

NBER WORKING PAPER SERIES

INFORMATION AVOIDANCE AND IMAGE CONCERNS

Christine L. Exley
Judd B. Kessler

Working Paper 28376
<http://www.nber.org/papers/w28376>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2021

For very helpful feedback on this paper, we thank Russell Golman, Zachary Grossman, Davide Pace, Matthew Rabin, Joshua Schwartzstein, Marta Serra Garcia, Joël van der Weele, and Roberto Weber. For funding, we thank Harvard Business School, the Wharton School, and the Wharton Behavioral Lab. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Christine L. Exley and Judd B. Kessler. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Information Avoidance and Image Concerns
Christine L. Exley and Judd B. Kessler
NBER Working Paper No. 28376
January 2021
JEL No. C9,C91

ABSTRACT

A rich literature finds that individuals avoid information, even information that is instrumental to their choices. A common hypothesis posits that individuals strategically avoid information to hold particular beliefs or to take certain actions—such as behaving selfishly—with lower image costs. Building off of the classic “moral wiggle room” design, this paper provides the first direct test of whether individuals avoid information because of image concerns. We analyze data from 4,626 experimental subjects. We find that image concerns play a role in driving information avoidance, but a role that is substantially smaller than the common approach in the literature would suggest. The large majority (66% to 81%) of information avoidance remains when image concerns cannot drive avoidance. We find evidence for other reasons why individuals avoid information, such as a desire to avoid interpersonal tradeoffs, a desire to avoid bad news, laziness, inattention, and confusion.

Christine L. Exley
Harvard Business School
25 Baker Way
Baker Library 449
Boston, MA 02163
clexley@gmail.com

Judd B. Kessler
The Wharton School
University of Pennsylvania
3620 Locust Walk
Philadelphia, PA 19104
and NBER
judd.kessler@wharton.upenn.edu

A data appendix is available at <http://www.nber.org/data-appendix/w28376>

1 Introduction

People frequently avoid information that could be instrumental to their decisions. They avoid medical tests that might encourage lifestyle changes (Thornton, 2008; Oster, Shoulson and Dorsey, 2013; Ganguly and Tasoff, 2017); investment outcomes that might influence financial strategies (Karlsson, Loewenstein and Seppi, 2009; Sicherman et al., 2016); thorough reviews of job applications that might discourage discrimination (Bartoš et al., 2016); nutritional facts that might encourage healthier eating (Thunström et al., 2016); details on how pay is determined that might encourage higher effort on a task (Huck, Szech and Wenner, 2017); news that might increase kindness towards refugees (Freddi, 2018); information about the environmental consequences of their actions that might encourage more eco-friendly choice (d’Adda et al., 2018); and information about how their actions influence others that might encourage less selfish behavior (Dana, Weber and Kuang, 2007).

Why do individuals avoid information in these contexts? A number of lines of research suggest that individuals avoid information in order to maintain certain beliefs (e.g., about themselves as healthy, financially responsible, politically enlightened, kind) even while taking actions that could suggest the opposite. Such explanations, however, rely on the sophistication of agents to strategically avoid information in order to maintain certain beliefs or in order to construct plausible deniability about their actions. In this paper, we explore the last context—on selfishness—to directly test whether individuals strategically avoid information because of image concerns by experimentally varying whether image concerns can drive information avoidance.

The context we explore has been the focus of a rich literature building off of Dana, Weber and Kuang (2007). That seminal paper generates a set of empirical results that have proven to be robust and influential. First, individuals frequently avoid information on how their choices influence the payoffs of others. Second, individuals make substantially more selfish decisions when they are able to avoid information than when the information cannot be avoided. Third, the fraction of individuals who avoid information is higher than the fraction of individuals who might be expected to avoid information because it will not affect their decision. These findings have been replicated many times (Larson and Capra, 2009; Matthey and Regner, 2011; Feiler, 2014; Grossman, 2014; Exley, 2016; Grossman and van der Weele, 2017), and they have raised an important debate about what drives passive information avoidance.¹

What drives this information avoidance? A leading explanation is image concerns—and, notably, self-image concerns about how one views oneself, given subjects’ anonymity in the prior work.² Significant empirical evidence supports the notion that image costs of acting selfishly are

¹See also conceptual replications in different paradigms (Kajackaite, 2015; Serra-Garcia and Szech, 2020).

²For important work on models of image concerns, see Rabin (1995); Bodner and Prelec (2003); Bénabou and Tirole (2004, 2006); Mijović-Prelec and Prelec (2010); Bénabou and Tirole (2011); Nyborg (2011); Grossman (2015); Grossman and van der Weele (2017); Bénabou, Falk and Tirole (2018); Foerster and van der Weele (2018).

smaller when individuals do not know for certain they are being selfish.³ If individuals recognize that being uninformed will decrease the image costs of selfishness, they may strategically avoid information so they can behave selfishly with a lower image cost.

In this paper, we provide the first direct empirical test of the popular hypothesis that individuals avoid information because of image concerns. We do this by comparing rates of information avoidance in the classic setting—built off of [Dana, Weber and Kuang \(2007\)](#)—to a new setting that only differs by removing selfish motives and thus, as detailed below, removes image concerns to *avoid* information. Across the two settings, we hold constant the structure of the decision, the content of the information, and the timing of information provision.⁴

In both settings, participants are asked to choose between A and B. In the classic setting, which we call the *Self/Other* condition, participants know that they earn more from choosing A but do not know whether A or B is better for another participant. They can avoid information and choose A or B directly, or they can learn which is better for the other participant before choosing. In this setting, participants might strategically avoid information so they can choose A directly, incurring lower image costs than if they chose A after learning for certain that A was worse for the other participant.

In our new setting, which we call the *Other/Other* condition, participants' decisions influence payoffs for two other participants but not for themselves. The payoffs, the structure of the decision, and the content of the information are the same as in the classic setting, except a different participant receives the payoff that would have gone to the decision maker. In this condition, image concerns cannot drive information avoidance. To see this, first note that image concerns about selfishness are clearly not relevant because selfish motives are removed. Moreover, even other image concerns—such as a desire to appear fair—cannot drive information *avoidance* in the *Other/Other* condition. Individuals with such image concerns should instead *acquire* information and choose the option aligned with those image concerns, which they can do without suffering a financial cost. That is, in the *Other/Other* condition there is no chance that acquiring information will force a tradeoff between choosing an option aligned with image concerns and an

³Evidence consistent with this has been shown in how individuals are viewed by others ([Conrads and Irlenbusch, 2013](#); [Krupka and Weber, 2013](#); [Bartling, Engl and Weber, 2014](#); [Grossman and van der Weele, 2017](#)) and by themselves ([Hsee, 1996](#); [Oberholzer-Gee and Eichenberger, 2008](#); [Haisley and Weber, 2010](#); [Falk and Szech, 2013](#); [Di Tella et al., 2015](#); [Danilov and Saccardo, 2016](#); [Exley, 2016](#); [Bénabou and Tirole, 2016](#); [Gino, Norton and Weber, 2016](#); [Bartling and Özdemir, 2017](#); [Falk and Szech, 2017](#); [Gneezy, Saccardo and van Veldhuizen, 2018](#); [Regner, 2018](#); [Olschewski et al., 2019](#); [Garcia, Massoni and Villeval, 2020](#); [Gneezy et al., 2020](#); [Pace and van der Weele, 2020](#)).

⁴The factors that we hold constant are important because changing the structure of the decision (e.g., going from a passive to active choice) may drive avoidance due to attention channels; changing the content of the information (e.g., varying whether it relates to own payoffs of others' payoffs) may drive avoidance due to subjects valuing different information content differently; and changing the timing of the information (e.g., whether it is revealed before or after a payoff choice is made) may drive information avoidance by varying the relevance and value of information. For empirical evidence showing that our identification strategy does not introduce other drivers of information avoidance, see Section 4.4.

option that benefits oneself, since no option benefits oneself.⁵

We run four studies, including 4,626 experimental subjects. In each of these studies, we replicate the results of Dana, Weber and Kuang (2007). In each of these studies, we also find that a subset of subjects indeed avoid information due to image concerns. In particular, in each of our four studies, there is significantly more information avoidance in the *Self/Other* condition than in the *Other/Other* condition.

