virtually free information, sometimes individuals prefer to avoid it, choosing, in effect, to remain in a state of willful ignorance. In this paper, we study whether this behavior is interdependent among individuals in a setting where remaining ignorant imposes negative payoff externalities on other members. Our experimental evidence shows that, on average, willful ignorance can be contagious: individuals are more likely to avoid payoff-improving and freely available information when they expect others to do so, and hence, when they expect the outcome in the bad state to be more severe (and news revealing a bad state to be more aversive). Our findings suggest that the interdependence is mediated by an increase in expected anticipatory disutility of discovering bad news as a result of others' information avoidance, and it is moderated by social preferences, particularly, by reciprocity. Beyond average effects, however, there is substantial heterogeneity in individual strategy types: while we empirically find individuals who exhibit strategic complementarity and who amplify willful blindness as in Bénabou (2013), we also document the existence of an equally sizeable share of strategic substitutes, who counter the spread of ignorance. This heterogeneity of reaction functions suggests that the degree of aggregate contagiousness of willful ignorance may depend crucially on both the composition of the group in terms of reaction functions and the distribution of subjective expectations about others' behavior. Simulations show that, in effect, changes in group composition generate a wide variety of information equilibria, ranging from full information acquisition (and maximum mitigation of the bad state outcome) to full willful ignorance (and maximum severity), and that the effects can be strong and non-linear.

Organization designers and policy-makers are often interested in promoting the take-up of information and learning within their organizations and populations. Our study adds nuance to the recommendations from the literature on information preferences, which typically focus on individual decision settings. Our findings suggest that incentivizing a person to acquire information (which may be beneficial if the individual is in isolation) may generate positive or negative externalities in the informational responses other members. Therefore, policy-makers might benefit from taking into account these "cognitive externalities" and general equilibrium effects when designing policies that aim at increasing the aggregate take-up of information in groups when individuals' outcomes are interlinked.

Our results suggest some implications for organization design. Taking the group as given, a potential way to encourage information take-up is by correcting potential misperceptions about how others deal with information. Another way is to design the group composition in terms of strategy types when assembling teams. Our simulations show that the group composition has a strong impact in the distribution of equilibrium shares of informed members, and suggest the equilibrium responses can be non-linear in the mixture of types, in a way that there are regions in which the same amount of change in composition delivers the strongest responses.

Our investigation has some limitations and motivates avenues for future research. First, because our experimental setup is based on the setup of Bénabou (2013) where there is a theoretical link between agents' decisions and the severity of the outcome, our design cannot disentangle whether (and how much of) the contagiousness arises from variations in the severity of the outcome ("severity effect") or from variations in the proportion of others who engage choose to remain ignorant ("herding effect"). This limitation applies both to the average treatment effect and the individual best-response functions. We view our study as a first step in establishing the empirical existence of interdependence in information decisions. Further research is needed to disentangle the severity effect from the herding effect. Second, we consider a static decision-making setting, where all individuals choose their information decisions simultaneously. In many applications, such decisions occur sequentially, motivating future studies on willful ignorance in groups in dynamic settings.

Our study highlights that understanding the ways in which individuals react to others' information decisions is crucial for designing environments that promote the acquisition of information, which, usually, improves subsequent decision-making in organizations.

References

Beaurenaut, M., Tokarski, E., Dezecache, G., & Grèzes, J. (2020). The 'Threat of Scream' paradigm: A tool for studying sustained physiological and subjective anxiety. *Scientific Reports*, 10(1), 12496.

Belloni, A., Chernozhukov, V., & Hansen, C. (2014). Inference on Treatment Effects after Selection among High-Dimensional Controls†. *The Review of Economic Studies*, 81(2), 608–650.

Bénabou, R. (2013). Groupthink: Collective Delusions in Organizations and Markets. *The Review of Economic Studies*, 80(2), 429–462.

Dana, J., Weber, R. A., & Kuang, J. X. (2007). Exploiting moral wiggle room: Experiments demonstrating an illusory preference for fairness. *Economic Theory*, 33(1), 67–80.

Engelmann, J. B., Lebreton, M., Salem-Garcia, N. A., Schwardmann, P., & van der Weele, J. J. (2024). Anticipatory Anxiety and Wishful Thinking. *American Economic Review*, 114(4), 926–960.

Falk, A., Becker, A., Dohmen, T., Enke, B., Huffman, D., & Sunde, U. (2018). Global Evidence on Economic Preferences*. *The Quarterly Journal of Economics*, 133(4), 1645–1692.

Falk, A. & Zimmermann, F. (2024). Attention and Dread: Experimental Evidence on Preferences for Information. *Management Science*, 70(10), 7090–7100.

Fischbacher, U., Gächter, S., & Fehr, E. (2001). Are people conditionally cooperative? Evidence from a public goods experiment. *Economics Letters*, 71(3), 397–404.

Fong, J. & Hunter, M. (2022). Can Facing the Truth Improve Outcomes? Effects of Information in Consumer Finance. *Marketing Science*, 41(1), 33–50.

Galai, D. & Sade, O. (2006). The "Ostrich Effect" and the Relationship between the Liquidity and the Yields of Financial Assets. *The Journal of Business*, 79(5), 2741–2759.

Ganguly, A. & Tasoff, J. (2017). Fantasy and Dread: The Demand for Information and the Consumption Utility of the Future. *Management Science*, 63(12), 4037–4060.

Giamattei, M., Yahosseini, K. S., Gächter, S., & Molleman, L. (2020). LIONESS Lab: A free web-based platform for conducting interactive experiments online. *Journal of the Economic Science Association*, 6(1), 95–111.

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

Haaland, I., Roth, C., & Wohlfart, J. (2023). Designing Information Provision Experiments. *Journal of Economic Literature*, 61(1), 3–40.

Hertwig, R. & Engel, C. (2016). Homo Ignorans: Deliberately Choosing Not to Know. *Perspectives on Psychological Science*, 11(3), 359–372.

Ho, E. H., Hagmann, D., & Loewenstein, G. (2021). Measuring Information Preferences. *Management Science*, 67(1), 126–145.

Kahneman, D. & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2), 263–291.

Karlsson, N., Loewenstein, G., & Seppi, D. (2009). The ostrich effect: Selective attention to information. *Journal of Risk and Uncertainty*, 38(2), 95–115.

Kreps, D. M. & Porteus, E. L. (1978). Temporal Resolution of Uncertainty and Dynamic Choice Theory. *Econometrica*, 46(1), 185.

Kumar, S., Forster, H., Bailey, P., & Griffiths, T. (2009). Mapping unpleasantness of sounds to their auditory representation. *The Journal of the Acoustical Society of America*, 124, 3810–7.

Masatlioglu, Y., Orhun, A. Y., & Raymond, C. (2023). Intrinsic Information Preferences and Skewness. 113(10), 2615–2644.

Meissner, T. & Pfeiffer, P. (2022). Measuring preferences over the temporal resolution of consumption uncertainty. *Journal of Economic Theory*, 200, 105379.

Momsen, K. & Ohndorf, M. (2023). Information avoidance: Self-image concerns, inattention, and ideology. *Journal of Economic Behavior & Organization*, 211, 386–400.

Olafsson, A. & Pagel, M. (2025). The Ostrich in Us: Selective Attention to Personal Finances. *The Review of Economics and Statistics*, (pp. 1–26).

Oprea, R. & Yuksel, S. (2022). Social Exchange of Motivated Beliefs. *Journal of the European Economic Association*, 20(2), 667–699.

Oster, E., Shoulson, I., & Dorsey, E. R. (2013). Optimal Expectations and Limited Medical Testing: Evidence from Huntington Disease. *American Economic Review*, 103(2), 804–830.

Pagel, M. (2018). A News-Utility Theory for Inattention and Delegation in Portfolio Choice. *Econometrica*, 86(2), 491–522.

Selten, R. (1967). Die Strategiemethode zur Erforschung des eingeschränkt rationalen Verhaltens im Rahmen eines Oligopolexperimentes. *Beiträge zur experimentellen Wirtschaftsforschung*.

Simonsohn, U., Simmons, J. P., & Nelson, L. D. (2020). Specification curve analysis. *Nature Human Behaviour*, 4(11), 1208–1214.

Suvorov, A., van de Ven, J., & Villeval, M. C. (2024). Selective Information Sharing and Group Delusion.

Thornton, R. L. (2008). The Demand for, and Impact of, Learning HIV Status. *American Economic Review*, 98(5), 1829–1863.

Tversky, A. & Kahneman, D. (1992). Advances in Prospect Theory: Cumulative Representation of Uncertainty. *Journal of Risk and Uncertainty*, 5(4), 297–323.

A Theoretical Appendix

Proof of Lemma 1 Consider v_i everywhere convex. Agent i prefers to avoid information if $\varphi_i > 0$, i.e.

$$v_i\left(pb + \left(1-p\right)g\right) > pv_i\left(b+\Delta\right) + \left(1-p\right)v_i\left(g\right)$$

Since $v_i\left(.\right)$ is strictly increasing, we have $v_i\left(b+\Delta\right) > v_i\left(b\right)$, hence we must have

$$v_i\left(pb + \left(1-p\right)g\right) > pv_i\left(b\right) + \left(1-p\right)v\left(g\right)$$

but this contradicts convexity. Therefore, the agent never avoids information, i.e., the agent is an Always Getter. ■

Proof of Lemma 2 Suppose v_i everywhere concave. First note that the net incentive to avoid information increases in the share of other group members who avoid information:

$$\begin{aligned}\frac{\partial\varphi_i}{\partial\lambda_{-i}} &= pv_i'\left(pb + \left(1-p\right)g\right)\frac{\partial b}{\partial\lambda_{-i}} - pv_i'\left(b+\Delta\right)\frac{\partial b}{\partial\lambda_{-i}}\\ &= \left[v_i'\left(pb + \left(1-p\right)g\right) - v_i'\left(b+\Delta\right)\right]p\frac{\partial b}{\partial\lambda_{-i}}\\ &> 0\end{aligned}$$

by concavity of v_i and assumption 1. It follows that:

- If $\varphi_i\left(\lambda_{-i}=0\right) > 0$ then by monotonicity $\varphi_i\left(\lambda_{-i}\right) > 0$ for any λ_{-i}, that is, the agent is an Always Avoider.
- If $\varphi_i\left(1\right) < 0$ then by monotonicity $\varphi_i\left(\lambda_{-i}\right) < 0$ for any λ_{-i}, that is, the agent is an Always Getter.
- If $\varphi_i\left(0\right) < 0$ and $\varphi_i\left(1\right) > 0$, then by continuity there exists an interior λ_{-i}^* such that the agent acquires information whenever $\lambda_{-i} < \lambda_{-i}^*$ and avoids it when $\lambda_{-i} > \lambda_{-i}^*$ (strategic complementarity). ■

Proof of Lemma 3 We assume throughout that $b(0)+\Delta < \widehat{U}_i$. Suppose that if everyone else avoids information, then at date 1 the expected utility under ignorance is so low that it falls below the reference level: $pb\left(1\right) + \left(1-p\right)g < \widehat{U}_i$. Then v_i is convex at $pb\left(1\right) + \left(1-p\right)g$. By convexity,

$$v_i\left(pb\left(1\right) + \left(1-p\right)g\right) < pv_i\left(b\left(1\right)\right) + \left(1-p\right)v_i\left(g\right)$$

so even without instrumental value of information Δ, avoiding information is dominated by acquiring it. Since v_i increasing, adding the instrumental value makes avoiding information even less desirable:

$$v_i\left(pb\left(1\right)+\left(1-p\right)g\right) < pv_i\left(b\left(1\right)\right)+\left(1-p\right)v_i\left(g\right) < pv_i\left(b\left(1\right)+\Delta\right)+\left(1-p\right)v_i\left(g\right)$$

