

Free to Fail? Paternalistic Preferences in the United States*

Björn Bartling, Alexander W. Cappelen, Henning Hermes,
Marit Skivenes & Bertil Tungodden

June 1, 2023

Abstract

We study paternalistic preferences in two large-scale experiments with participants from the general population in the United States. Spectators decide whether to intervene to prevent a stakeholder, who is mistaken about their choice set, from making a choice that is not aligned with the stakeholders' own preferences. We find causal evidence for the nature of the intervention being of great importance for the spectators' willingness to intervene. Only a minority of the spectators implement a hard intervention that removes the stakeholder's freedom to choose, while a large majority implement a soft intervention that provides information without restricting the choice set. This finding holds regardless of the stakeholder's responsibility for being mistaken about the choice set – whether the source of mistake is internal or external – and in different subgroups of the population. We introduce a theoretical framework with two paternalistic types – libertarian paternalists and welfarists – and show that the two types can account for most of the spectator behavior. We estimate that about half of the spectators are welfarists and that about a third are libertarian paternalists. Our results shed light on attitudes toward paternalistic policies and the broad support for soft interventions.

JEL classifications: C91, C93, D69, D91

Keywords: paternalism, libertarian paternalism, welfarism, freedom to choose

*We thank Sandro Ambuehl, Ceren Ay, Jill D. Berrick, Stefano DellaVigna, Martin Dufwenberg, Mathias Ekström, Dirk Engelmann, Christine Exley, Eleonora Freddi, Uri Gneezy, Hege S. Helland, Holger Herz, Ariel Kalil, Judd Kessler, Botond Köszegi, Dorothea Kübler, Mathea Loen, Audun Løvlie, George Loewenstein, Matt Lowe, Muriel Niederle, Nick Netzer, Christopher Roth, Barbara Ruiken, Erik Ø. Sørensen, Krishna Srinivasan, Matthias Sutter, Georg Weizsäcker, Ludger Woessmann and numerous seminar and conference participants for very helpful comments. Camilla Allocchio provided excellent research support. We gratefully acknowledge funding from the Norwegian Research Council through its Centers of Excellence Scheme project 262675 (FAIR) and Research Grants 262636 and 325134, ERC Consolidator Grant 724460 and ERC Advanced Grant 788433. This study was preregistered in the AEA RCT Registry (AEARCTR-0004630) and administered by FAIR-The Choice Lab.

1 Introduction

People sometimes make choices that are detrimental to their welfare. This creates opportunities for paternalistic interventions (Camerer, Issacharoff, Loewenstein, O’Donoghue, and Rabin, 2003; Thaler and Sunstein, 2003). The extent to which such opportunities should be used is a key issue in the relationship between the state and its citizens (Dworkin, 1972; Arneson, 1980; Le Grand and New, 2015). The role of paternalistic interventions is also at the heart of many interpersonal relationships, such as the relationship between parents and their children, experts and laypeople, employers and employees, and donors and recipients (Jacobsson, Johannesson, and Borgquist, 2007; Doepke and Zilibotti, 2017; Gangadharan, Grossman, Jones, and Leister, 2018; Kassirer, Levine, and Gaertig, 2020; Kiessling, Chowdhury, Schildberg-Hörisch, and Sutter, 2021; Beshears, Choi, Laibson, Madrian, and Skimmyhorn, 2022).

Opportunities for paternalistic interventions raise two fundamental normative questions. First, is it acceptable to restrict an individual’s freedom to choose to promote their welfare? Second, who should judge whether an intervention improves the welfare of an individual? These questions have shaped an extensive normative literature, both in economics and in the social sciences more broadly, on how to justify paternalistic interventions (Nussbaum, 2001; Sen, 2004; Kaplow and Shavell, 2006; Hausman and McPherson, 2009; Le Grand and New, 2015; Thaler and Sunstein, 2021).

Libertarian Paternalism, which has become an influential position in recent years, can be characterized in terms of these two questions (Thaler and Sunstein, 2021). It only justifies interventions that (i) do not restrict people’s freedom to choose and (ii) promote their welfare, as judged by themselves. Informed by insights from the behavioral sciences, a large and growing literature on libertarian paternalism has advocated interventions that manipulate the choice architecture without reducing the choice set of individuals, such as default options or information provision. These interventions aim to nudge people to make choices that are aligned with their own preferences (Thaler and Sunstein, 2021). The idea of libertarian paternalism has received considerable attention by policy makers and business leaders, as evidenced by the many behavioral insights teams established by governments and corporations across the world (OECD, 2017; DellaVigna and Linos, 2022; List, Rodemeier, Roy, and Sun, 2023).