Across our studies, however, we find that, at most, 19%–34% of information avoidance in the classic paradigm is due to image concerns. The 66%–81% residual is substantial, suggesting that the large majority of information avoidance in the classic paradigm cannot be attributed to image concerns. The remaining information avoidance arises for other reasons, potentially including a desire to avoid interpersonal tradeoffs, a desire to avoid learning bad news (e.g., that you cannot achieve your preferred payoffs), laziness, inattention, and confusion.⁶

We explore these other causes of information avoidance by subtly changing the payoffs and the choice architecture of the information avoidance decision. In additional treatments, we find that a fraction of information avoidance is due to a desire to avoid interpersonal tradeoffs or to avoid learning “bad news” about the state.⁷ In a setting where we have removed image reasons to avoid information—as well as concerns about making interpersonal tradeoffs or learning bad news—we find that an active choice frame substantially mitigates information avoidance, suggesting that some of the information avoidance in the classic paradigm reflects laziness, inattention, and confusion—*independent of how these factors may interact with image concerns.*⁸

The main contribution of this paper is our ability to directly test whether individuals strategically avoid information *because of* image concerns. We are able to do this by introducing a control condition—the *Other/Other* condition—that removes image reasons to avoid information but holds constant the choice architecture, the content of information, and the timing of information provision. Given the prevalence of information avoidance across domains and the

⁵Section 2.1 further details why such image concerns cannot drive information *avoidance* in the *Other/Other* condition, even if they cause participants to acquire information and influence whether they choose A or B.

⁶For example, participants may have a direct non-image-related preference to avoid interpersonal tradeoffs; prefer the payoffs they can achieve in one state and dislike learning that they are in their less-preferred state; be too lazy to click to acquire information; be so inattentive that they avoid information by randomly choosing Option A or Option B directly; or be so confused that they do not understand the value of acquiring the information. We provide evidence for these channels in what follows and hope future work will further narrow in on these (or other specific) channels.

⁷The finding that individuals may avoid information in the Dana, Weber and Kuang (2007) paradigm due to an aversion to learning bad news relates to the work on information avoidance in health, financial, and individual performance domains (Eil and Rao, 2011; Thornton, 2008; Karlsson, Loewenstein and Seppi, 2009; Sicherman et al., 2016; Oster, Shoulson and Dorsey, 2013; Mobius et al., 2014; Ganguly and Tasoff, 2017; Golman et al., 2020), which often appeal to this type of explanation for information avoidance.

⁸In this setting, we build off of the excellent work showing how choice architecture reduces information avoidance (Grossman, 2014). We add to this work by testing how choice architecture affects information avoidance only after we have removed the role of image concerns in driving avoidance, allowing us to estimate the effect of choice architecture on avoidance absent any effects that might work through image concerns.

many rich lines of research on information avoidance, we see the use of such a control condition as an important methodological advance that could be applied more widely.

To see why the control condition is necessary, consider the common approach in the prior literature for assessing the role of image concerns in driving information avoidance. It compares the rate of information avoidance when payoffs are unknown to the rate of *selfishness* when payoffs are known. The latter represents the fraction of subjects who may avoid information because they do not value it (since they will act selfishly regardless). But this difference does not identify the extent of information avoidance that is due to image concerns. First, selfish subjects could avoid information because they do not value it, as suggested above, or they could avoid information strategically to mitigate the image costs of their selfishness. Second, non-selfish subjects may avoid information for *non-image* reasons—such as laziness, inattention, or confusion—rather than image reasons. This is true even if these subjects end up acting selfishly when uninformed, and even if they do so because of decreased image costs; they could have avoided information for a non-image reason and then been happily surprised by the opportunity to benefit themselves without knowing for certain they were being selfish. We further discuss these points, and explain why other prior work is also unable to identify the role of image concerns on information avoidance, in Section 2.3.

We generate three main empirical takeaways. The first takeaway is evidence in support of the popular hypothesis that image concerns do indeed drive some information avoidance. This finding reinforces the proposed explanations in the extant literature, including the model and empirical evidence in support of the model presented in Grossman and van der Weele (2017).

The second takeaway is that the amount of information avoidance that can be attributed to strategic image considerations may be misestimated absent a control condition such as our *Other/Other* condition. As shown in Section 3.2, our direct test finds that significantly less information avoidance is due to strategic image concerns than one would have assumed based on the common approach in the literature discussed above. In particular, our approach estimates the role of strategic image concerns to be less than half of what the common approach would have suggested in our setting.

The third takeaway is that the large majority of information avoidance cannot be attributed to strategic image concerns. This takeaway matters for three reasons. First, it makes clear the value of considering additional explanations for information avoidance beyond standard payoff motives and beyond belief management and signaling motives. Second, it provides evidence in support of one class of reasons why individuals do not suffer as big an image signaling cost when avoiding information as when behaving selfishly in the presence of information (Conrads and Irlenbusch, 2013; Krupka and Weber, 2013; Bartling, Engl and Weber, 2014; Grossman and van der Weele, 2017). Since individuals avoid information when avoidance cannot be driven by selfish motives (as in our *Other/Other* condition), there must be factors beyond selfishness

that drive information avoidance, making it a weaker signal of selfishness than selfishness when informed. Third, it matters for policy. Interventions to encourage prosocial behaviors that assume individuals strategically avoid information might attempt to have individuals make their information acquisition choices before they realize that the information is designed to encourage prosocial behavior (Exley and Petrie, 2018). However, if most information avoidance is not related to this strategic motive, interventions that tackle other drivers of information avoidance—such as those that focus on financial incentives to acquire information (Cain and Dana, 2012; van der Weele, 2014; Feiler, 2014; Grossman and van der Weele, 2017; Serra-Garcia and Szech, 2020)—may prove more effective.⁹

We build off of the Dana, Weber and Kuang (2007) paradigm, and we replicate its findings and the findings of the literature that follows. That prior literature provides compelling evidence that the ability to act selfishly without knowing that an act was selfish facilitates more selfish behavior. To examine the extent to which the ability to avoid information influences selfish behavior, those prior studies have exactly the right set of treatments: one where information can be avoided and one where information cannot be avoided. We pursue a different identification approach because we are interested in a different question. We study *why* individuals avoid information, rather than the consequences of information avoidance.

2 Design

This section describes the design of our main treatment conditions. Additional treatment conditions are introduced later. All instructions and decision screens are shown in Appendix A.

2.1 Overview of our Identification Strategy

A high-level view of our identification strategy is as follows. Our paradigm is built off of one of the experiments in Dana, Weber and Kuang (2007).¹⁰ In that experiment, and in the long line of papers that follow, a decision maker chooses between two payoff options that determine the payoffs for themselves and someone else. We call this original version of the game the *Self/Other* condition to emphasize that the decision maker determines the payoff for themselves (i.e., *Self*) and for another participant (i.e., *Other*). To explore the role of image concerns in driving information avoidance, we design a new version of this game in which image concerns can no longer drive information avoidance. In this new version, the decision maker chooses between two options that determine the payoffs for two other participants. We call our new version of the

⁹Policymakers seeking to influence behavior may also want make information unavoidable. Even when information cannot be avoided, however, how individuals process that information may depend on their awareness of incentives to favor certain beliefs (Babcock et al., 1995; Gneezy, Saccardo and van Veldhuizen, 2018; Schwardmann, Tripodi and van der Weele, 2019; Gneezy et al., 2020; Saccardo and Serra-Garcia, 2020).

¹⁰We use their *Baseline* treatment and their *Hidden Information* treatment (with baseline or alternative payoffs) in our design. In what follows, however, we adopt slightly different terminology than their paper. We maintain the *Hidden Information* treatment terminology. We refer to their “Baseline” treatment as the *Known Information* treatment, their “baseline” payoffs as occurring in the *Unaligned* state, and their “alternative” payoffs as occurring in the *Aligned* state. We will also refer to their Players X and Y as Players 1 and 2, respectively.

game the *Other/Other* condition to emphasize that decisions only influence the payoffs of other participants.

The *Other/Other* condition removes selfish motives, which means image concerns about selfishness cannot drive information avoidance. The removal of selfish motives also prevents image concerns unrelated to selfishness—such as a desire to appear fair—from driving information avoidance in the *Other/Other* condition. To see this, note that a participant in the *Self/Other* condition may avoid information to avoid facing a tradeoff between appearing fair and money for themselves. A participant in the *Other/Other* who values appearing fair does not face this tradeoff between financial incentives and image concerns. This participant can simply acquire the information and then choose the fair outcome. Consequently, while image concerns may cause participants to acquire information and influence whether participants choose A or B in the *Other/Other* condition, image concerns cannot cause participants to *avoid* information in the *Other/Other* condition.¹¹

Moreover, since the choice architecture, the content of the information, and the timing of information provision are all the same across the *Self/Other* and *Other/Other* conditions, we attribute any difference in information avoidance between these two conditions—and, specifically, the extent to which information avoidance is higher in the *Self/Other* condition—as being a consequence of image concerns. See Section 4.4 for a further discussion and empirical validation of this approach.