Therefore, $\varphi^i\left(1\right) < 0$. When all other agents acquire information, agent i prefers to avoid. Suppose that if everyone else acquires information, then at date 1 the expected utility under ignorance goes above the reference level: $pb\left(0\right)+\left(1-p\right)g > \widehat{U}_i$. Then v_i is concave at $pb\left(0\right)+\left(1-p\right)g$. By concavity,

$$v_i\left(pb\left(0\right)+\left(1-p\right)g\right) > pv_i\left(b\left(0\right)\right)+\left(1-p\right)v_i\left(g\right)$$

For avoidance to be optimal, we need

$$v_i\left(pb\left(0\right)+\left(1-p\right)g\right) > pv_i\left(b\left(0\right)+\Delta\right)+\left(1-p\right)v_i\left(g\right)$$

Define the difference in expected utilities $d \equiv v_i\left(g\right)-v_i\left(pb\left(0\right)+\left(1-p\right)g\right)$. Then, the avoidance condition can be re-expressed as:

$$\begin{aligned} v_i\left(g\right)-d &> pv_i\left(b\left(0\right)+\Delta\right)+\left(1-p\right)v_i\left(g\right) \\ \text{i.e. } p\left[v_i\left(g\right)-v_i\left(b\left(0\right)+\Delta\right)\right] &> d \end{aligned}$$

When $d \to 0$, the condition above always holds. Therefore, by continuity, there exists a value $d^* > 0$ such that for $d < d^*$ the agent avoids information. If $d < d^*$, the agent avoids information when all other agents in the group acquire it ($\varphi_i\left(\lambda_{-i}=0\right) > 0$) and acquires information when all others avoid it ($\varphi_i\left(\lambda_{-i}=1\right) < 0$). By continuity of $\varphi_i\left(\lambda_{-i}\right)$, there exists an interior threshold λ^{**} such that agent i strictly prefers to avoid information if $\lambda_{-i} < \lambda_i^{**}$ and strictly prefers to acquire information if $\lambda_{-i} > \lambda_i^{**}$ (strategic substitutability). ■

B Additional analyses

All subsequent analyses are referenced in the document. Brief descriptions are provided for each analysis.

B.1 Descriptive statistics

Descriptive statistics of the experimental sample are reported below.

	N	Mean	SD	Median	Min	Max
Demographic characteristics						
Age	460	22.509	3.713	22.000	18.000	47.000
Gender (Female=1)	465	0.583	0.494	1.000	0.000	1.000
Has work experience	460	0.154	0.362	0.000	0.000	1.000
General information preferences						
IPS (overall scale)	464	3.116	0.414	3.111	1.000	4.000
Social and economic preferences						
Patience	464	-0.000	1.000	0.201	-3.493	1.124
Risk seeking	464	-0.000	1.000	0.296	-2.678	1.570
Positive reciprocity	464	-0.000	0.803	-0.017	-3.911	1.244
Negative reciprocity	464	0.000	0.836	0.117	-1.833	1.946
Altruism	464	0.000	0.836	0.096	-2.163	1.947
Trust	464	-0.000	1.000	-0.178	-1.662	2.048
Field of study						
Medicine	465	0.213	0.410	0.000	0.000	1.000
STEM	465	0.200	0.400	0.000	0.000	1.000
Humanities	465	0.228	0.420	0.000	0.000	1.000
Economics	465	0.108	0.310	0.000	0.000	1.000
Business	465	0.211	0.408	0.000	0.000	1.000
Other fields	465	0.056	0.230	0.000	0.000	1.000

Notes: The table shows descriptive statistics. The number of observations is lower for some variables elicited at the end of the session (e.g., age) because a few participants did not respond. For other variables (e.g., the IPS score and social preferences), one observation is missing due to a participant leaving the session early to attend a lecture.

Table B.1: Sample summary statistics

B.2 Are screams a bad? Preferences over screams and volume

We classify “Volume Haters” (resp., “Lovers”) as subjects who report higher (lower) utility from hearing screams at volume 50 than at volume 100. Similarly, we classify “Scream Haters” (“Lovers”) as subjects who report a higher (lower) utility from learning that the state is Quiet than that it is Screams. To capture the intensity of these preferences, we compute for each individual the difference between the utility of volume 50 and volume 100, and the difference between the utility of learning that the state is Quiet state relative to Screams.

Figure B.1 shows the joint distribution of subjects' reported utilities of hearing screams at different volumes, distinguishing each type of volume preference. In our sample, 88% of subjects are Volume Haters, 3% are Volume Indifferent, and 9% are Volume Lovers.

Figure B.2 presents the joint distribution of subjects' reported utilities of learning each state, distinguishing each type of scream preference. In our sample, 81% of subjects are Scream Haters, 6% are Scream Indifferent, and 14% are Scream Lovers.

Figure B.1: Preferences over volume of screams: reported utilities and preference types

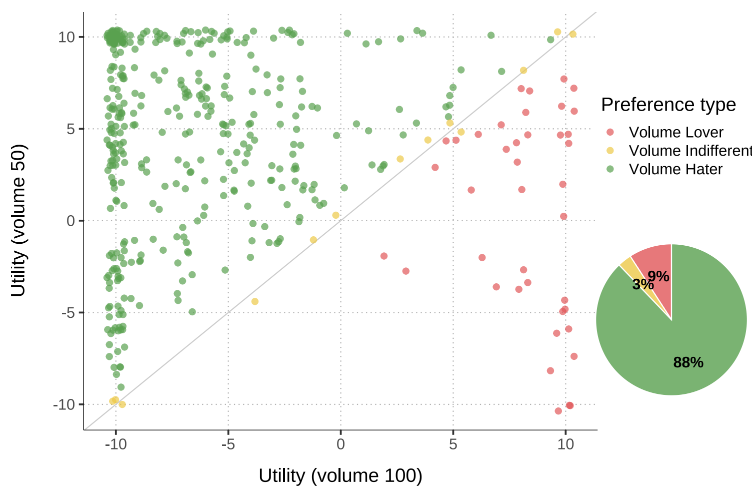

Notes: The figure displays reported utilities when screams were played at different volumes. Each point (jittered) corresponds to a participant's reported utilities for volume 50 (y-axis) and volume 100 (x-axis). The gray 45° line indicates indifference between volume levels. Observations are classified into three preference types: Volume Haters (green), who prefer quieter screams; Volume Indifferents (yellow), who report no difference; and Volume Lovers (red), who prefer louder screams. The pie chart summarizes the prevalence of each preference type in the sample.

The joint distribution of preference types over screams and volume is shown in Table B.2.

B.3 Balance checks

Table B.3 reports the treatment balance checks.

Figure B.2: Preferences over screams: reported utilities upon learning each state

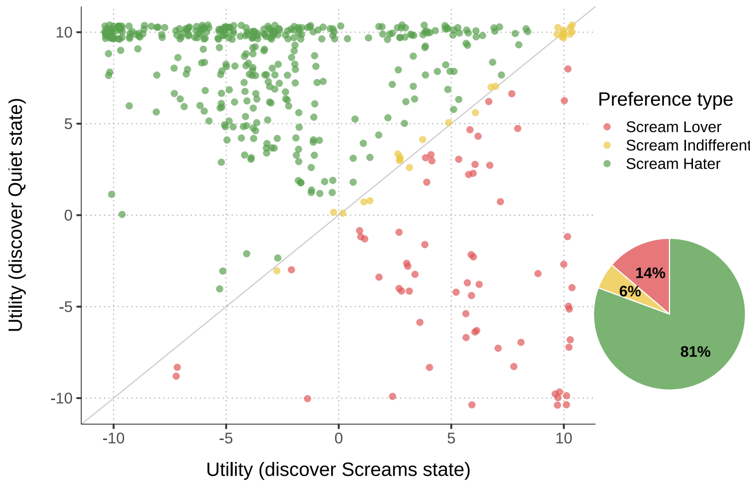

Notes: The figure displays reported expected utilities upon learning the state of the world. Each point (jittered) corresponds to a participant's reported utilities for learning the Quiet state (y-axis) and the Screams state (x-axis). The gray 45° line indicates indifference across states. Observations are classified into three preference types: Scream Haters (green), who prefer to discover the Quiet state; Scream Indifferents (yellow), who report no difference; and Scream Lovers (red), who prefer to discover the Scream state. The pie chart summarizes the prevalence of each preference type in the sample.

	Volume Preference Type				
Scream Preference Type	**Volume Lover**	**Volume Indif-ferent**	**Volume Hater**		**Total**
Scream Lover	23	3	34	4	64
Scream Indifferent	1	5	19	1	26
Scream Hater	16	5	330	24	375
Total	40	13	383	29	465
Scream Lover	4.9%	0.6%	7.3%	0.9%	13.8%
Scream Indifferent	0.2%	1.1%	4.1%	0.2%	5.6%
Scream Hater	3.4%	1.1%	71.0%	5.2%	80.6%
Total	8.6%	2.8%	82.4%	6.2%	100.0%

Table B.2: Distribution of Scream- and Volume-preference types

	(1) Few Avoided	(2) Many Avoided	(3) Difference (2)-(1)
Demographic characteristics			
Age	22.53	22.49	-0.04
	(3.98)	(3.46)	(0.35)
Gender (Female=1)	0.57	0.60	0.03
	(0.50)	(0.49)	(0.05)
Information preferences			
IPS, overall scale	3.10	3.13	0.02
	(0.42)	(0.41)	(0.04)
Health subscale	3.04	3.08	0.04
	(0.72)	(0.68)	(0.06)
Finance	2.88	2.90	0.02
	(0.67)	(0.63)	(0.06)
Personal	3.11	3.14	0.03
	(0.55)	(0.55)	(0.05)
General	3.04	3.08	0.03
	(0.70)	(0.70)	(0.06)
Occupational	3.29	3.30	0.01
	(0.48)	(0.50)	(0.05)
Field of study			
Medicine	0.19	0.23	0.04
	(0.40)	(0.42)	(0.04)
STEM	0.18	0.21	0.03
	(0.39)	(0.41)	(0.04)
Humanities	0.19	0.26	0.07*
	(0.40)	(0.44)	(0.04)
Economics	0.14	0.08	-0.06**
	(0.35)	(0.27)	(0.03)
Business	0.25	0.18	-0.07*
	(0.43)	(0.38)	(0.04)
Other fields	0.06	0.05	-0.01
	(0.24)	(0.22)	(0.02)
Work experience			
Has work experience	0.16	0.15	-0.01
	(0.37)	(0.36)	(0.03)
Observations	223	242	465

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Table B.3: Balance checks

B.4 Treatment effect on belief about prevalence of avoidance

Table B.4 reports estimates from linear regressions of beliefs on the treatment dummy and different sets of controls. Column (1) reproduces the results in Figure 4: on average, the expected share of avoiders among other group members is 21.9 percentage points higher in *Many Avoided* than in *Few Avoided*. This difference is highly significant and robust to the inclusion of the following controls: strata fixed effects (column 2), field of study fixed effects and individual characteristics (age and gender; column 3), and economic and social preferences (altruism, reciprocity, risk-seeking, patience, trust), as well as the IPS (column 4). Overall, the evidence confirms that the treatment successfully shifted beliefs about the prevalence of avoidance among other members.

	Dep. var: Belief about share of avoiders			
	(1)	(2)	(3)	(4)
Many Avoided	0.219***	0.215***	0.218***	0.221***
	(0.021)	(0.025)	(0.025)	(0.025)
Control group mean	0.371***	0.471***	0.394***	0.494***
	(0.016)	(0.072)	(0.137)	(0.163)
Strata FEs		✓	✓	✓
Field of study FEs			✓	✓
Individual characteristics			✓	✓
Preferences (economic, social, IPS)				✓
Observations	465	465	460	460

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Notes: The table shows the estimated coefficients of a regression of beliefs about the share of information avoiders among other members on the treatment dummy. The dependent variable is a continuous measure of beliefs, ranging from 0 (0% of the subject's group will avoid information") to 1 (100%). The treatment dummy is equal to one if the subject was in the *Many Avoided* condition and zero if in the *Few Avoided* condition. Individual characteristics include age and gender. "Preferences" include economic and social preferences (risk-seeking, time preferences, altruism, reciprocity, and trust) and the Information Preferences Scale (Ho et al., 2021), measuring an individual's general preference for obtaining information. Field of study fixed effects include dummies for fields of study. Strata-fixed effects include dummies for week of year and time of day. The number of observations drops in columns (3) and (4) because five participants did not provide information about their age.