The focus on the freedom to choose in libertarian paternalism contrasts the other main position in the normative literature, the classical *Welfarism* approach, which assesses policies only in terms of how they affect people’s welfare (Sen, 2004). The welfarism approach finds interventions that restrict people’s choice set acceptable as long as the interventions promote their welfare. The most prominent view of welfarism in economics overlaps with libertarian paternalism in considering an individual’s own preferences to be the appropriate basis for welfare evaluations (Kaplow and Shavell, 2000, 2001). However, the welfarism approach can also be combined with other conceptions

of welfare that do not rely on satisfying an individual's own preferences (Nussbaum, 2001; Sen, 2004; Hausman and McPherson, 2009; Le Grand and New, 2015).

To shed light on people's paternalistic preferences and the extent to which they are consistent with libertarian paternalism or welfarism, we implement a novel experimental design in two large-scale experiments with more than 14,000 participants from the general population in the United States in the role of a spectator. In both experiments, the spectators make consequential decisions for another individual, the stakeholder. The first experiment, Study 1, explores whether a concern for the stakeholder's freedom to choose affects the spectator's willingness to intervene by manipulating the nature of the intervention. The second experiment, Study 2, investigates both the spectators' intervention decisions and their welfare evaluations, which allows us to provide estimates of the prevalence of libertarian paternalists and welfarists in the general population.

In Study 1, each spectator is matched with a stakeholder who will receive a monetary bonus. There are two bonus options in the stakeholder's choice set, a safe option and a risky option. Absent an intervention, the stakeholder will make a choice between the two options in a non-transparent choice environment. The spectator is informed that the non-transparent choice environment leads the stakeholder to be mistaken about the odds of the risky option and that, as a consequence, the stakeholder prefers the risky option to the safe option. The spectator is also informed that the matched stakeholder would prefer the safe option over the risky option if they were not mistaken about the odds of the risky option. In one set of treatments, the spectator is given the opportunity to implement a *hard intervention*, which removes the stakeholder's freedom to choose and gives them the safe option. In another set of treatments, the spectator is given the opportunity to implement a *soft intervention*, which does not restrict the stakeholder's choice set but informs the stakeholder about the correct odds of the risky option.

Study 1 also includes a second treatment dimension where we vary the reason why the stakeholder is mistaken about the choice set. A growing literature on social preferences documents that the willingness to redistribute depends on the extent to which individuals are seen as responsible for their situation (Konow, 2000; Fong, 2001; Cappelen, Drange Hole, Sørensen, and Tungodden, 2007; Alesina, Stantcheva, and Teso, 2018; Almås, Cappelen, and Tungodden, 2020). Likewise, the willingness to make a paternalistic intervention may depend on whether individuals are seen as responsible for making choices that are not aligned with their own preferences. We therefore study the role of the source of the stakeholder's mistake for the spectator's willingness to intervene. In one set of treatments, the spectator is informed that the stakeholder has made an incorrect calculation, which we refer to as a situation with *internal responsibility*. In another set of treatments, the spectator is informed that the stakeholder received incorrect information, which we refer to as a situation with *external responsibility*.

Study 1 provides three main findings. First, we document that, in line with libertarian paternalism, the nature of a paternalistic intervention is a major causal determinant of

the spectators' willingness to intervene. While a large majority of about 85 percent of the spectators are willing to implement the soft intervention that does not restrict the stakeholder's choice set, only about a third of the spectators are willing to implement the hard intervention that removes the stakeholder's freedom to choose. Second, we find that the source of the stakeholder's mistake is of minor importance for the spectators' willingness to intervene. Third, the heterogeneity analysis shows that the estimated treatment effects are robust across different subgroups of the general population in the United States.

Taken together, Study 1 shows that, both for the internal and the external source of mistake, only a minority of the spectators are willing to implement the hard intervention, while a large majority are willing to implement the soft intervention. To guide the interpretation of this empirical pattern, we develop a theoretical framework in which a spectator's willingness to intervene is determined by two key features of paternalistic preferences: whether the spectator cares about the stakeholder's freedom to choose and how the spectator conceptualizes the stakeholder's welfare. We impose minimal assumptions on the spectator's behavior capturing that the spectators do not make dominated choices. The assumptions imply that the share of spectators whose welfare evaluations are aligned with the stakeholder's preferences should be (i) weakly larger than the share of spectators implementing the hard intervention and (ii) weakly smaller than the share of spectators implementing the soft intervention. Further, we formalize two paternalistic types, libertarian paternalists and welfarists, and show how the prevalence of these paternalistic types relates to key parameters in the theoretical framework.