2.2 Details of our Design

A subject chooses between two options: Option A and Option B. The two options determine payoffs for two players, Player 1 and Player 2. While Player 1 always earns more from Option A than from Option B, the payoff for Player 2 depends on the “state.” Player 2 earns more from Option A in the *Aligned* state but earns more from Option B in the *Unaligned* state. Thus, in the unaligned state (and only the unaligned state), the subject faces a tradeoff in terms of benefiting Player 1 or benefiting Player 2.

The conditions under which a subject chooses between Option A and Option B vary according to the experimental treatment. In particular, in our Study 1, subjects are randomly assigned to:

1. the *Aligned* or *Unaligned* state,
2. the *Hidden Information* or *Known Information* condition, and

¹¹These arguments apply to many types of image concerns discussed in the prior literature that broadly relate to other-regarding preferences (e.g., such as a desire to appear nice or fair). One could, of course, imagine other types of image concerns that might drive information avoidance in the *Other/Other* condition. For instance, if one wants to signal that they *do not care* about the payoffs of others—and by “do not care” we mean signal indifference, since a desire to harm others would also push towards acquiring information—they may want to avoid information in both the *Other/Other* and *Self/Other* conditions. But image concerns like these are not posited in the prior literature and seem unlikely to be empirically relevant.

3. the *Self/Other* or *Other/Other* condition.

How choices map to payoffs depends on the random assignment in (1). As shown in the top panel of Table 1, the higher payoff for Player 1 is always from Option A, while the higher payoff for Player 2 depends on the state. We call these the “Classic Payoffs,” as we adopt the same payoff structure as the classic experiment in [Dana, Weber and Kuang \(2007\)](#).¹² In later studies, we alter the payoffs for some treatments. These payoffs are shown in the middle and bottom panels of Table 1 and will be described in detail when we discuss those studies.

Table 1: Payoffs for (Player 1, Player 2)

Classic Payoffs with Online Participants (used in Studies 1, 2, and 3)		
	Unaligned State	Aligned State
Option A	(\$0.60, \$0.10)	(\$0.60, \$0.50)
Option B	(\$0.50, \$0.50)	(\$0.50, \$0.10)
Classic Payoffs with Penn Undergraduates (used in Study 4)		
	Unaligned State	Aligned State
Option A	(\$6, \$1)	(\$6, \$5)
Option B	(\$5, \$5)	(\$5, \$1)
New Payoffs with Online Participants (used in Studies 2 and 3)		
	Aligned State 1	Aligned State 2
Option A	(\$0.50, \$0.10)	(\$0.50, \$0.50)
Option B	(\$0.50, \$0.50)	(\$0.50, \$0.10)

Each cell denotes the payoffs given to (Player 1, Player 2) according to whether Option A or Option B is chosen by the decision maker and to the state. In the *Self/Other* condition, the decision maker knows that Players 1 and 2 are themselves and another participant, respectively. In the *Other/Other* condition, the decision maker knows that Players 1 and 2 are two other participants.

How information on payoffs is presented depends on the random assignment in (2). In the *Known Information* condition, subjects are directly informed of the state and the associated payoffs and are asked to choose between Option A and Option B directly. By contrast, in the *Hidden Information* condition, subjects are informed of how the payoffs depend on the state and are informed that there is an equal chance that they have been assigned to either state. They are then asked whether they would like to: (i) choose Option A, (ii) choose Option B, or (iii) reveal which state they are in before choosing between Option A and Option B. While subjects

¹²As discussed in Section 2.4, our experimental subjects are recruited from Amazon’s Mechanical Turk, which has lower average hourly wages than the typical on-campus experimental lab. The lower wage and our desire to recruit many subjects for our study versions suggested we use smaller payoffs than the original study. We choose to divide the payoffs by 10 (e.g., paying \$0.60 rather than \$6 as in [Dana, Weber and Kuang \(2007\)](#)). In Section 3.3, we show that we replicate our results in a traditional on-campus experimental lab with same monetary payoffs as in [Dana, Weber and Kuang \(2007\)](#).

avoid information if they choose (i) or (ii), subjects acquire information (i.e., reveal the state) if they choose (iii).

Whether the information avoidance in the *Hidden Information* condition may be driven by image concerns depends on their random assignment in (3). In the *Self/Other* condition, subjects know that they earn the Player 1 payoffs and another participant earns the Player 2 payoffs, implying that Option A always benefits themselves. In the *Other/Other* condition, subjects know that two other participants earn the Player 1 and Player 2 payoffs, implying that neither option benefits themselves.

2.3 Why is the *Other/Other* condition necessary?

As detailed in the prior sections, we assess the extent to which information avoidance may be driven by image concerns by comparing the rates of information avoidance in the *Self/Other* condition to the *Other/Other* condition.

One may wonder whether we could have instead inferred the relevance of image concerns using data from the *Hidden Information* and *Known Information* conditions of the *Self/Other* condition only. Indeed, prior work often compares the rate of information avoidance in the *Hidden Information* condition to the rate of selfishness in the unaligned state in the *Known Information* condition. This prior work consistently finds that information avoidance is more common than selfishness and suggests that this difference between information avoidance and selfishness could be due to image concerns. This approach has—importantly—raised the debate about the motives for information avoidance. But, as noted in the Introduction, there are two reasons why this difference does not identify the role of image concerns in driving information avoidance.

The first is that individuals may avoid information because of image concerns even in settings when the information would not affect their choice. Consider an agent who always makes the most selfish choice (i.e., chooses Option A). She may still decide to avoid information in the *Hidden Information* condition to appear less selfish, even though it does not change her behavior. Assuming that the difference (i.e., between information avoidance and selfishness) is due to image concerns ignores this possibility and could underestimate the extent to which image concerns drive information avoidance.¹³

The second is that individuals who avoid information—and behave more selfishly as a result—in the *Hidden Information* condition could do so for reasons unrelated to image concerns. An agent who acts generously in the *Known Information* condition but who avoids information and acts selfishly in the *Hidden Information* condition may not have avoided information because of

¹³Note that it may be important to identify whether selfish individuals avoid information for image reasons. If selfish individuals suffer image costs that they wish to mitigate through information avoidance, they will prefer environments where they can avoid information to those where they are informed (even if they choose to behave selfishly in both cases). Such individuals may seek out opportunities to avoid information or avoid information even when doing so comes at a financial cost.

image concerns. She could have decided to avoid information for other reasons (e.g., inattention, confusion, laziness, a desire to avoid interpersonal tradeoffs, a desire to avoid bad news) but then acted selfishly (perhaps even because of the decreased image concerns that came from being uninformed). Assuming that the difference (i.e., between information avoidance and selfishness) is due to image concerns ignores this possibility and could overestimate the extent to which image concerns drive information avoidance.

In addition, one may wonder whether other approaches in the prior literature have been able to isolate the impact of image concerns on information avoidance. Prior work has compellingly shown that rates of information avoidance can be affected by choice architecture, the content of information, and the timing of information provision. But this work does not isolate the role of image concerns in driving information avoidance. Changing the choice architecture, such as making the information acquisition an active choice, could matter for image reasons—if avoiding information under an active choice is a stronger signal of selfishness—or for non-image reasons, such as laziness, inattention, and confusion.¹⁴ Changing the content of the information, as frequently occurs when considering avoidance across contexts, could matter for image reasons—if image concerns are content-dependent—or because subjects differentially value different types of information.¹⁵ Changing the timing of information provision, such as providing it after the relevant decision has been made, could affect information avoidance for image reasons—if image costs of learning information are different when the information cannot be instrumental—or because of associated changes in choice architecture and the value of information.¹⁶ An important feature of our approach is that our control condition (i.e., the *Other/Other* condition) holds constant the choice architecture, the content of the information, and the timing of information provision while varying whether image concerns can drive information avoidance.

2.4 Implementation Details

Subjects were recruited on Amazon Mechanical Turk to complete Study 1 in July 2019.¹⁷ We recruited 800 subjects, and approximately 100 were randomly assigned to each of the eight treatment conditions described above (i.e., resulting from random assignment to: (i) the *Aligned* or *Unaligned* state, (ii) the *Hidden Information* or *Known Information* condition, and (iii) the

¹⁴After documenting that requiring an active choice about information avoidance substantially reduces information avoidance, Grossman (2014) notes that one possible explanation—that should be examined in future work—is that the “active pursuit of ignorance” could be viewed as “more inappropriate socially.” Moreover, prior work finds that individuals may use the possibility of inattention and confusion as excuses to behave selfishly (Exley and Kessler, 2019).