Table B.4: Manipulation check (regressions)

B.5 Robustness to exclusion of Scream- and Volume-Lovers

Here, we show that the main regression results are robust to the exclusion of Scream- and Volume-Lovers. Figure B.3 shows the treatment effect on information avoidance by type of preference for scream (panel A) and for volume (panel B). The treatment effect is positive for scream- and volume-haters, but the direction of the effect is, if anything, the opposite for scream- and volume-lovers. This suggests that, by not excluding scream- and volume-lovers from our main regressions, the treatment effect is rather attenuated.

Figure B.3: Treatment effect on information avoidance by preference over screams and volume

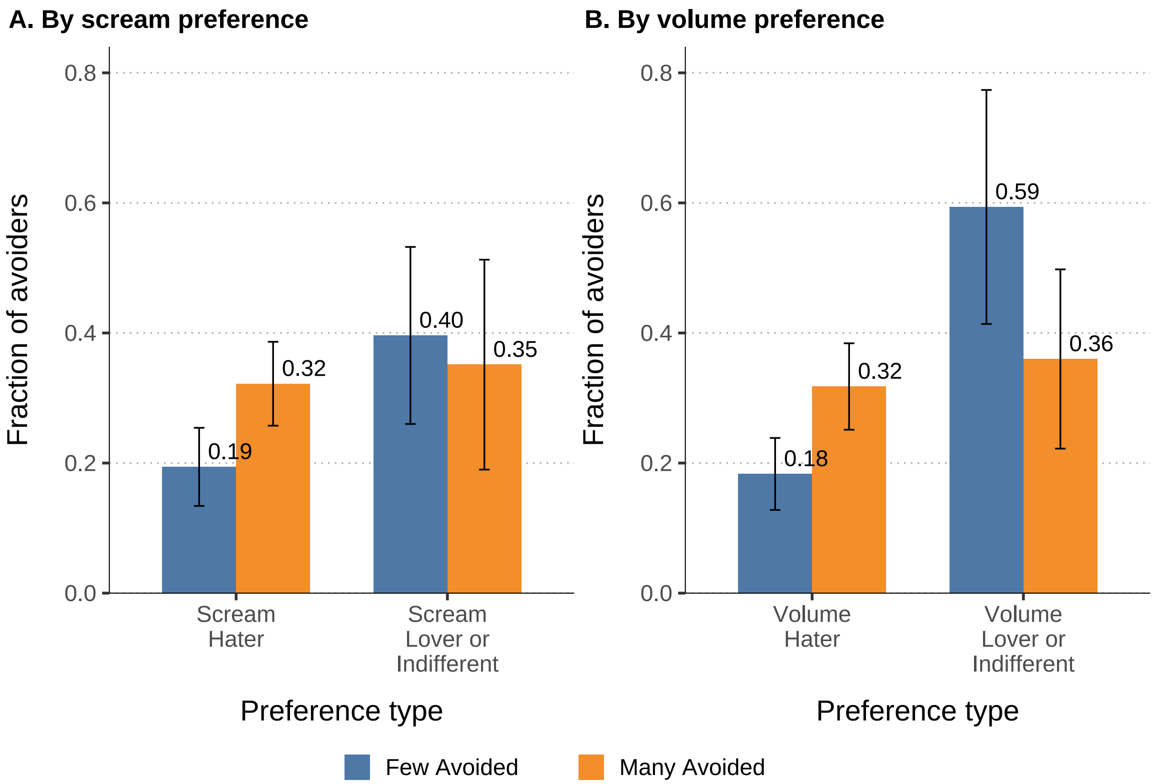

Notes: The figure shows treatment effects on information avoidance (chose Later) conditional on individual preference types. Panel A splits participants by scream preference: Scream-Haters versus Scream-Lovers or -Indifferents. Panel B splits by volume preference: Volume-Haters versus Volume-Lovers or -Indifferents. For each subgroup, the bars report the fraction of avoiders under the *Few Avoided* (blue) and *Many Avoided* (orange) treatments. Error bars represent 95% confidence intervals.

Table B.5 shows our main regression results if we exclude Scream-Lovers. Table B.6 shows our main regression results if we exclude Volume-Lovers. In both subsamples, the magnitude of the estimated treatment effect increases, compared to the full sample (reported in Table 1), for all the reported specifications. The treatment coefficient becomes significant at 1% for all specifications.

	Dependent variable: Avoided information				
	(1)	(2)	(3)	(4)	(5)
Many Avoided	0.128***	0.162***	0.157***	0.169***	0.170***
	(0.045)	(0.052)	(0.052)	(0.052)	(0.053)
Information Preferences Scale			-0.126**	-0.124**	-0.114*
			(0.056)	(0.059)	(0.060)
Social preferences					
Altruism				-0.049	-0.072
				(0.031)	(0.047)
× Many Avoided					0.050
					(0.064)
Positive reciprocity				-0.030	-0.061
				(0.032)	(0.046)
× Many Avoided					0.064
					(0.066)
Negative reciprocity				0.031	0.037
				(0.028)	(0.037)
× Many Avoided					-0.008
					(0.054)
Constant	0.194***	-0.041	0.188	0.123	0.112
	(0.031)	(0.240)	(0.324)	(0.340)	(0.352)
Strata FEs		✓	✓	✓	✓
Field of study FEs		✓	✓	✓	✓
Individual characteristics			✓	✓	✓
Observations	375	375	371	371	371

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Notes: The table reports estimates from a linear probability model regressing information avoidance on the treatment dummy, restricted to the subsample of Scream-Haters. See notes in Table 1.

Table B.5: Treatment effect on information avoidance (subsample of Scream-Haters)

	Dependent variable: Avoided information				
	(1)	(2)	(3)	(4)	(5)
Many Avoided	0.134***	0.159***	0.161***	0.166***	0.164***
	(0.044)	(0.050)	(0.049)	(0.050)	(0.051)
Information Preferences Scale			-0.133**	-0.132**	-0.127**
			(0.054)	(0.057)	(0.058)
Social preferences					
Altruism				-0.062**	-0.076*
				(0.030)	(0.045)
× Many Avoided					0.033
					(0.061)
Positive reciprocity				-0.023	-0.037
				(0.030)	(0.041)
× Many Avoided					0.032
					(0.063)
Negative reciprocity				0.021	0.007
				(0.027)	(0.035)
× Many Avoided					0.024
					(0.054)
Constant	0.183***	-0.136	0.015	-0.088	-0.094
	(0.028)	(0.162)	(0.270)	(0.284)	(0.297)
Strata FEs		✓	✓	✓	✓
Field of study FEs		✓	✓	✓	✓
Individual characteristics			✓	✓	✓
Observations	383	383	380	380	380

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Notes: The table reports estimates from a linear probability model regressing information avoidance on the treatment dummy, restricted to the subsample of Volume Haters. See notes in Table 1

Table B.6: Treatment effect on information avoidance (subsample of Volume-Haters)

B.6 Robustness to alternative specifications

Specification curve analyses To assess the robustness of our main treatment effect, we conducted a specification curve analysis (Simonsohn et al., 2020) by systematically varying sample restrictions, control sets, and model specifications. Specifically, we estimated the treatment effect across all combinations of four subsamples (full sample, excluding Scream-Lovers, excluding Volume-Lovers, excluding either) and twelve covariate sets (ranging from no controls to different combinations of demographics, field-of-study fixed effects, GPS and IPS measures). The results are presented in Figure B.4. Each point in the top panel represents the estimated effect size under an alternative specification, sorted by magnitude. Black dots indicate that the estimate is statistically significant at 5% level. The lower panels display the corresponding sample, control, and model choices. The main specification is highlighted with a red square. Figure B.5 show a similar plot but at level $\alpha = 0.01$.

Across specifications, the estimated average treatment effect remains consistently positive and is statistically significant in the majority of cases. Excluding subgroups (e.g., Volume-Lovers, Scream-Lovers) or adding or removing controls (demographics, IPS, GPS scales) does not qualitatively change the result. These findings indicate that the estimated treatment effect remains robust to a broad set of alternative specifications.

Post-double selection The specification curve analyses above rely on pre-defined subsets of covariates that are either included or excluded jointly as blocks. This design tests robustness across broad groups of control variables but lacks some flexibility: for instance, social preferences may be either included or excluded as a block, without considering combinations where altruism and reciprocity enter separately. To allow for such flexibility, we employ the post-double selection procedure (Belloni et al., 2014), a machine-learning–based, data-driven method for selecting controls from a potentially large set of covariates with the goal of causal estimation of parameters.

Results are reported in Table B.7. We estimate the treatment effect on information avoidance using both the plug-in method (columns 1–4) and cross-validation (columns 5–8). All specifications include strata fixed effects, and we progressively add four sets of controls: (i) baseline controls (economic and social preferences; IPS scales, subscales, and individual questions; field of study dummies; age; gender; and work experience), (ii) experimental context and comprehension controls (group size, perceived ease of instructions, and quiz difficulty), (iii) quiz performance until the information decision (questions answered and number correct), and (iv) quiz performance at the end of the experimental session (questions answered and number correct, performance level, and earnings).

Across all specifications, the estimated treatment effect remains positive and statistically significant, with magnitudes similar to those in the main regressions. These findings confirm that the main result (higher information avoidance in the *Many Avoided* condition) is robust to alternative specifications where selection of controls is data-driven.

Figure B.4: Specification Curve Analysis (α = 0.05)

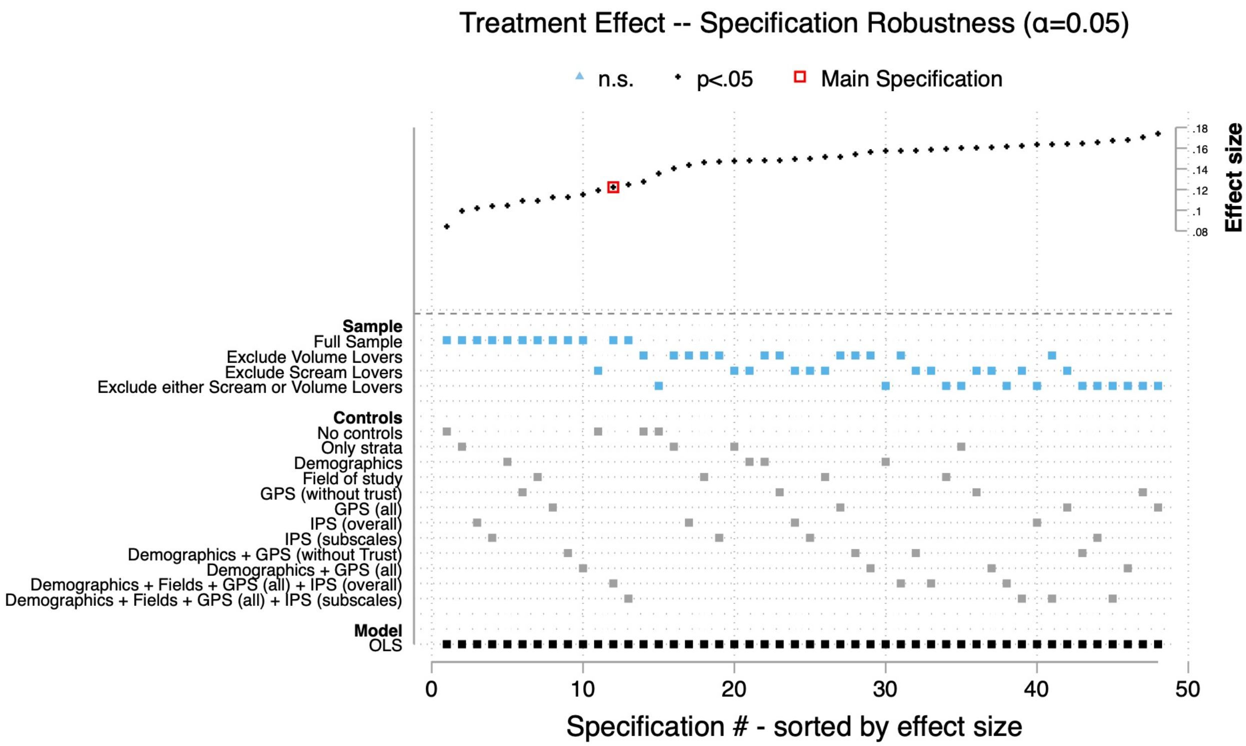

Notes: The figure reports a specification curve analysis of the treatment effect. Each point represents the estimated effect under an alternative model specification, sorted by effect size. The top panel plots estimated effect sizes, with black dots indicating specifications significant at the 5% level, blue triangles indicating non-significant estimates, and the red square denoting the main specification reported in Section 5.2. The lower panels display the sample restrictions, control sets, and model choices corresponding to each specification. Effect sizes remain consistently positive across specifications, indicating robustness of the main result.