¹⁵However, as made clear in Serra-Garcia and Szech (2020), variations like these can be quite informative for other purposes, such as examining the elasticity of information avoidance.

¹⁶However, as made clear in Grossman and van der Weele (2017), variations like these can be quite informative for other purposes, such as providing empirical evidence in support of models with image concerns.

¹⁷To be eligible for any of our Amazon Mechanical Turk studies, subjects needed to have a 95% approval rating or better from at least 100 prior HITs and a US IP addresses. For this 5-minute study, they received a 50 cent completion fee plus any additional payoffs detailed above.

Self/Other or *Other/Other* condition). As is described in what follows, we twice replicated the results from Study 1 by recruiting, for the same eight treatment conditions, an additional 807 subjects in September 2019 (as part of Study 2) and an additional 796 subjects in February 2020 (as part of Study 3). These results from Studies 1–3 are detailed next in Sections 3.1 and 3.2. In Section 3.3, we describe another replication with 222 undergraduate students from the University of Pennsylvania (as part of Study 4).¹⁸

Prior to making any decision, subjects receive detailed instructions and must correctly answer understanding questions that require them to report that Player 1 payoffs are higher for Option A than for Option B and that Player 2 payoffs are higher for either Option A or Option B (and that this depends on the state if assigned to the *Hidden Information* condition). To mitigate the relevance of direct reciprocity concerns, subjects are informed that they are randomly assigned into groups of three participants, that one member of their group will be randomly selected as the decision maker, and that only the choice of the decision maker will determine additional payoffs for the group. That is, the decision maker determines additional payoffs for themselves and another group member in the *Self/Other* condition or for both of their other group members in the *Other/Other* condition. After subjects make their decisions, they fill out a short demographic survey. See Appendix A for full experimental instructions and decision screens.

3 Results

In this section, we present results from Study 1 and the treatments from Studies 2–4 that replicate Study 1. Specifically, in Section 3.1, we show that our *Self/Other* condition replicates the original results from Dana, Weber and Kuang (2007), demonstrating that our online sample yields the same results as prior literature. In Section 3.2, we compare behavior in our *Self/Other* and *Other/Other* conditions to present new results on the extent to which information avoidance is driven by image concerns. We find that some subjects avoid information because of image concerns but that this represents only a small portion of the information avoidance in the *Self/Other* condition and substantially less than what the common approach from prior literature would infer from our data. In Section 3.3, we show that we replicate the findings from Section 3.2 in a traditional undergraduate subject pool with a design that involves higher payoffs.

3.1 Replicating the original moral wiggle room findings

In this section, we focus on results from the *Self/Other* condition and show they are consistent with results from Dana, Weber and Kuang (2007). We report on data from all three of our online studies.

Consistent with prior literature, we find a large fraction of subjects avoid information and that this fraction exceeds the rate of selfishness when information is known. Across the three

¹⁸In Section 4, we present additional design details and results, including treatment variations from Studies 2 and 3, involving another 2,003 subjects. In total, 4,626 subjects participated in our four studies.

studies, 0.67 (Study 1), 0.72 (Study 2), and 0.65 (Study 3) of subjects in the *Hidden Information* condition avoid information. Across the three studies, 0.32 (Study 1), 0.33 (Study 2), and 0.33 (Study 3) of subjects choose Option A—the selfish option—in the unaligned state of the *Known Information* condition.

Also replicating prior literature, we find that the ability to avoid information leads to more selfish behavior. As shown in Table 2, which focuses on results from the unaligned state, the rates of choosing Option A increase by at least 20 percentage points in the *Hidden Information* condition when compared to the *Known Information* condition. Across the three studies, 0.56 (Study 1), 0.59 (Study 2), and 0.53 (Study 3) chose Option A in the unaligned state of the *Hidden Information* condition.

Table 2: Linear probability model of the likelihood of choosing Option A

	Study 1	Study 2	Study 3
<i>Hidden Information</i>	0.23*** (0.07)	0.27*** (0.07)	0.20*** (0.07)
Constant	0.33*** (0.05)	0.32*** (0.05)	0.33*** (0.05)
N	199	200	200

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are robust and shown in parentheses. The results are from a linear probability model of the likelihood to choosing Option A, where *Hidden Information* is an indicator for being the *Hidden Information* condition. In all columns, the data are restricted to the decisions made in the *Unaligned* state of the *Self/Other*, *Hidden Information* condition or the *Unaligned* state of the *Self/Other*, *Known Information* condition.

These findings match the results from the prior literature discussed in the Introduction, demonstrating that in all three of these studies we successfully replicate Dana, Weber and Kuang (2007), and the literature that follows, in our online experimental paradigm.

3.2 Do individuals avoid information because of image concerns?

When considering results from the *Self/Other* condition in Studies 1–3, the prior section found high rates of information avoidance and found large increases in the rate of choosing Option A when information could be avoided. To what extent can this be explained by subjects in the *Self/Other* condition avoiding information *because of* image concerns? To answer this question, we compare rates of information avoidance in the *Self/Other* condition to the *Other/Other* condition in which image concerns cannot drive information avoidance. In this section, we discuss results from our online Studies 1–3. In the next section, we discuss parallel results from Study 4 involving Penn undergraduate students.

Table 3 shows results from all of our *Hidden Information* conditions. It presents a linear probability model of whether subjects avoid information on an indicator for whether subjects are randomly assigned to the *Other/Other* condition. Four main results follow.

Table 3: Linear probability model of the likelihood of avoiding information

	Study 1	Study 2	Study 3	Study 4
Other/Other	-0.13*** (0.05)	-0.17*** (0.05)	-0.14*** (0.05)	-0.21*** (0.07)
Constant	0.67*** (0.03)	0.72*** (0.03)	0.65*** (0.03)	0.62*** (0.05)
N	397	399	386	222

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are robust and shown in parentheses. The results are from a linear probability model of avoiding information, where *Other/Other* is an indicator for being the *Other/Other* condition. In all columns, the data are restricted to the decisions made in the unaligned or aligned state of the *Hidden Information* condition. In columns 1, 2, 3, and 4, the data are restricted to the decisions made in Study 1, 2, 3, or 4, respectively.

First, the coefficient estimates on the constant show the rates of information avoidance in the *Self/Other* conditions. As noted in the prior section, these rates of information avoidance are high (0.67 in Study 1, 0.72 in Study 2, and 0.65 in Study 3).

Second, as shown by the significant negative coefficient estimates on a dummy for being in the *Other/Other* condition, we document significantly less information avoidance when image concerns cannot drive such avoidance. Compared to the *Self/Other* condition, information avoidance in the *Other/Other* conditions falls by 13 percentage points in Study 1, 17 percentage points in Study 2, and 14 percentage points in Study 3. This implies that the minority of information avoidance in the *Self/Other* condition is due image concerns. The percentage due to image concerns is $\frac{0.13}{0.67} = 19\%$ in Study 1, $\frac{0.17}{0.72} = 24\%$ in Study 2, and $\frac{0.14}{0.65} = 22\%$ in Study 3.

Third, our identification strategy suggests that a smaller fraction of information avoidance is due to image concerns than we would have guessed if we had subtracted the rate of choosing the selfish option in the unaligned state of the *Known Information* condition from the rate of information avoidance and assumed that the difference represented subjects who were motivated by image concerns (i.e., the approach discussed in Section 2.3). In the *Self/Other* condition, the fraction of participants who avoid information in the *Hidden Information* condition minus the fraction of participants who choose Option A in the unaligned state of the *Known Information* condition is $0.67 - 0.33 = 0.34$ (Study 1), $0.72 - 0.32 = 0.40$ (Study 2), and $0.65 - 0.32 = 0.33$ (Study 3). These calculations would imply that image concerns account for a majority of information avoidance: $\frac{0.34}{0.67} = 51\%$ in Study 1, $\frac{0.40}{0.72} = 56\%$ in Study 2, and $\frac{0.33}{0.65} = 51\%$ in Study 3. These percentages are all over twice as large as the percentages of information avoidance that we attribute to image concerns in the prior paragraph. Thus, not only is our comparison of the *Self/Other* to *Other/Other* conceptually different than this alternative approach, it is a difference that proves empirically important.

Fourth, a large fraction of subjects avoid information in the *Other/Other* condition. As seen by adding the coefficient estimates on the constant and *Other/Other* in Table 3, or as shown

directly in the “Classic *O/O*” column of Table 4, these rates are 0.54 in Study 1, 0.55 in Study 2, and 0.51 in Study 3. These high rates of information avoidance in the *Other/Other* condition highlight that the large majority of information avoidance observed in the *Self/Other* condition is *not* due to image concerns: $\frac{0.54}{0.67} = 81\%$ in Study 1, $\frac{0.55}{0.72} = 76\%$ in Study 2, and $\frac{0.51}{0.65} = 78\%$ in Study 3.

In light of the large share of information avoidance that is not due to image concerns, we consider additional drivers of information avoidance in Section 4.

3.3 Do our results replicate with undergraduate students?

Before we explore the causes of the information avoidance that is not due to image concerns, we turn to whether our results from Studies 1–3 are driven by the subjects being recruited from Amazon Mechanical Turk. First, we emphasize that many concerns about the subjects being “different” than the traditional undergraduate subject pool are alleviated by us replicating all three prior findings on information avoidance and the impact of information avoidance from Dana, Weber and Kuang (2007) and the literature that follows (as shown in Section 3.1). Second, we emphasize that other concerns about our results are alleviated by us replicating our results from Study 1 twice (i.e., in Study 2 and Study 3). Questions around the attentiveness of our subjects are further addressed in Section 4.3.

Nonetheless, to more directly assess whether our results replicate with undergraduate students, we ran Study 4, which recruited 222 undergraduate students from the Wharton Behavioral Lab at the University of Pennsylvania. Experimental instructions and decision screens for Study 4 are shown in Appendix A.4. There are two main differences with the design of Study 4 relative to the design detailed in Section 2 for Studies 1–3. First, as shown in the middle panel of Table 1, we increased the value of payoffs to match the typical values used in this literature for undergraduate student subjects. Second, given the limited subject pool size, all subjects were assigned to one of the *Hidden Information* conditions (i.e., we excluded the *Known Information* conditions).¹⁹

As shown in the Study 4 column of Table 3, our results persist. Information avoidance is substantial when subjects have a selfish motive to choose Option A. The coefficient estimates on the constant show that 0.62 of subjects avoid information in the *Self/Other* condition. In addition, the coefficient on *Other/Other* is large, negative and statistically significant ($p < 0.01$); information avoidance is 21 percentage points less likely when image concerns cannot drive avoidance. As can be seen from adding these coefficients, or directly in the “Classic *O/O*” column of Table 4, 0.41 of subjects avoid information in the *Other/Other* condition.

We replicate both that there is information avoidance due to image concerns and that the

¹⁹That said, to get a sense of the results in the *Known Information* conditions, we ask all subjects two questions from the *Known Information* condition (i.e., what they would do in the unaligned state and what they would do in the aligned state) after they complete the *Hidden Information* condition.

majority of information avoidance persists in the *Other/Other* condition when image concerns cannot drive information avoidance. Based on our estimates, $\frac{0.21}{0.62} = 34\%$ of information avoidance from the *Self/Other* condition is attributable to image concerns and $\frac{0.41}{0.62} = 66\%$ arises for other reasons.

4 Discussion and Additional Results

We find that the majority of information avoidance in the classic paradigm cannot be explained by image concerns. As detailed in our Introduction, this finding is important to our understanding of information avoidance when considering motivated reasoning, and it is relevant for policy.

What drives this remaining information avoidance? The following subsections report on additional treatments from Study 2 and 3, and additional data from all of our studies, to explore the drivers of this information avoidance. Table 4 summarizes the rates of information avoidance across all of our *Hidden Information* conditions in all of our studies. The results shown in the first two columns, under “Classic,” are the results we have already discussed in Section 3. Those under the *S/O* and *O/O* columns refer to results from the *Self/Other* and *Other/Other* conditions, respectively. Results from the additional columns will be discussed in what follows.

Table 4: Fraction avoiding information in *Hidden Information* conditions

Payoffs:	Classic		New		New, Active Choice	
	<i>S/O</i>	<i>O/O</i>	<i>S/O–New</i>	<i>O/O–New</i>	<i>S/O–Active</i>	<i>O/O–Active</i>
Study 1	0.67	0.55
Study 2	0.72	0.55	0.44	0.43	.	.
Study 3	0.65	0.52	0.47	0.45	0.25	0.20
Study 4	0.62	0.41
N	698	706	400	391	199	197

The first pair of columns involve the “Classic Payoffs” shown in the top panel of Table 1. The middle and last pair of columns involve the “New Payoffs” shown in the bottom panel of Table 1. The last pair of columns involve treatments where participants must actively choose whether or not to acquire information before having the ability to choose Option A or Option B. Within each pair of columns, results are split according to whether participants were randomly assigned to one of the conditions involving payoffs for themselves and another participant (i.e., the *Self/Other*, *Self/Other–New*, or *Self/Other–Active* condition) or one of the conditions involving payoffs for two other participants (i.e., the *Other/Other*, *Other/Other–New*, or *Other/Other–Active* condition).

4.1 Considering additional explanations: tradeoff aversion and bad news aversion

Why might individuals avoid information in the *Other/Other* condition even though they do not have any monetary payoffs at stake? Two initial explanations may come to mind. First, participants may be averse to making decisions that involve tradeoffs between people—even if

those tradeoffs do not affect themselves—and so may not want to be put into a position (like in the unaligned state) where they have to make a tradeoff between the two participants. Second, participants may favor the payoffs they can achieve in one of the two states and thus prefer that state. Subjects may want to believe that they are in their preferred state to avoid learning the “bad news” that they are actually in their less-preferred state. How a desire to avoid bad news can generate information avoidance is further developed and discussed in [Golman, Hagmann and Loewenstein \(2017\)](#) and [Golman and Loewenstein \(2018\)](#).

To what extent do these two explanations drive the residual information avoidance observed in the *Other/Other* condition? After observing the large residual in the *Other/Other* condition of Study 1, we introduced new treatment conditions in Study 2 and Study 3 to answer this question. As shown in the bottom panel of Table 1, the “New Payoffs” are the same as the “Classic Payoffs” except for one change: Option A gives \$0.50, rather than \$0.60 cents, to Player 1. This change implies that the payoffs for the two players are always (weakly) aligned, eliminating concerns about tradeoff aversion, and that the two states are identical in what payoffs can be achieved, eliminating concerns that individuals may prefer one state to the other.²⁰

Consistent with an aversion to tradeoffs or to bad news, Table 4 shows that the rates of information avoidance are 7–12 percentage points lower with the new payoffs (compare the rates in the *Other/Other–New* condition with “New” payoffs to those in the *Other/Other* condition with “Classic” payoffs). These differences are statistically significant in Study 2 (0.55 vs. 0.43, $p < 0.01$) but only suggestive in Study 3 (0.52 vs. 0.45, $p = 0.23$). Combining data from Study 2 and 3 yields a significant difference (0.54 vs. 0.44, $p < 0.01$).

4.2 Considering additional explanations: how information is acquired

Jointly considering our main results and the additional results in Section 4.1 shows that substantial information avoidance cannot be attributed to image concerns, an aversion to making interpersonal tradeoffs, or the prospect of learning bad news. What can explain the remaining information avoidance?

We explore the possibility that inattention, confusion, and laziness could be driving the residual information avoidance by introducing a treatment that is likely to reduce the role of these three potential drivers. After observing the information avoidance among subjects facing the “New Payoffs” in Study 2, we introduced an *Active Choice* version of the *Hidden Information* condition in Study 3. In this version, subjects again face the “New Payoffs,” but prior to choosing Option A or Option B, subjects first have to actively choose whether to reveal or not reveal the state to which they are assigned (see the screenshot in Appendix Figure A.38).

As compared to the standard *Hidden Information* condition, the *Active Choice* version may reduce information avoidance for reasons surrounding confusion, inattention, and laziness. The

²⁰Subjects facing the “New Payoffs” are randomly assigned to either: *Aligned State 1* or *Aligned State 2*; the *Hidden Information* or *Known Information* condition; and the *Self/Other–New* or *Other/Other–New* condition.

active choice version makes the decision simpler, so confused subjects might better understand the value of revealing information in the active choice version. Inattentive subjects, such as subjects who choose somewhat randomly, should be less likely to avoid information in the active choice version where 1 of 2 options reveal information, rather than 1 of 3 in the standard version. Lazy subjects who avoided information in the standard version—by choosing Option A or B directly—to avoid having to click to a new screen and otherwise think more about the decision should be less likely to avoid information in the active choice version since they have no way of skipping the subsequent decision screen.

The active choice version might decrease information avoidance through any of these channels, but it is an appealing intervention to test because, in the original [Dana, Weber and Kuang \(2007\)](#) paradigm, information avoidance has been shown to be sensitive to choice architecture. [Grossman \(2014\)](#) shows less information avoidance when subjects face an active choice about acquiring information or if information acquisition—rather than avoidance—is the default.