Figure B.5: Specification Curve Analysis ($\alpha = 0.01$)

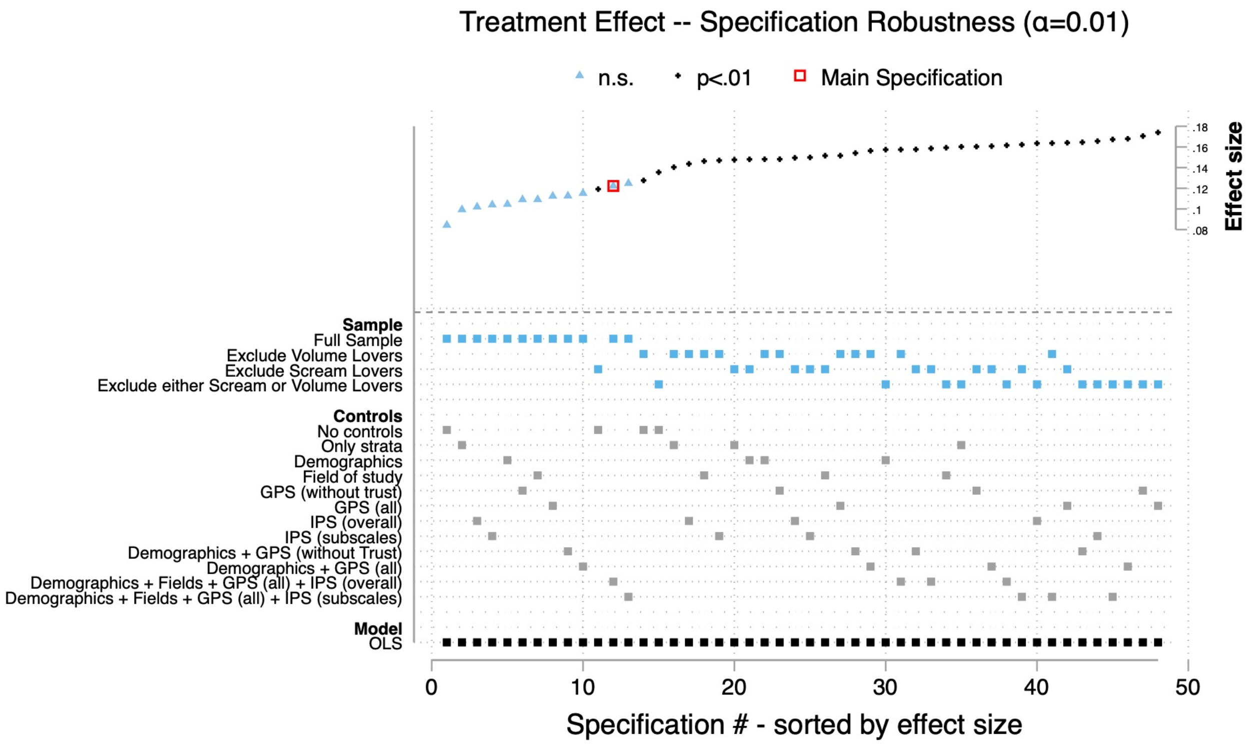

Notes: The figure reports a specification curve analysis of the treatment effect. Each point represents the estimated effect under an alternative model specification, sorted by effect size. The top panel plots estimated effect sizes, with black dots indicating specifications significant at the 1% level, blue triangles indicating non-significant estimates, and the red square denoting the main specification reported in Section 5.2. The lower panels display the sample restrictions, control sets, and model choices corresponding to each specification. Effect sizes remain consistently positive across specifications, indicating robustness of the main result.

	Method: Plug-in				Cross-Validation			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Many Avoided	0.095*	0.101*	0.101*	0.101*	0.123**	0.109*	0.110*	0.110*
	(0.050)	(0.056)	(0.056)	(0.056)	(0.049)	(0.056)	(0.056)	(0.057)
Strata FEs	✓	✓	✓	✓	✓	✓	✓	✓
Set 1: Baseline controls	✓	✓	✓	✓	✓	✓	✓	✓
Set 2: Context & comprehension		✓	✓	✓		✓	✓	✓
Set 3: Quiz performance A			✓	✓			✓	✓
Set 3: Quiz performance B				✓				✓

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Notes: The table reports post-double selection estimates of the effect of being in the Many Avoided condition on information avoidance. Columns (1)–(4) present results using the plug-in method, while columns (5)–(8) use cross-validation. All specifications include strata fixed effects and progressively add four sets of controls to the covariate selection algorithm: (i) baseline controls (age, gender, field of study dummies, social and economic preferences, Information Preferences Scale (IPS), and work experience), (ii) context and comprehension controls (group size, perceived ease of instructions, and quiz difficulty), (iii) quiz performance before making information decision (number of questions answered and number correct), and (iv) quiz performance at the end of experiment (number of questions answered and number correct, performance level, and earnings).

Table B.7: Post-double selection estimates

B.7 Alternative outcome: continuous measure of avoidance (IA Scale)

As a robustness check, we replicate our main analysis using the continuous measure of information avoidance (the "IA Scale") rather than the binary indicator. Recall that the IA Scale captures the strength of preference for the chosen information decision, ranging from –10 (strong preference for acquiring information) to +10 (strong preference for avoiding).

Table B.8 reports the results. The overall pattern closely mirrors that of the main regressions with the binary outcome (Table 1): signs of coefficients and statistical significance remain unchanged. Across specifications, the treatment effect of *Many Avoided* remains positive and significant. The coefficient is stable with the inclusion of controls and fixed effects, indicating that our main finding —–that information avoidance increases under *Many Avoided*—– remains robust to the choice of outcome measure.

Moreover, the results with IA Scale also echo the findings on social preferences documented earlier: altruism significantly predicts lower avoidance, while reciprocity remains directionally similar as before, but without though reaching statistical significance. The IPS remains negatively and significantly associated with strength of preference for avoiding information.

	Dependent variable: IA Scale				
	(1)	(2)	(3)	(4)	(5)
Many Avoided	1.271**	1.610**	1.677**	1.846**	1.838**
	(0.624)	(0.746)	(0.740)	(0.743)	(0.753)
Information Preferences Scale			-1.441*	-1.487*	-1.389*
			(0.772)	(0.807)	(0.834)
Social preferences					
Altruism				-1.094**	-1.168*
				(0.428)	(0.638)
× Many Avoided					0.291
					(0.873)
Positive reciprocity				-0.326	-0.855
				(0.444)	(0.620)
× Many Avoided					1.098
					(0.937)
Negative reciprocity				0.224	0.120
				(0.405)	(0.586)
× Many Avoided					0.181
					(0.820)
Constant	-3.821***	-4.525	-2.052	-2.815	-2.761
	(0.431)	(3.463)	(4.706)	(4.889)	(5.065)
Strata FEs		✓	✓	✓	✓
Field of study FEs		✓	✓	✓	✓
Individual characteristics			✓	✓	✓
Observations	465	465	460	460	460

* p < 0.1, ** p < 0.05, *** p < 0.01. Robust standard errors in parentheses.

Notes: The table reports estimates from a linear regression of the IA Scale on the treatment dummy. The IA Scale captures the strength of preference for the chosen information decision, ranging from –10 (strong preference for acquiring information) to +10 (strong preference for avoiding). The treatment dummy equals one if the subject was in the *Many Avoided* condition and zero in the *Few Avoided* condition. Regressions include strata fixed effects, demographic controls, field of study, and measures of economic and information preferences (see notes in Table B.4). The number of observations drops in column (3) to (5) since five participants did not provide information about their age.

Table B.8: Treatment effect on IA Scale

B.8 Instrumental variable regressions (2SLS)

Table B.9 reports instrumental variable regressions (2SLS) of information avoidance on beliefs about the share of avoiders. Beliefs are instrumented with the random treatment assignment. The model specifications (columns 1-5) match those of the main regressions (Table 1), except that the treatment dummy is replaced by the instrumented belief.

Once again, the overall pattern is very similar to that of the main regressions (Table 1). Across specifications (columns 1–5), a higher belief that others avoid information significantly increases the likelihood of avoiding information. The estimated coefficient on beliefs is stable, positive, and statistically significant across specifications (ranging from $\hat{\beta}^{2SLS} = 0.384$ to 0.604).

The coefficients associated to social preference measures remain directionally similar. Altruism is associated with lower avoidance, while positive reciprocity reduces avoidance in the Few Avoided condition but reverses under the Many Avoided condition, although coefficients lose statistical significance relative to the main regressions.

Finally, the IPS is negatively associated with the decision to avoid information.

	Dependent variable: Avoided information				
	(1)	(2)	(3)	(4)	(5)
Belief about share of avoiders	0.384**	0.504**	0.523**	0.556***	0.604**
	(0.181)	(0.213)	(0.206)	(0.201)	(0.239)
Information Preferences Scale			-0.123***	-0.122***	-0.065
			(0.045)	(0.046)	(0.060)
Social preferences					
Altruism				-0.072***	-0.076
				(0.025)	(0.144)
× Belief					0.030
					(0.314)
Pos. reciprocity				-0.032	-0.296*
				(0.028)	(0.165)
× Belief					0.635
					(0.386)
Neg. reciprocity				-0.000	-0.080
				(0.024)	(0.124)
× Belief					0.091
					(0.275)
Constant	0.100	-0.130	0.117	0.006	-0.209
	(0.088)	(0.232)	(0.282)	(0.288)	(0.359)
Strata FEs		✓	✓	✓	✓
Field of study FEs		✓	✓	✓	✓
Individual characteristics			✓	✓	✓
Observations	465	465	460	460	460
First Stage F-stat	104.71	75.93	76.49	81.34	12.77

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Notes: The table reports the second stage 2SLS estimates from regressing the information avoidance indicator (equal to one if the subject chose Later and zero otherwise) on the subject's reported belief about the share of information avoiders in his group, where the belief is instrumented by the treatment assignment dummy (equal to one if the subject was in the *Many Avoided* condition and zero otherwise). Depending on the specification, listed in columns 1 to 5, regressions include strata fixed effects, demographic controls, field of study, and measures of economic and information preferences (see notes in Table B.4). The number of observations drops in column (3) to (5) since five participants did not provide information about their age.