When image concerns can drive information avoidance, however, the choice architecture of the decision may directly influence the image cost associated with avoidance as well as affecting information avoidance through the channels noted above. We only introduce the active choice version with the new payoffs, once we have ruled out information avoidance being driven by image concerns, concerns about interpersonal tradeoffs, or the prospect of learning bad news. This allows us to add to this prior literature by specifically exploring the extent to which choice architecture affects information avoidance through confusion, inattention, and laziness (i.e., absent any effect choice architecture might have through image concerns, concerns about interpersonal tradeoffs, or the prospect of learning bad news).

This change in the choice architecture proves powerful in our setting. As seen by comparing the “New” columns of Study 3 to the “New, Active Choice” columns of Study 3 in [Table 4](#), information avoidance is substantially lower when an active choice is required (0.25 vs. 0.47, $p < 0.01$, in the *Self/Other* condition; and 0.20 vs. 0.45, $p < 0.01$, in the *Other/Other* condition).

That the active choice frame significantly reduces information avoidance points to a combination of motives for information avoidance that involve inattention, confusion, and laziness. We view these possibilities—and, in particular, how they may interact with image concerns about selfishness—as promising avenues to consider in future work.

4.3 Considering the role of information processing

While the prior section suggests a possible role of inattention, confusion, and laziness in driving information avoidance, results from the *Known Information* conditions suggest a natural limit to how much behavior might be motivated by inattention, confusion, or indifference.

[Table 5](#) shows the percentage of subjects who choose Option A in each condition. First, consider results with the classic payoffs. The top panel of the fourth column of [Table 5](#) shows that nearly all subjects—98% of subjects (pooling across all studies and conditions)—choose

Option A in the *Aligned* state of the *Known Information* condition. That is, nearly all subjects choose the weakly higher payoff option when they are directly informed of the payoff information, regardless of whether they are in the *Self/Other* or *Other/Other* condition.²¹

Table 5: Fraction choosing Option A

	Unaligned State		Aligned State	
	Hidden Information	Known Information	Hidden Information	Known Information
Classic Payoffs				
Study 1: <i>Self/Other</i>	0.56	0.33	0.87	1.00
Study 1: <i>Other/Other</i>	0.32	0.19	0.71	0.97
Study 2: <i>Self/Other</i>	0.59	0.32	0.84	0.98
Study 2: <i>Other/Other</i>	0.28	0.12	0.69	0.98
Study 3: <i>Self/Other</i>	0.53	0.33	0.78	0.95
Study 3: <i>Other/Other</i>	0.26	0.18	0.70	0.95
Study 4: <i>Self/Other</i>	0.73	.	0.96	.
Study 4: <i>Other/Other</i>	0.34	.	0.82	.
	Aligned State 2		Aligned State 1	
	Hidden Information	Known Information	Hidden Information	Known Information
New Payoffs				
Study 2: <i>Self/Other</i>	0.19	0.07	0.69	0.93
Study 2: <i>Other/Other</i>	0.11	0.15	0.69	0.98
Study 3: <i>Self/Other</i>	0.22	0.08	0.68	0.93
Study 3: <i>Other/Other</i>	0.14	0.11	0.71	0.91
New Payoffs with Active Choice				
Study 3: <i>Self/Other</i>	0.25	.	0.79	.
Study 3: <i>Other/Other</i>	0.29	.	0.86	.
N	1296	1018	1295	1019

The above results show the fraction of participants choosing Option A, according to the treatment condition to which they were assigned.

Second, consider results with the new payoffs. The middle panel of the fourth column of Table 5 shows similar results: 94% of subjects (pooling across studies and conditions) choose

²¹This is also true for 98%–99% of the additional decisions subjects make in Study 4—detailed in footnote 19—when they are asked about what they would do if assigned to the *Aligned* state condition. We do not include these additional decisions in Table 5 because they could be confounded by subjects in Study 4 first having made decisions in the *Hidden Information* condition.

Option A in *Aligned State 1* of the *Known Information* condition. This fraction is consistent across the *Self/Other* and *Other/Other* conditions. The second column of Table 5 shows that this pattern is not unique to choosing Option A. With the new payoffs, Option A in *Aligned State 1* of the *Known Information* condition is identical to Option B in *Aligned State 2* of the *Known Information* condition, and Option B in *Aligned State 2* of the *Known Information* condition is chosen 90% of the time (pooling across all studies and conditions).

These results suggest that, when asked directly about whether they want to select an option that benefits other participants, at most 10% of subjects decline to do so, either because they do not care about the payoff information, are choosing randomly, or fail to process the payoff information.

While the information avoidance decision may be more cognitively difficult than the direct choice of Option A or Option B in the *Known Information* conditions, it is clear from these *Known Information* choices that very few subjects are completely inattentive, confused, or indifferent.

4.4 Does the *Other/Other* condition induce additional motives for information avoidance?

As detailed in Section 2.1, the choice architecture, the content of information (i.e., the payoffs for Player 2 that depend on the state), and the timing of information provision are all identical in the *Self/Other* and *Other/Other* conditions. We hold these features constant and compare the rates of information avoidance when image concerns can cause subjects to avoid information in the *Self/Other* condition to the rates of information avoidance when image concerns cannot cause subjects to avoid information in the *Other/Other* condition. More specifically, the only change we make across conditions involves switching the Player 1 payoffs from the decision maker in the *Self/Other* condition to another subject in the *Other/Other* condition.

While we view this as a minimal change that ensures image concerns cannot drive information avoidance—and believe this methodology represents an important advance for the literature—one may wonder if we introduced additional, empirically-relevant motives to avoid information. For example, one may wonder if participants avoid information more in the *Other/Other* condition simply because they do not pay attention whenever the game does not involve payoffs for themselves. That we require participants to correctly answer understanding questions about the payoffs (see screenshot shown in Appendix Figure A.7) should help to guard against this possibility. Additionally, one may wonder if participants avoid information more in the *Other/Other* condition because the marginal value of the information is lower in the *Other/Other* condition, even though the content of the information is the same. However, since individuals systematically value money for themselves more than money for others (Exley, 2016; Exley and Kessler, 2019; Exley, 2020), we suspect that, if anything, the opposite is true.²²

²²In the *Self/Other* condition, subjects likely have a stronger preference over Player 1 payoffs (i.e., their money)

More fundamentally, results from our experiment provide empirical evidence against the possibility that simply switching the payoffs of Player 1 from the decision maker in the *Self/Other* condition to another subject in the *Other/Other* condition induces more information avoidance for reasons noted above or otherwise. First, consider the results from the *Self/Other–New* condition and the *Other/Other–New* condition. Comparing the information avoidance in these two conditions allows us to isolate the role of switching the payoffs of Player 1 from the decision maker to another subject in a setting where image concerns *cannot* drive information avoidance. As shown in the middle two columns of Table 4, there is no difference in information avoidance across these conditions. If anything, information avoidance is directionally *lower* in the *Other/Other–New* condition than the *Self/Other–New* condition in both Study 2 and Study 3. Second, consider the results from the *Self/Other–Active* condition and the *Other/Other–Active* condition. As shown in the right two columns of Table 4, the same pattern arises when comparing information avoidance across these conditions. If anything, information avoidance is directionally *lower* in the *Other/Other–Active* condition than the *Self/Other–Active* condition.

This section highlights that replacing payoffs for the decision maker with payoffs for another participant, as we do going from the *Self/Other* condition to the *Other/Other* condition, does not introduce additional information avoidance. The results presented here suggest, if anything, the extent to which image concerns drive information avoidance that we identify in our main results may be an upper bound.

5 Conclusion

Our experiment explores the extent to which information avoidance is driven by image concerns. We focus on the classic [Dana, Weber and Kuang \(2007\)](#) paradigm. We provide evidence of more information avoidance when image concerns could motivate information avoidance, highlighting that some subjects indeed avoid information because of image concerns. But, we also show how the common approach in the literature misestimates the role of image concerns in driving information avoidance and that the majority of information avoidance cannot be attributed to image concerns. We find evidence for other drivers of information avoidance, including a desire to avoid interpersonal tradeoffs, a desire to avoid bad news, laziness, inattention, and confusion.

Our exploration of information avoidance opens up additional questions for future work, two of which we note here. First, our results suggest that it might be worthwhile to revisit the relevance of both image-driven concerns and non-image-driven concerns in a range of other contexts in

and thus the information about Player 2 payoffs (i.e., someone else’s money) is likely to matter less for their decision making. Thus, if subjects acquire information less often in the *Other/Other* condition than in the *Self/Other* condition, this is unlikely to be because they value the information less in the *Other/Other* condition. Consequently, to the extent that one is worried that information has a different marginal value in the *Self/Other* and *Other/Other* conditions, the difference in information avoidance between the conditions likely provides an upper bound on the amount of information avoidance that can be driven by image concerns. This is why, unlike in our prior work ([Exley, 2016](#); [Exley and Kessler, 2019, 2020](#)), we do not believe it is necessary to calibrate payoff options across conditions so that they are equivalently valued by the decision maker.

which information avoidance is prevalent (see [Golman, Hagmann and Loewenstein \(2017\)](#) for an excellent review of information avoidance across contexts).²³ While we were surprised by the extent of information avoidance in the classic paradigm that could not be attributed to image concerns, we also suspect that there are many contexts where one may be surprised by the extent to which image concerns drive information avoidance. We hope future work jointly considers reasons related to image concerns and not related to image concerns to bolster our understanding of information avoidance and other avoidance decisions (e.g., when individuals avoid the ask or avoid making decisions).²⁴ Control conditions—such as our *Other/Other* condition—may prove of particular use to this future work that seeks to narrow in on the role of image concerns in driving information avoidance.²⁵

Second, and related, our work suggests gains from further exploring inattention, laziness, or confusion as potentially important drivers of information avoidance across a number of domains. It is possible that people rationally avoid information in response to problem complexity as in models of rational inattention and sparsity ([Sims, 2003](#); [Gabaix, 2014, 2017](#)), that they avoid information because they look at problems the wrong way (see [Handel and Schwartzstein \(2018\)](#) for an excellent review), or even that the ability to avoid information provides individuals with an “excuse” not to fully think through decisions. While we have shown that image concerns can explain part of the information avoidance in a classic paradigm, much information avoidance remains. We see great promise in exploring the other drivers of information avoidance across domains.²⁶

²³Interesting questions also remain on how individuals *seek* information (see, e.g., [Spiekermann and Weiss \(2016\)](#)) for image and non-image reasons.

²⁴For work related to how people avoid opportunities to be generous see, e.g., [Dana, Cain and Dawes \(2006\)](#); [Broberg, Ellingsen and Johannesson \(2007\)](#); [Jacobsen et al. \(2011\)](#); [DellaVigna, List and Malmendier \(2012\)](#); [Lazear, Malmendier and Weber \(2012\)](#); [Trachtman et al. \(2015\)](#); [Andreoni, Rao and Trachtman \(2016\)](#); [Lin, Schaumberg and Reich \(2016\)](#).

²⁵Of course, there are other ways to vary the role of image concerns. Experimentalists might add social-image concerns or enhance self-image concerns to investigate whether these changes increase information avoidance. We chose instead to use the *Other/Other* condition to be able to speak to the extent of image concerns in driving information avoidance in the classic setting, which requires us to remove image concerns from that setting.

²⁶More generally, in light of the findings in [van der Weele et al. \(2014\)](#) that suggest excuses for selfish behavior—and related image concerns—are less relevant in reciprocal contexts, many interesting questions remain about the conditions under which image concerns prove relevant.

References

- Andreoni, James, Justin M. Rao, and Hannah Trachtman.** 2016. “Avoiding the ask: A field experiment on altruism, empathy, and charitable giving.” *Journal of Political Economy*.
- Babcock, Linda, George Loewenstein, Samuel Issacharoff, and Colin Camerer.** 1995. “Biased Judgments of Fairness in Bargaining.” *The American Economic Review*, 85(5): 1337–1343.
- Bartling, Björn, and Yagiz Özdemir.** 2017. “The Limits to Moral Erosion in Markets: Social Norms and the Replacement Excuse.” *Working Paper*.
- Bartling, Björn, Florian Engl, and Roberto A. Weber.** 2014. “Does willful ignorance deflect punishment? – An experimental study.” *European Economic Review*, 70(0): 512 – 524.
- Bartoš, Vojtěch, Michal Bauer, Julie Chytilová, and Filip Matějka.** 2016. “Attention Discrimination: Theory and Field Experiments with Monitoring Information Acquisition.” *American Economic Review*, 106(6): 1437–1475.
- Bénabou, Roland, and Jean Tirole.** 2004. “Willpower and personal rules.” *Journal of Political Economy*, 112(4): 848–886.
- Bénabou, Roland, and Jean Tirole.** 2006. “Incentives and Prosocial Behavior.” *American Economic Review*, 96(5): 1652–1678.
- Bénabou, Roland, and Jean Tirole.** 2011. “Identity, Morals, and Taboos: Beliefs as Assets.” *Quarterly Journal of Economics*, 126: 805–855.
- Bénabou, Roland, and Jean Tirole.** 2016. “Mindful Economics: The Production, Consumption, and Value of Beliefs.” *Journal of Economic Perspectives*, 30(3): 141–164.
- Bénabou, Roland, Armin Falk, and Jean Tirole.** 2018. “Narratives, Imperatives and Moral Reasoning.” *Working Paper*.
- Bodner, Ronit, and Drazen Prelec.** 2003. “Self-signaling and diagnostic utility in everyday decision making.” *The psychology of economic decisions*, 1: 105–26.
- Broberg, Tomas, Tore Ellingsen, and Magnus Johannesson.** 2007. “Is generosity involuntary?” *Economics Letters*, 94(1): 32–37.
- Cain, Daylian M., and Jason Dana.** 2012. “Paying People to Look at the Consequences of Their Action.” *Working Paper*.

- Conrads, Julian, and Bernd Irlenbusch.** 2013. “Strategic ignorance in ultimatum bargaining.” *Journal of Economic Behavior and Organization*, 92(C): 104–115.
- d’Adda, Giovanna, Yu Gao, Russell Golman, and Massimo Tavoni.** 2018. “It’s so Hot in Here: Information Avoidance, Moral Wiggle Room, and High Air Conditioning Usage.” *Working Paper*.
- Dana, Jason, Daylian M. Cain, and Robyn M. Dawes.** 2006. “What you don’t know won’t hurt me: Costly (but quiet) exit in dictator games.” *Organizational Behavior and Human Decision Processes*, 100: 193–201.
- Dana, Jason, Roberto A. Weber, and Jason Xi Kuang.** 2007. “Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness.” *Economic Theory*, 33: 67–80.
- Danilov, Anastasia, and Silvia Saccardo.** 2016. “Disguised Discrimination.” *Working Paper*.
- DellaVigna, Stefano, John List, and Ulrike Malmendier.** 2012. “Testing for Altruism and Social Pressure in Charitable Giving.” *Quarterly Journal of Economics*, 127(1): 1–56.
- Di Tella, Rafael, Ricardo Perez-Truglia, Andres Babino, and Mariano Sigman.** 2015. “Conveniently Upset: Avoiding Altruism by Distorting Beliefs about Others’ Altruism.” *American Economic Review*, 105(11): 3416–42.
- Eil, David, and Justin M. Rao.** 2011. “The Good News-Bad News Effect: Asymmetric Processing of Objective Information about Yourself.” *American Economic Journal: Microeconomics*, 3(2): 114–138.
- Exley, Christine L.** 2016. “Excusing Selfishness in Charitable Giving: The Role of Risk.” *Review of Economic Studies*, 83(2): 587–628.
- Exley, Christine L.** 2020. “Using Charity Performance Metrics as an Excuse Not To Give.” *Management Science*, 66(2): 553–563.
- Exley, Christine L., and Jeffrey K. Naecker.** 2017. “Observability Increases the Demand for Commitment Devices.” *Management Science*, 63(10): 3262–3267.
- Exley, Christine L., and Judd B. Kessler.** 2019. “Motivated Errors.” *Working Paper NBER No. 26595*.
- Exley, Christine L., and Judd B. Kessler.** 2020. “Equity Concerns are Narrowly Framed.” *Working Paper, No. 25326, 20 NBER Working Paper No. 25326*.