Table B.9: 2SLS estimates: effect of beliefs on information avoidance, using treatment assignment as instrument

B.9 Treatment effect on the distribution of strategy types (extensive margin)

	Dependent variable: Subject belongs to strategy type			
	(1) Always Getter	(2) Always Avoider	(3) Complement	(4) Substitute
Many Avoided	-0.065	0.014	0.009	0.036
	(0.046)	(0.026)	(0.038)	(0.037)
Control group mean	0.457***	0.081***	0.202***	0.175***
	(0.034)	(0.018)	(0.027)	(0.026)
Observations	465	465	465	465

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Notes: The table examines whether the treatment impacts the distribution of strategy types. It reports estimated coefficients from linear regressions where each column corresponds to a different strategy type. The dependent variable is a dummy equal to 1 if the subject's best-response schedule matches the given type and 0 otherwise. The independent variable is the treatment dummy (equal to 1 for the *Many Avoided* condition and 0 for the *Few Avoided* condition).

Table B.10: Treatment effect on distribution of strategy types (extensive margin)

B.10 Mechanisms: supplementary tables

	Dependent variable: Avoided information				
	(1)	(2)	(3)	(4)	(5)
Belief about share of avoiders	0.090	0.122	0.098	0.085	0.063
	(0.173)	(0.175)	(0.176)	(0.173)	(0.169)
Pr(screams)	-0.194	-0.176	-0.151	-0.156	-0.208
	(0.150)	(0.148)	(0.148)	(0.144)	(0.145)
Belief about share of avoiders × Pr(screams)	0.852***	0.800***	0.787***	0.826***	0.909***
	(0.294)	(0.294)	(0.297)	(0.293)	(0.292)
Constant	0.111	0.136	-0.116	0.125	0.029
	(0.076)	(0.166)	(0.215)	(0.279)	(0.285)
Strata FEs		✓	✓	✓	✓
Field of study FEs			✓	✓	✓
Individual characteristics				✓	✓
Information preferences				✓	✓
Social and economic preferences					✓
Observations	465	465	465	460	460

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Notes: The table reports estimates from a linear probability model regressing information avoidance on the subject's belief about the share of information avoiders among others (i.e., who would choose Later), the subjective belief that the state is Screams ("Pr(screams)"), and their interaction. The regressions include a variety of controls; see notes in Table 1.

Table B.11: Interaction of beliefs and information avoidance

C Group composition and equilibrium — additional analyses

C.1 Best-response schedules

Figure C.1: Effect of group composition on aggregate best-response schedules

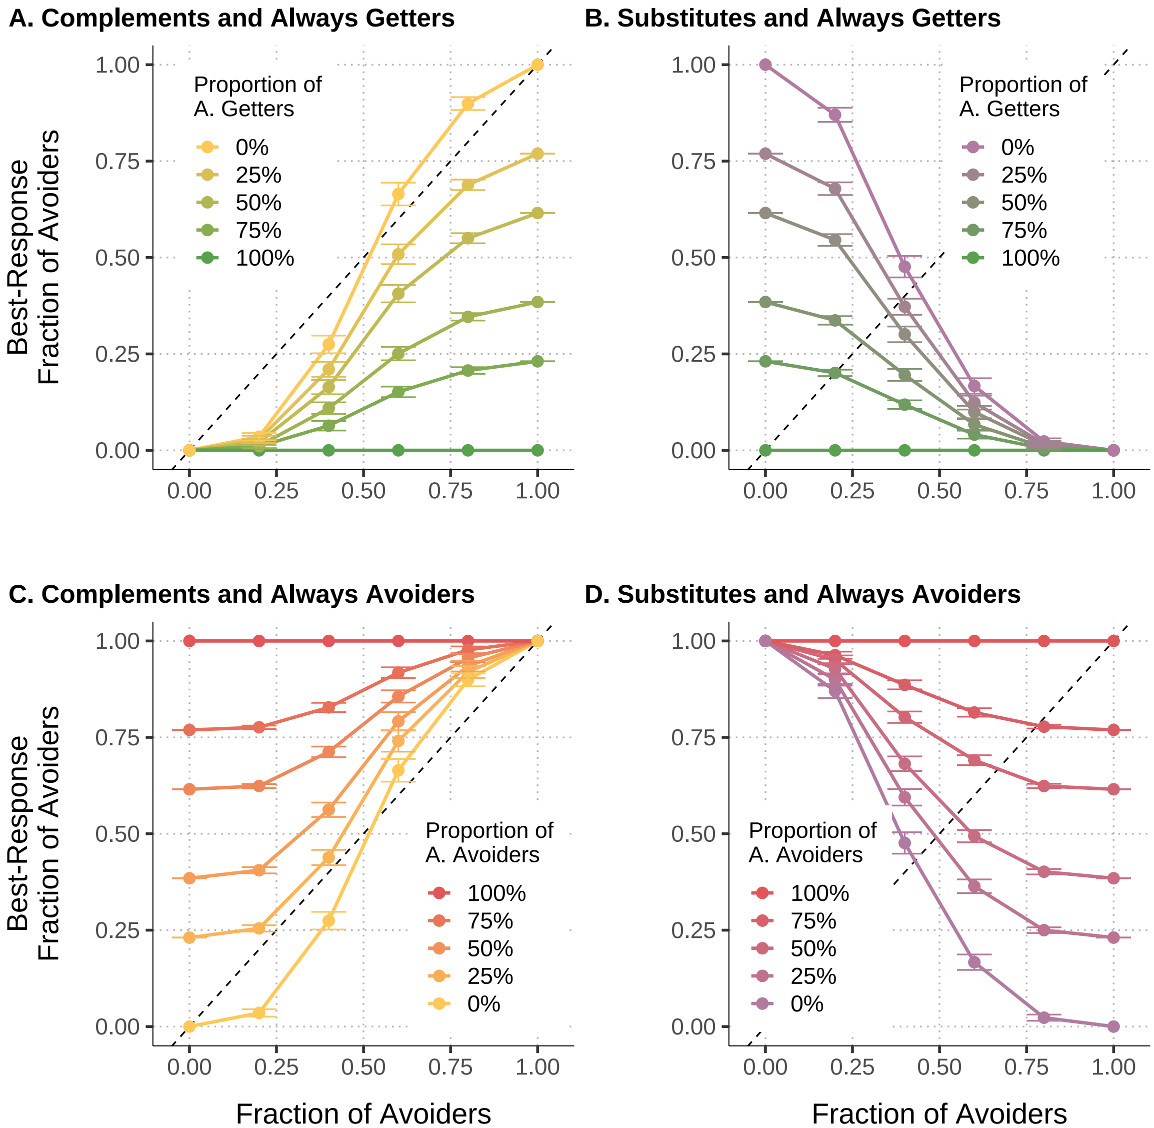

Notes: The figure shows aggregate best-response schedules by group composition. Each panel focuses on mixtures of two strategy types. Lines show the average (over bootstrapped groups) aggregate best-response schedule of a given composition. Groups composed entirely of conditional types are shown in yellow (Strategic Complements) and purple (Strategic Substitutes). Within each panel, additional lines represent increasing shares of unconditional types—Always Getters (green) or Always Avoiders (red). Error bars denote 95% confidence intervals.

C.2 Higher-dimensional group compositions

In the main analysis, we focused on "restricted" sequences of group compositions, where only two strategy types were mixed. We now turn to richer cases involving more than two types. To visualize these, it is possible to map the richer composition space onto a unit probability triangle. Since Substitutes are highly likely to produce unstable equilibria, we omit them in this analysis. This allows redefining compositions as 3-dimensional probability vectors, which can be fully represented in the triangle.

Figure C.2 illustrates this representation. The proportion of Always Getters is represented on the x-axis; and Always avoiders on the y-axis.[36] Each cell corresponds to a unique group composition, associated with a distribution of fixed points (with examples visualized in the histograms above). Panel A presents, for each cell, the average of the set of stable equilibria. A group of pure Complements (the cell at $(0, 0)$) has an average close to 0.5. As we increase the share of Always Getters (moving horizontally along the x-axis) the average equilibrium share of avoiders decreases smoothly towards zero. Conversely, increasing the share of Always Avoiders (moving upward from (0,0) along the y-axis) pushes the average smoothly toward one. Overall, across the triangle, average equilibrium outcomes respond smoothly —although nonlinearly— to group composition.

Panel B shows the maximum equilibrium in the distribution, for each cell. Intuitively, if a policy-maker regards avoidance of information as an undesirable behavior, panel B shows the "worst" equilibrium in the distribution, for each cell. For groups composed entirely of Complements, the worst equilibrium corresponds to full avoidance, and the addition of Always Avoiders does not change this. By contrast, introducing Always Getters improves the worse-case outcome: the maximum equilibrium share of avoiders declines monotonically. Notably, a strong discontinuity emerges around $p_{AG} = 0.5$. More generally, such discontinuities —where a fixed amount of change in composition produces large, abrupt shifts in the worst-case equilibrium— appear across the composition space, and these are stronger the higher the proportion of Complements. These discontinuities cluster around $p_{AG} = 0.5$: when the proportion of Always Getters is sufficiently high, the aggregate best-response function shifts downward enough to eliminate any stable interior equilibria, generating the jump in the maximum. The discontinuities weaken as the share of Always Avoiders increases, as their presence helps sustain interior equilibria.

Panel C shows the minimum among the equilibria (the "best" equilibrium in the distribution), with patterns analogous to those for the maximum. Here, the relevant probability is p_{AA}. Finally, panel D exhibits the variance of the equilibrium set, which may matter for a risk-averse policy-

[36]The proportion of Complements is implied by the complementary probability and visible as the horizonal (or vertical) distance from the cell to the hypothenuse of the triangle.

maker. Naturally, the variance is highest in groups consisting of pure Complements, where multiple equilibria exist. Adding unconditional types flattens the aggregate schedule, progressively collapsing the distribution of equilibria (as shown in Figure 11) and hence reducing variance.

Figure C.2: Group composition and summary features of equilibrium distributions

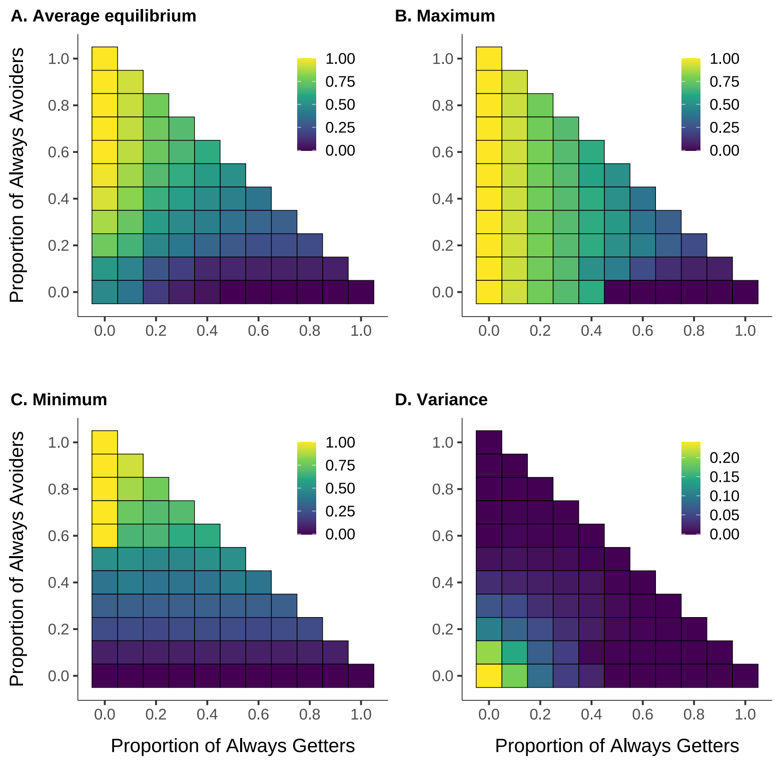

Notes: The figure illustrates how group composition affects summary features of the distribution of information equilibria. Each cell in the unit probability triangle represents a distinct group composition, defined by the proportion of Always Getters (x-axis), Always Avoiders (y-axis), and Strategic Complements (implied by the vertical, or horizontal, distance to the hypotenuse). Each panel reports a different feature of the equilibrium distribution —the average (A), maximum (B), minimum (C), or variance (D)— with values represented by the color gradient. Only stable equilibria are considered; Strategic Substitutes are excluded ($p_{SS} = 0$), as their equilibria are typically unstable.

D Instructions and other screens

Preliminary instructions

At the start of the session, all participants received a printed copy of the instructions below, which were also read aloud by the experimenter.

> Welcome to the CeDEx Lab!
>
> Thank you very much for participating in this scientific experiment.
>
> This is an experiment on decision-making. By carefully reading the instructions that will be shown on your computer, you can —depending on the decisions you make— earn an appreciable amount of money. It is therefore important that you pay close attention to the instructions.
>
> We are interested in your individual choices. Therefore, communication is not allowed during the experiment. If you have a question, please just raise your hand.
>
> In addition, there are some ground rules that need to be complied with to ensure fair conditions for all participants:
>
> - **Focus on your own screen**. You should neither interfere with other participants nor look at their workstations (including their screens). Please do not lean backwards and always keep your chair close to your desk.
> - **Be patient**. There will be some periods in the experiment in which you will be asked to wait for some time (for example, you might finish your tasks earlier than other people; or you might be asked to just plainly wait). In any case, please remain in your workstation, waiting in silence, without disturbing other participants.
> - **No phones**. Please, put your phone(s) in silence mode and put them away (for instance, put them in your coat or bag): do not leave them on your desk. Please do it now. During this session, you should not check or use your phone at any time. If you check or use your phone, you will be asked to leave the session. This is to ensure fair play conditions for all participants.
> - **Headphones at all times**. In this experiment, you are asked to wear the headphones provided next to your computer at all times. This is because sounds may be played to you as we will explain later. If you don't properly wear your headphones, you will be asked to leave the experiment.

If you do not comply with any of these ground rules, you will be asked to leave the experiment. Throughout the experiment, some of us, experimenters, will be walking around invigilating for these conditions to be met.

ou will soon be given instructions about today's experiment. All instructions will be provided on your computer screen.

Note that participation in this study is entirely voluntary. You can leave the experiment at any time you want. If you leave, we will take into account your progress up to that point to calculate your payment and will pay you accordingly.

In case you have any questions, feel free to ask the experimenter at any time, in a very quiet voice so that no one else is disturbed.

Now, please put on your headphones and then we will be ready to start the experiment.

After these instructions were read aloud, subsequent instructions were presented on participants' computer stations, as described in the next subsection.

Main instructions and decision screens

The following text reproduces the instructions displayed to participants via the computer stations. Each new screen in the instructions is marked below by a horizontal line. We provide clarifying notes in square brackets in italics.

General instructions for participants

Welcome and thank you for participating in this scientific experiment.

First, you will be asked for some contact information. Then, you will be given the instructions on how to earn money in this experiment.

Contact Information

In the following, we will ask for the email address associated with your PayPal account, and your university email address.

This information will be stored only for payment and research purposes.

Please, **be sure that your email addresses are typed correctly. Otherwise, we may not be able to make the payment**.

Please input the email address associated to your PayPal account. [Open text box]

Please type it again. [Open text box]

Now input your university email address. [Open text box]

Input it again. [Open text box]

Please input your **seat number** (this is the number written in the plastic card that you have drawn when you entered into the lab).

[Open text box]

Your consent to participate

Please be informed that the present study involves the possibility of listening to sounds that have been rated as aversive (i.e., unpleasant or disturbing) by participants in other studies. The safety of the participants is a priority for the CeDEx research centre. Therefore, technical measures have been put in place to ensure the hearing safety of the participants.

In order to participate in today's study, your explicit consent to participate is needed. Please note that participation in this study is entirely voluntary.

If you do not give your explicit consent, you will withdraw from this study. There will be no further consequences, and we will keep inviting you to other studies. We will give you £3 as a token of appreciation for coming today.

○ "I do NOT GIVE my explicit consent to participate"

○ "I GIVE my explicit consent to participate"

In this experiment, you and other participants in the room will be part of the same group.

Groups are typically composed of around 30 participants.

Introduction

In today's experiment, you will be asked to do some tasks, which will be explained as you progress through the session. Depending on your performance (and your performance alone) in those tasks, your earnings may increase.

For your participation in today's session, you will receive £5, plus additional earnings based on your performance in the tasks.

Instructions: The Quiz (1/2)

In this experiment, you will be participating in a Quiz. The Quiz questions cover different topics (for instance, sports, geography, history, arts, music, etc.).

For each question, you will be provided with 4 possible answers, out of which exactly one will be correct.

In this Quiz task, **the more questions you answer correctly, the more you earn**.

The Quiz consists of **two parts**, namely Part 1 and Part 2.

- **Quiz Part 1**: You start with Part 1 of the quiz. You have 3 minutes to answer as many quiz questions as possible.

- **Quiz Interruption and Decisions**: After those 3 minutes, the Quiz will be briefly interrupted. During this interruption, you will be asked to make some decisions, as will be explained on the following screens.
- **Quiz Part 2**: After the interruption, you will answer Part 2 of the quiz, for 4 minutes.

The following figure illustrates the timeline:

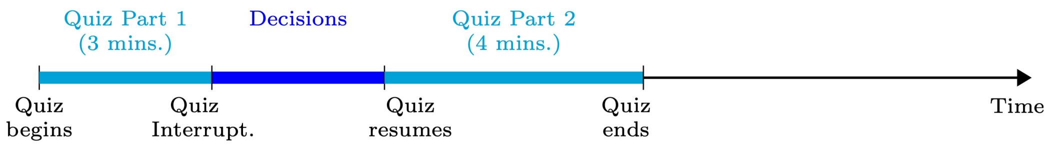

The Quiz (2/2)

Below we explain how your earnings are calculated. The key feature is that **you should try to maximise the number of questions answered correctly**.

The number of correct answers in Parts 1 and 2 will be added up to calculate your earnings.

Your earnings increase in accordance with the level you reach. You start at Level 1. To reach Level 2, you need to answer 10 questions correctly (you do not need to correctly answer 10 questions in a row, what counts is the total number of correct answers). Then, to reach each additional level you need to answer correctly 20 extra questions per level. Thus, your earnings from the quiz are determined as follows:

- Level 1 (0-9 correct answers) = £5
- Level 2 (10-29 correct answers) = £6
- Level 3 (30-49 correct answers) = £7
- Level 4 (50-69 correct answers) = £9
- Level 5 (70-89 correct answers) = £13
- Level 6 (90+ correct answers) = £20

Your earnings from the quiz will be added to the £5 you will earn just for participating in this experiment.

Example: suppose you answer correctly 25 questions in Part 1 and 35 in Part 2. Then the total number of correct answers is 60 (=25+35), so the level you reach is 4. Therefore, your earnings from the quiz would be £9. Adding up the £5 from your participation in this experiment, your total earnings would be £14 (=£9 + £5).

You can see that your earnings increase substantially, the higher the level that you reach.

The Risk Period

After the quiz finishes, a period called the "Risk Period" will start. The Risk Period lasts for 4 minutes.

During the Risk Period, a series of events may or may not happen as explained on the next screens.

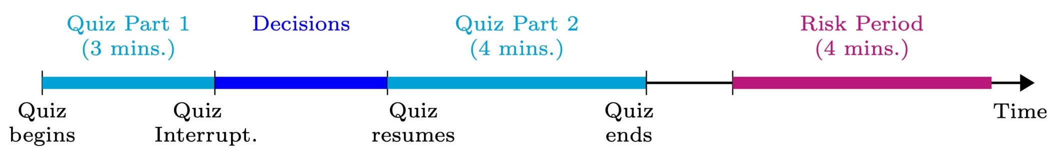

The Outcome

In the Risk Period, one of two outcomes will happen:

- ***Screams*** outcome: During the Risk Period, your group (including you) will hear a series of **unpleasant distress screams** from your headphones. These will happen at *unpredictable* moments during the Risk Period.

  Previous research has shown that people consider similar screams as more aversive (disturbing) than the sound of nails sliding on a chalkboard.

- ***Quiet*** outcome: you and your group will not hear any screams during the whole experiment.

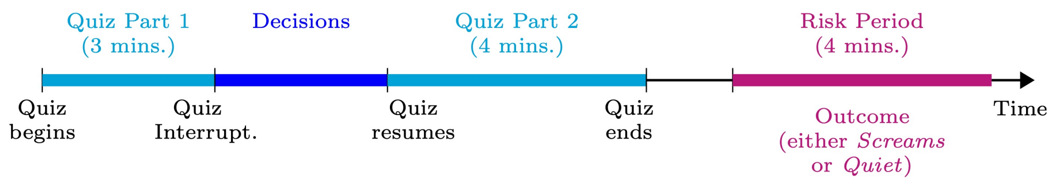

The Lottery

The Outcome (*Screams* or *Quiet*) in the Risk Period is the **same for everyone** in your group, and will be randomly determined by a computerised lottery.

Revelation of Outcome

The Outcome (i.e., *Screams* or *Quiet*) will be **revealed to everyone**, just a **few screens before the Risk Period**.

Therefore, everybody will already know the Outcome before the Risk Period starts.

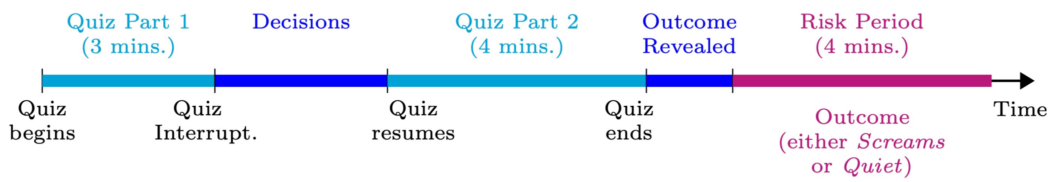

Your Decision (1/3)

However, you can choose whether you want to know about the Outcome (*Screams* or *Quiet*) earlier or whether you prefer not to be aware of the Outcome until the beginning of the Risk Period.

At the end of Part 1 of the quiz, you will choose one of two options:

- Option "**Now**" implies that you will learn the Outcome (Screams or Quiet) **immediately** (i.e., during the interruption of the quiz).
- Option "**Later**" implies that you will **not** learn about the Outcome (Screams or Quiet) immediately, and rather you will learn it just before the Risk Period.

  This means that, during Part 2 of the quiz, you will not know whether in the forthcoming Risk Period you will hear the screams or not.

So the difference between "Now" and "Later" is **whether in Part 2 of the Quiz you are aware** of the Outcome or not. However, the Outcome is not affected by the option you choose.

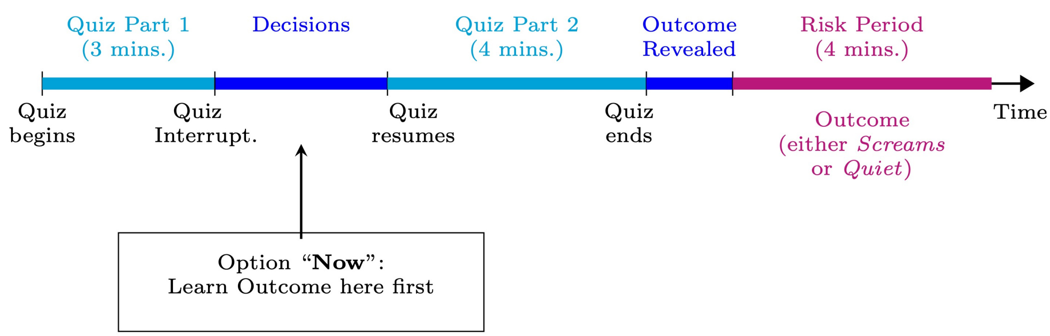

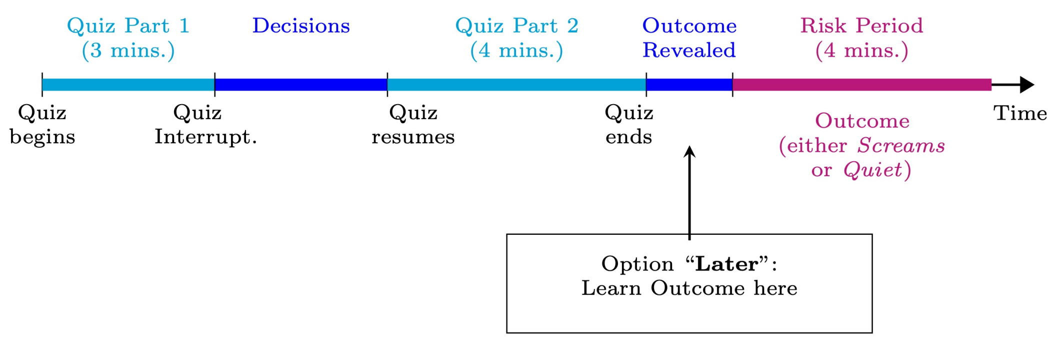

The Volume of Screams

Although the Outcome is not affected by the option you choose, the **volume** of the screams (in case the Outcome is *Screams*) is affected by the choices that you and everyone else make:

- Each group member who chooses **"Later" increases the volume level by about 2 points out of 100**.