- Exley, Christine L., and Ragan Petrie.** 2018. “The Impact of a Surprise Donation Ask.” *Journal of Public Economics*, 158(152-167).
- Falk, Armin, and Nora Szech.** 2013. “Morals and markets.” *Science*, 340(6133): 707–711.
- Falk, Armin, and Nora Szech.** 2017. “Diffusion of Being Pivotal and Immoral Outcomes.” *Working Paper*.
- Feiler, Lauren.** 2014. “Testing Models of Information Avoidance with Binary Choice Dictator Games.” *Journal of Economic Psychology*.
- Foerster, Manuel, and Joel J van der Weele.** 2018. “Denial and Alarmism in Collective Action Problems.” *Working Paper*.
- Freddi, Eleonora.** 2018. “Do People Avoid Morally Relevant Information? Evidence from the Refugee Crisis.” *Working Paper*.
- Gabaix, Xavier.** 2014. “A sparsity-based model of bounded rationality.” *The Quarterly Journal of Economics*, 129(4): 1661–1710.
- Gabaix, Xavier.** 2017. “Behavioral Inattention.” In *Handbook of Behavioral Economics: Applications and Foundations 1*. Vol. 2, 261–343. Elsevier.
- Ganguly, Ananda, and Joshua Tasoff.** 2017. “Fantasy and dread: The demand for information and the consumption utility of the future.” *Management Science*, 63(12).
- Garcia, Thomas, Sébastien Massoni, and Marie Claire Villeval.** 2020. “Ambiguity and excuse-driven behavior in charitable giving.” *European Economic Review*, 124.
- Gino, Francesca, Michael I. Norton, and Roberto A. Weber.** 2016. “Motivated Bayesians: Feeling Moral While Acting Egoistically.” *Journal of Economic Perspectives*, 30(3): 189–212.
- Gneezy, Uri, Silvia Saccardo, and Roel van Veldhuizen.** 2018. “Bribery: Behavioral Drivers of Distorted Decisions.” *Journal of the European Economic Association*.
- Gneezy, Uri, Silvia Saccardo, Marta Serra-Garcia, and Roel van Veldhuizen.** 2020. “Bribing the Self.” *Games and Economic Behavior*, 120: 917–946.
- Golman, Russell, and George Loewenstein.** 2018. “Information gaps: A theory of preferences regarding the presence and absence of information.” *Decision*, 5(3).
- Golman, Russell, David Hagmann, and George Loewenstein.** 2017. “Information Avoidance.” *Journal of Economic Literature*, 55(1): 1–40.

- Golman, Russell, George Loewenstein, Andras Molnar, and Silvia Saccardo.** 2020. “The demand for, and avoidance of, information.” *Working Paper*.
- Grossman, Zachary.** 2014. “Strategic ignorance and the robustness of social preferences.” *Management Science*, 60(11): 2659–2665.
- Grossman, Zachary.** 2015. “Self-signaling and social-signaling in giving.” *Journal of Economic Behavior & Organization*, 117(0): 26–39.
- Grossman, Zachary, and Joël J van der Weele.** 2017. “Self-image and willful ignorance in social decisions.” *Journal of the European Economic Association*, 15(1).
- Haisley, Emily C., and Roberto A. Weber.** 2010. “Self-serving interpretations of ambiguity in other-regarding behavior.” *Games and Economic Behavior*, 68: 614–625.
- Handel, Benjamin, and Joshua Schwartzstein.** 2018. “Frictions or Mental Gaps: What’s Behind the Information We (Don’t) Use and When Do We Care?” *Journal of Economic Perspectives*, 32(1): 155–178.
- Hsee, Christopher K.** 1996. “Elastic justification: How unjustifiable factors influence judgments.” *Organizational Behavior and Human Decision Processes*, , (1).
- Huck, Steffen, Nora Szech, and Lukas M Wenner.** 2017. “More effort with less pay: On information avoidance, optimistic beliefs, and performance.” *Working Paper*.
- Jacobsen, Karin J, Kari H Eika, Leif Helland, Jo Thori Lind, and Karine Nyborg.** 2011. “Are nurses more altruistic than real estate brokers?” *Journal of Economic Psychology*, 32(5): 818–831.
- Kajackaite, Agne.** 2015. “If I close my eyes, nobody will get hurt. The effect of ignorance on performance in a real effort experiment.” *Journal of Economic Behavior and Organization*, 116: 518–524.
- Karlsson, Niklas, George Loewenstein, and Duane Seppi.** 2009. “The ostrich effect: Selective attention to information.” *Journal of Risk and Uncertainty*, 38(2).
- Krupka, Erin L., and Roberto A. Weber.** 2013. “Identifying social norms using coordination games: Why does dictator game sharing vary?” *Journal of the European Economic Association*, 11(3): 495–524.
- Larson, Tara, and Monica C. Capra.** 2009. “Exploiting moral wiggle room: Illusory preference for fairness? A comment.” *Judgment and Decision Making*, 4(6): 467–474.

- Lazear, Edward P., Ulrike Malmendier, and Roberto A. Weber.** 2012. "Sorting in experiments with application to social preferences." *American Economic Journal: Applied Economics*, 4(1): 136–163.
- Lin, Stephanie C., Rebecca L. Schaumberg, and Taly Reich.** 2016. "Sidestepping the rock and the hard place: The private avoidance of prosocial requests." *Journal of Experimental Social Psychology*, 35–40.
- Matthey, Astrid, and Tobias Regner.** 2011. "Do I really want to know? A cognitive dissonance-based explanation of other-regarding behavior." *Games*, 2(1): 114–135.
- Mijović-Prelec, Danica, and Drazen Prelec.** 2010. "Self-deception as self-signalling: a model and experimental evidence." *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365(1538): 227–240.
- Mobius, Markus M., Muriel Niederle, Paul Niehaus, and Tanya S. Rosenblat.** 2014. "Managing Self-Confidence: Theory and Experimental Evidence." *Working Paper*.
- Nyborg, Karine.** 2011. "I don't want to hear about it: Rational ignorance among duty-oriented consumers." *Journal of Economic Behavior & Organization*, 79(3): 263–274.
- Oberholzer-Gee, Felix, and Reiner Eichenberger.** 2008. "Fairness in Extended Dictator Game Experiments." *The B.E. Journal of Economic Analysis & Policy*, 8(1).
- Olschewski, Sebastian, Marius Dietsch, Elliot A Ludvig, et al.** 2019. "Anti-social motives explain increased risk aversion for others in decisions from experience." *Judgment and Decision Making*, 14(1): 58–71.
- Oster, Emily, Ira Shoulson, and E Dorsey.** 2013. "Optimal expectations and limited medical testing: evidence from Huntington disease." *American Economic Review*, 103(2): 804–30.
- Pace, Davide D., and Joël J van der Weele.** 2020. "Curbing Carbon: An Experiment on Uncertainty and Information about CO2 emissions." *Working Paper*.
- Rabin, M.** 1995. "Moral preferences, moral constraints, and self-serving biases." *Working paper*.
- Regner, Tobias.** 2018. "Reciprocity under moral wiggle room: Is it a preference or a constraint?" *Experimental Economics*, 1–18.
- Saccardo, Silvia, and Marta Serra-Garcia.** 2020. "Cognitive Flexibility or Moral Commitment? Evidence of Anticipated Belief Distortion." *Working Paper*.

- Schwardmann, Peter, Egon Tripodi, and Joël J. van der Weele.** 2019. “Self-Persuasion: Evidence from Field Experiments at Two International Debating Competitions.” *Working Paper*.
- Serra-Garcia, Marta, and Nora Szech.** 2020. “The (in)elasticity of moral ignorance.” *Working Paper*.
- Sicherman, Nachum, George Loewenstein, Duane J Seppi, and Stephen P Utkus.** 2016. “Financial attention.” *The Review of Financial Studies*, 29(4): 863–897.
- Sims, Christopher A.** 2003. “Implications of rational inattention.” *Journal of Monetary Economics*, 50(3): 665–690.
- Spiekermann, Kai, and Arne Weiss.** 2016. “Objective and subjective compliance: A norm-based explanation of ‘moral wiggle room’.” *Games and Economic Behavior*, 96: 170–183.
- Thornton, Rebecca L.** 2008. “The demand for, and impact of, learning HIV status.” *American Economic Review*, 98(5).
- Thunström, Linda, Jonas Nordström, Jason F Shogren, Mariah Ehmke, and Klaas van’t Veld.** 2016. “Strategic self-ignorance.” *Journal of Risk and Uncertainty*, 52(2).
- Trachtman, Hannah, Andrew Steinkruger, Mackenzie Wood, Adam Wooster, James Andreoni, James J. Murphy, and Justin M. Rao.** 2015. “Fair weather avoidance: unpacking the costs and benefits of “Avoiding the Ask”.” *Journal of the Economic Science Association*, 1–7.
- van der Weele, Joël J.** 2014. “Inconvenient Truths: Determinants of Strategic Ignorance in Moral Dilemmas.” *Working Paper*.
- van der Weele, Joël J., Julija Kulisa, Michael Kosfeld, and Guido Friebel.** 2014. “Resisting Moral Wiggle Room: How Robust Is Reciprocal Behavior?” *American Economic Journal: Microeconomics*, 6(3): 256–264.