- Hence, if everybody chooses "Later", then the screams will be played at maximum volume (we call this level of volume "level 100").
- If nobody chooses "Later", then the screams will be played at volume "level 50".

This means that, if the Outcome is *Screams*, your choice between "Now" and "Later" would generate a difference in volume of about **2 points, out of 100**.

If instead the Outcome is *Quiet*, your choice has no implications beyond the timing in which you will find out the Outcome.

Check your Understanding

To check your understanding of the instructions, please imagine you are in each of the situations below and answer the following questions.

Situation 1a) Suppose that the Outcome is ***Screams***. In addition, suppose that **everybody** in your group (including you) chooses "**Later**". What would be the volume of the screams?

○ Level 50

○ Level 52

○ Level 98

○ Level 100

○ There would be no screams

Situation 1b) Suppose that the Outcome is ***Screams***. In addition, suppose that **everybody else** in your group chooses "**Later**", but **you** choose "**Now**". What would roughly be the volume of the screams?

○ Level 50

○ Level 52

○ Level 98

○ Level 100

○ There would be no screams

Situation 2a) Suppose that the Outcome is ***Screams***. In addition, suppose that **everybody** in your group (including you) chooses "**Now**". What would be the volume of the screams?

- ○ Level 50
- ○ Level 52
- ○ Level 98
- ○ Level 100
- ○ There would be no screams

Situation 2b) Suppose that the Outcome is ***Screams***. In addition, suppose that **everybody else** in your group chooses "**Now**", but **you** choose "**Later**". What would roughly be the volume of the screams?

- ○ Level 50
- ○ Level 52
- ○ Level 98
- ○ Level 100
- ○ There would be no screams

Situation 3) Suppose that the Outcome is ***Screams***. In that case:

- Your decision between "Now" and "Later" will...
  - ○ Only affect the volume of the screams to be heard by you
  - ○ Only affect the volume of the screams to be heard by participants other than you
  - ○ Affect the volume of the screams to be heard by everyone (including you)
- The decisions between "Now" and "Later" **by participants other than you** will...
  - ○ Only affect the volume of the screams to be heard by you
  - ○ Only affect the volume of the screams to be heard by participants other than you
  - ○ Affect the volume of the screams to be heard by everyone (including you)

Situation 4) Suppose that the Outcome is ***Quiet***. In addition, suppose that **everybody** in your group (including you) chooses "**Later**". What would roughly be the volume of the screams?

- ○ Level 50
- ○ Level 52
- ○ Level 98
- ○ Level 100
- ○ There would be no screams at all

When you are ready, press the "Continue" button below.

(In case you are not taken to the next screen: some of your answers to the questions above might be wrong, please review them).

Instructions: Your Decision (2/3)

As mentioned, the difference between “Now” and “Later” is whether in Part 2 of the Quiz you are aware of the Outcome or not, and whether the volume of the screams (if the Outcome is Screams) changes by 2 points out of 100.

However, your choice between “Now” and “Later” does **not** change any of the following:

- the duration of the session;
- the sequence of the quiz;
- the amount of time you can spend on the quiz;
- when or whether you may potentially hear the screams.

Finally, note that for this choice between “Now” and “Later”, there is no right or wrong choice. You should simply decide based on your own preferences.

Instructions: Your Decision (3/3)

Recall that if you choose “Now”, then you will be shown a screen that informs you about the Outcome.

Upon being informed about the Outcome, you will be asked to select a button that correctly reflects it. You will need to do it to proceed to the next screens, which include Part 2 of the Quiz.

Differently, if you choose "Later", then you will be shown a screen that will ask you to wait, and you will not learn the Outcome until Quiz Part 2 is finished, just before the Risk Period.

Control Questions

To briefly check that you understood the instructions correctly, please answer the following questions.

- If you choose option "Now"...
  - ○ You will learn the Outcome immediately
  - ○ You will not learn the Outcome immediately, but rather just before the Risk Period
  - ○ You will never learn the Outcome
- If you choose option "Later"...
  - ○ You will learn the Outcome immediately
  - ○ You will not learn the Outcome immediately, but rather just before the Risk Period
  - ○ You will never learn the Outcome
- If the Outcome is *Screams*, the **timing** in which you will hear the screams...
  - ○ will be earlier if you choose option "Now"
  - ○ will be earlier if you choose option "Later"
  - ○ will be the same independently of whether you choose "Now" or "Later"
- If the Outcome is *Screams*, the **volume** of the screams...
  - ○ will be 2 points higher if you choose option "Later"
  - ○ will be 2 points higher if you choose option "Now"
  - ○ will be the same independently of whether you choose "Now" or "Later"

- Just before the Risk Period starts, you will be informed about the Outcome...
  - ○ Only if you have chosen “Now”
  - ○ Only if you have chosen “Later”
  - ○ In any case, independently of whether you have chosen “Now” or “Later”
  - ○ You will not be informed about the Outcome at that point

When you are ready, press the "Continue" button below.

(In case you are not taken to the next screen: some of your answers to the questions above might be wrong, please review them).

A Pilot Study

[Note: The text below was shown to participants in the Many Avoided condition. The version for the Few Avoided condition differed, with deviations indicated in square brackets.]

The information below gives you an idea of the choices of other people in a similar situation.

We ran a pilot study similar to this experiment, and found that, out of a set of 10 University of Nottingham students, **8 [2] chose the option “Later”**, preferring not to find out the Outcome until the end of Part 2. Therefore, in that group, **the overwhelming majority [minority] chose “Later”**.

If people in your group behave similarly to those in the pilot study, you can expect around 23 [6] people other than you to choose Later, and around 6 [23] to choose Now. In that case, the resulting volume level would roughly be 89 [60] if you choose Now, and 91 [62] if you choose Later.

Out of any 10 members of your group (*other than you*), how many do you believe will choose “Later”?

[Slider ranging from “0 out of 10 will choose “Later”” to “10 out of 10 will choose “Later””]

You are now ready to start the quiz

From the next screen, the quiz will start.

When you feel ready to start, press "Continue", and both the quiz and the 3-minute timer will start at once.

Quiz

Please, answer the following questions to achieve a higher level and earn a higher payment.

[List of multiple choice quiz questions is displayed.]

Your quiz time has ended.

Please, press the Continue button.

The quiz has been interrupted. You will now be asked to make a series of decisions, including your preferred choice between knowing your lottery Outcome "Now" and "Later".

What do you think the Outcome of your group's lottery will be? (*Screams* or *Quiet*)

[Radiobuttons (10 levels), with extremes labeled as "Definitely Quiet" and "Definitely Screams".]

[Note: this screen elicits the Kreps-Porteus questions (KP questions)]

Before choosing between "Now" and "Later", please answer the following questions about how happy you would be in different (hypothetical) situations.

- Imagine that you choose "Now", and then a moment later discover that the Outcome is *Screams*. How happy would you be in that moment *thinking about the Outcome ahead*?

  [Slider ranging from -10="Very unhappy" to 10="Very happy".]

- Imagine that you choose "Now", and then a moment later discover that the Outcome is *Quiet*. How happy would you be in that moment *thinking about the Outcome ahead*?

  [Slider ranging from -10="Very unhappy" to 10="Very happy".]

- Imagine that you choose "Later", and therefore you do not learn the Outcome until later. How happy would you be in that situation *thinking about the Outcome ahead*?

  [Slider ranging from -10="Very unhappy" to 10="Very happy".]

Do you want to know the Outcome "Now" or "Later"?

Choose whether you want to know "Now" or "Later" whether the screams are upcoming.

○ Later

○ Now

How strong is your preference for your selected option, over the other?

[Radiobuttons (10), ranging from "I am **indifferent** between the two options" to "I have a **very strong** preference for the selected option".]

A quick question

You have chosen the option ["Now"/"Later"]. Why did you prefer this option?

Please take a moment to answer this question.

[Open response textbox.]

Instructions: The Decisions Table

One participant will have the opportunity to **condition their decision on what the other participants chose** in the previous screen. This participant will be selected randomly by the computer.

Suppose you are that randomly selected participant. Then, the decision between "Now" and "Later" that you complete in the table on the following screen will be implemented for you.

See the image below for an example. In the table, you have the possible choices that other participants have made in the previous screen. Each row represents a different situation. **For each row, you can decide which option you would prefer**. Then, according to what the other players have *actually* chosen in the previous screen, the relevant row of the table will be determined, and the option selected in that row of the table will actually be implemented for you.

You have a positive chance of being selected as this player. Therefore, you must complete the table carefully according to your preferences.

(Please scroll down to see the 'Continue' button)

Out of 10 other participants, those who chose *Later* are...	In that situation, you prefer to know the Outcome...	Volume of screams if you choose Now	Volume of screams if you choose Later
0	Now Later	50	52
2	Now Later	60	62
4	Now Later	69	71
6	Now Later	79	81
8	Now Later	89	91
10	Now Later	98	100

Please, select your preferred option for each of the following situations.

Out of 10 other participants, those who chose *Later* are...	In that situation, you prefer to know the Outcome...	Volume of screams if you choose Now	Volume of screams if you choose Later
0	Now Later	50	52
2	Now Later	60	62
4	Now Later	**69**	**71**
6	Now Later	79	81
8	Now Later	89	91
10	Now Later	98	100

You will now be asked a series of questions. To answer them, please imagine you are in each of the situations presented.

Questions (1/5)

[Note: This and following screens elicit the Information Preferences Scale (IPS; Ho et al., 2021)]

Please, imagine you are in each of the following situations, and answer the questions:

1. As part of a medical checkup, your doctor asks you a series of questions. The answers to these questions can be used to estimate your life expectancy (the age you are predicted to live to). Do you want to know how long you can expect to live?
   - ○ Definitely don't want to know
   - ○ Probably don't want to know
   - ○ Probably want to know
   - ○ Definitely want to know

2. You provide some genetic material to a testing service to learn more about your ancestors. You are then told that the same test can, at no additional cost, tell you whether you have an elevated risk of developing Alzheimer's (a brain disorder that slowly destroys memory and thinking skills and, eventually, the ability to carry out the simplest tasks). Do you want to know whether you have a high risk of developing Alzheimer's?

   ○ Definitely don't want to know
   ○ Probably don't want to know
   ○ Probably want to know
   ○ Definitely want to know

3. At your medical checkup, you are given the option to see the results of a diagnostic test, which can identify the extent to which your body has suffered long-term effects from stress. Do you want to know how much lasting damage your body has suffered from stress?

   ○ Definitely don't want to know
   ○ Probably don't want to know
   ○ Probably want to know
   ○ Definitely want to know

Questions (2/5)

4. Ten years ago, you had the opportunity to invest in two retirement funds: Fund A and Fund B. For the past 10 years, you have invested all your retirement savings in Fund A. Do you want to know the balance you would have if you had invested in Fund B instead?

   ○ Definitely don't want to know
   ○ Probably don't want to know
   ○ Probably want to know
   ○ Definitely want to know

5. You decide to go to the theater for your birthday and give your close friend (or partner) your credit card so they can purchase tickets for the two of you, which they do. You aren't sure but suspect that the tickets may have been expensive. Do you want to know how much the tickets cost?
   - ○ Definitely don't want to know
   - ○ Probably don't want to know
   - ○ Probably want to know
   - ○ Definitely want to know
6. You bought an electronic device at a store at what seemed like a reasonable, though not particularly low, price. A month has passed, and the item is no longer returnable. You see the same device displayed in another store with a sign announcing "SALE." Do you want to know the price you could have bought it for?
   - ○ Definitely don't want to know
   - ○ Probably don't want to know
   - ○ Probably want to know
   - ○ Definitely want to know

Questions (3/5)

7. You gave a close friend one of your favorite books for her birthday. Visiting her apartment a couple of months later, you notice the book on her shelf. She never said anything about it; do you want to know if she liked the book?
   - ○ Definitely don't want to know
   - ○ Probably don't want to know
   - ○ Probably want to know
   - ○ Definitely want to know
8. Someone has described you as quirky, which could be interpreted in a positive or negative sense. Do you want to know which interpretation he intended?
   - ○ Definitely don't want to know
   - ○ Probably don't want to know

- ○ Probably want to know
- ○ Definitely want to know

9. You gave a toast at your best friend's wedding. Your best friend says you did a good job, but you aren't sure if he or she meant it. Later, you overhear people discussing the toasts. Do you want to know what people really thought of your toast?
   - ○ Definitely don't want to know
   - ○ Probably don't want to know
   - ○ Probably want to know
   - ○ Definitely want to know
10. As part of a fundraising event, you agree to post a picture of yourself and have people guess your age (the closer they get, the more they win). At the end of the event, you have the option to see people's guesses. Do you want to learn how old people guessed that you are?
    - ○ Definitely don't want to know
    - ○ Probably don't want to know
    - ○ Probably want to know
    - ○ Definitely want to know
11. You have just participated in a psychological study in which all of the participants rate others' attractiveness. The experimenter gives you an option to see the results for how people rated you. Do you want to know how attractive other people think you are?
    - ○ Definitely don't want to know
    - ○ Probably don't want to know
    - ○ Probably want to know
    - ○ Definitely want to know

Questions (4/5)

12. Some people seek out information even when it might be painful. Others avoid getting information that they suspect might be painful, even if it could be useful. How would you describe yourself?

- ○ Definitely don't want to know
- ○ Probably don't want to know
- ○ Probably want to know
- ○ Definitely want to know

13. If people know bad things about my life that I don't know, I would prefer not to be told.
    - ○ Definitely don't want to know
    - ○ Probably don't want to know
    - ○ Probably want to know
    - ○ Definitely want to know

Questions (5/5)

The questions below ask about what you would decide in the workplace. If you have previously worked in a company, imagine yourself in it and answer the questions. If you had not worked for a company, read the following questions as imagining you are working in a company.

14. Employers commonly encourage workers to complete health risk assessments that gauge their health characteristics and risk. These are usually used to estimate workplace-wide risks the company may face. If you had the opportunity, would you like to get your individual risk assessment?
    - ○ Definitely don't want to know
    - ○ Probably don't want to know
    - ○ Probably want to know
    - ○ Definitely want to know
15. Most people spend some time at work on activities they are not proud of (e.g., browsing social media). Suppose you could compose a list of these activities and have your computer track how much time you spend doing them. The information would not be shared with your employer. Would you like to know what fraction of each hour at work you spend on these activities?

- ○ Definitely don't want to know
- ○ Probably don't want to know
- ○ Probably want to know
- ○ Definitely want to know

16. As part of a unit-wide evaluation, you and your coworkers are asked how easy it is to get along with each other and what your strengths and weaknesses are. Would you like to know how your coworkers rated you?
    - ○ Definitely don't want to know
    - ○ Probably don't want to know
    - ○ Probably want to know
    - ○ Definitely want to know
17. The rise of artificial intelligence is likely to lead to job losses in a wide range of occupations, affecting workers in industries from long-distance trucking to healthcare. Would you like to know whether your job is at risk owing to automation in the next 10 years?
    - ○ Definitely don't want to know
    - ○ Probably don't want to know
    - ○ Probably want to know
    - ○ Definitely want to know
18. One option in your employer's retirement plan allows you to compare your investment return with that of your coworkers. Would you like to know how your investment return compares with the average and highest investment returns earned by employees at your firm?
    - ○ Definitely don't want to know
    - ○ Probably don't want to know
    - ○ Probably want to know
    - ○ Definitely want to know

Please answer the following questions about how happy you would be in different (hypothetical) situations.

- Imagine that the Outcome is *Screams*, and that the **volume of the screams** will be **level 50**. How happy would you thinking about the Outcome ahead?
  [Slider from -10="Very unhappy" to 10="Very happy"]
- Imagine that the Outcome is *Screams*, and that the **volume of the screams** will be **level 100**. How happy would you thinking about the Outcome ahead?
  [Slider from -10="Very unhappy" to 10="Very happy"]

Thank you for answering the questions.

In the next screen, you will get to know the Outcome if you selected option Now, or remain without knowing it, if you selected option Later. If you were selected as the random player, then the option you selected in the relevant row of the table will be implemented.

---

*[Note: this screen was displayed to participants whose implemented option was **Now**]*

Based on your previous decisions, the choice implemented for you is "Now".

The Outcome is Screams [Quiet].

Therefore, you **will hear the screams [not hear any screams]** during the Risk Period.

Please, indicate below which is the Outcome.

○ Screams

○ Quiet

---

*[Note: this screen was displayed to participants whose implemented option was **Later**]*

Based on your previous decisions, the choice implemented for you is "Later".

You may or may not hear the screams during the Risk Period. You will be informed about the Outcome later, just before the Risk Period starts.

Press Continue to proceed to Quiz Part 2.

Quiz

Please, answer the following questions to achieve a higher level and earn a higher payment.

[The same list of multiple choice quiz questions is displayed.]

The Quiz has ended.

You will be asked a series of questions, and then move to the Risk Period.

Questions (1/5)

[Note: This and following screens elicit social and economic preferences from the Global Preferences Survey (GPS; Falk et al., 2018)]

We now ask for your **willingness to act** in a certain way in four different areas.

Please indicate your answer on a scale from 0 to 10.

How willing are you to give up something that is beneficial for you today in order to benefit more from that in the future? [Slider: 0="Completely unwilling to do so", 10="Very willing to do so".]

In general, how willing or unwilling are you to take risks? [Slider: 0="Completely unwilling to take risks", 10="Very willing to take risks".]

Questions (2/5)

1. When someone does me a favor I am willing to return it.

   [Slider: 0="Does not describe me at all", 10="Describes me perfectly".]

2. Please think about what you would do in the following situation:

   You are in an area you are not familiar with, and you realize you lost your way. You ask a stranger for directions. The stranger offers to take you to your destination. Helping you costs the stranger about £20 in total. However, the stranger says he or she does not want any money from you. You have six presents with you. The cheapest present costs £5, the most expensive one costs £30. Do you give one of the presents to the stranger as a "thank-you" gift? If so, which present do you give to the stranger?

   ○ No present
   ○ Present worth £5
   ○ Present worth £10
   ○ Present worth £15
   ○ Present worth £20
   ○ Present worth £25
   ○ Present worth £30

Questions (3/5)

If I am treated very unjustly, I will take revenge at the first occasion, even if there is a cost to do so. [Slider: 0="Does not describe me at all", 10="Describes me perfectly".]

How willing are you to punish someone who treats others unfairly, even if there may be costs for you? [Slider: 0="Completely unwilling to do so", 10="Very willing to do so".]

Questions (4/5)

Imagine the following situation: Today you unexpectedly received £1,000. How much of this amount would you donate to a good cause? [Open text box; inputs restricted to numbers between 0 and 1000]

How willing are you to give to good causes without expecting anything in return? [Slider: 0="Completely unwilling to do so", 10="Very willing to do so".]

Questions (5/5)

“I assume that people have only the best intentions.” [Slider: 0=“Does not describe me at all”, 10=“Describes me perfectly”.]

Instructions: the Risk Period (1/2)

Thanks for kindly answering the questions.

You are now about to enter the Risk Period.

Before the Risk Period starts, two things will happen:

1. In the next screen, you will be provided with a sample scream, to get an idea of the kind of screams that may happen during the Risk Period.
2. After that, everybody will be informed about: a) which is the Outcome (*Screams* or *Quiet*); and b) the volume of the screams as the result of the group’s choices, if the Outcome is Screams.

After that, the Risk Period will start. It will last 4 minutes, and you will be provided with the answers to the quiz to review if you want to.

Pre-experiencing a sample scream (1/2)

To give you an idea of what the screams sound like, on the next screen, you will hear a sample scream, similar to the screams you would hear if your lottery outcome were *Screams*.

On the next screen, you will hear only 1 scream, whereas during the Risk Period you may hear more than 1 scream if your lottery Outcome is *Screams*, and 0 if your lottery outcome is *Quiet*.

To ensure your safety, the volumes of the screams have been calibrated to comply with the safe hearing levels recommended in the UK.

Please, keep your headphones on at all times. You are also reminded you may leave the session at any time you would like.

Please, try to remain silent throughout.

Now, if you wish to proceed, press Continue.

Pre-experiencing a sample scream (2/2)

To hear the sample scream, press the "Play" button. From the moment you press the "Play" button below, the sample scream will be played at a random moment within 30 seconds.

You will be able to hear the sample scream only once. In other words, the button will not do anything from the second time you press it.

Please, try to remain silent throughout.

After the 30 seconds, the Continue button will appear. You will be able to proceed by pressing it.

[Play button]

The Outcome, revealed

[If the lottery Outcome was Screams, the following text was displayed:]

Your group's lottery Outcome is ***Screams***.

Therefore, your group **will hear the screams** during the Risk Period.

The volume level of the screams is [**volume level**]. This is the result of the choices of all participants.

[If the lottery Outcome was Quiet, the following text was displayed:]

Your lottery outcome is ***Quiet***.

Therefore, you will **not hear any screams** during the Risk Period.

Risk Period

The Risk Period has started. Please, remain in silence until it ends.

In the meantime, you can read the list of quiz questions and answers below.

At the end of the Risk Period, a "Continue" button will appear, at the bottom. (You may need to scroll down)

Quiz Answers

Below is the list of quiz questions. The correct answers to these are underlined.

[List of all quiz questions and their answers is displayed]

Thank you for your patience. The Risk Period has ended.

You will now be asked a few final questions.

Please, press the Continue button to proceed.

Final Questions (1/4)

Please, kindly answer the following questions.

- How easy to understand were the instructions? [Radiobuttons (5): "Very easy" to "Very hard"]
- How hard to answer were the quiz questions? [Radiobuttons (5): "Very easy" to "Very hard"]
- Do you have any experience at work, even casually? ["Yes"/"No"]

Final Questions (2/4)

- All in all, how do you feel about this study? [Open text box]
- Is there anything else you would want to tell us? [Open text box]

Final Questions (4/4)

- What is your age? [Open text box, restricted to numbers]
- What is your gender? [Female/Male/Other]

You have answered [Q] questions correctly. Therefore, you have reached performance level [P]. In consequence, you have earned £[e] from the quiz.

Adding the £5 for your participation, your **total earnings** from the experiment are £[$e + 5$].

You will be contacted about your payment in the next couple of days.

Thanks for your participation

We appreciate that you took the time to come and participate in today's experiment. Thanks a lot, and hope to see you around soon!

Your session is ending now. You can now leave the lab. Please, do it in silence, so other participants are not disturbed.