The empirical pattern observed in Study 1 is consistent with the theoretical framework: the share of spectators implementing the hard intervention is smaller than the share of spectators implementing the soft intervention. However, Study 1 does not provide us with a measure of the share of spectators whose welfare evaluations are aligned with the stakeholder's preferences and thus does not allow us to fully test the two implications of the framework. Consequently, we conduct Study 2 that does not only manipulate the nature of the intervention but also elicits whether the spectator's welfare evaluation is aligned with the stakeholder's preferences. Specifically, Study 2 introduces a new treatment, in which the spectator decides whether to allocate the safe option or the risky option to the stakeholder, without having the opportunity to give the stakeholder the freedom to choose. Study 2 thus provides a stricter test of the theoretical framework and enables us to estimate the share of libertarian paternalists and welfarists based on both key features of paternalistic preferences.

Study 2 replicates the empirical pattern established in Study 1. We find again that a large majority of more than 80 percent of the spectators are willing to implement the soft intervention, while only about a third of the spectators are willing to implement the hard intervention. Hence, we establish in two independent large-scale samples of the general population in the United States that the nature of an intervention is of great

importance for the willingness to intervene. Further, in the new treatment, we find that a large majority of the spectators, about 70 percent, allocate the safe option to the stakeholder, aligned with the stakeholder's preferences. Taken together, the findings are in line with the theoretical framework: the share of spectators whose welfare evaluations are aligned with the stakeholder's preferences is (i) larger than the share of spectators implementing the hard intervention and (ii) smaller than the share of spectators implementing the soft intervention. Hence, Study 2 provides evidence that the willingness to intervene is determined by how an intervention affects the stakeholder's welfare and the stakeholder's freedom to choose. Finally, we estimate the prevalence of the two paternalistic types. We find that the majority of the spectators are welfarists, 52.9 percent, but also that a significant share of the spectators are libertarian paternalists, 34.4 percent. In terms of the two views on welfarism, we find that 35.4 percent of the spectators are welfarists whose welfare evaluations are aligned with the stakeholder's preferences and 17.5 percent are welfarists whose welfare evaluations are not aligned with the stakeholder's preferences. Overall, our findings show that the two paternalistic types can rationalize the behavior of most of the spectators in both Study 1 and Study 2.

Our results shed light on why soft interventions have gained strong support in recent years (Thaler and Sunstein, 2021). In line with the attention given to soft interventions in the policy debate (OECD, 2017), we find that the vast majority of the general population in the United States is willing to implement a soft intervention that promotes the welfare of a stakeholder, as judged by themselves. However, this should not be interpreted as evidence of most Americans being libertarian paternalists. In fact, in our studies, a large part of the support for the soft intervention comes from welfarists who are willing to implement both the soft intervention and the hard intervention. Hence, the popularity of soft interventions can be explained by these interventions attracting support both from welfarists who respect the preferences of the stakeholder and from libertarian paternalists. In the same way, resistance to hard interventions that restrict people's choice set may not only be driven by libertarian paternalists who respect people's freedom to choose, but may, as in our study, also reflect that a significant share of people are welfarists who believe that the hard intervention does not promote the welfare of the stakeholder. An interesting implication of the estimated prevalence of the different paternalistic types is that the libertarian paternalists are part of the majority coalition on the acceptability of the soft intervention and also part of the majority coalition on the non-acceptability of the hard intervention. As a result, even though we estimate libertarian paternalists to comprise only about a third of the population in the United States, they may trigger both political support for implementing soft interventions and political resistance against hard interventions.

The paper contributes to the growing literature on paternalism by being the first experimental study on the role of the nature of a paternalistic intervention for people's willingness to intervene. Several survey-based studies have shown that a majority

of the population in various countries approve of a broad range of soft interventions (Diepeveen, Ling, Suhrcke, Roland, and Marteau, 2013; Reisch and Sunstein, 2016; Evers, Marchiori, Junghans, Cremers, and De Ridder, 2018; Sunstein, Reisch, and Rauber, 2018), even though there is some resistance to soft interventions that are considered to be manipulative (Jung and Mellers, 2016; Tannenbaum, Fox, and Rogers, 2017; Arad and Rubinstein, 2018). Further, in a recent laboratory experiment with students, Ambuehl, Bernheim, and Ockenfels (2021) show that spectators project their own time preferences onto the stakeholders and are frequently willing to restrict the stakeholders' choice sets by removing impatient choice options. We extend this literature by showing in large-scale experiments with general population samples in the United States that the majority of people are willing to implement a soft intervention but not willing to implement a hard intervention. A novel feature of our setting is that the spectator is informed about the stakeholder's preferences, which allows us to observe the spectator's intervention decision in isolation of uncertainty about the stakeholder's preferences. Further, we provide evidence showing that the lower willingness to implement the hard intervention than the soft intervention is driven by a significant share of the population being libertarian paternalists who respect people's freedom to choose. Finally, we contribute to the literature by formalizing the two most prominent paternalistic types in the normative literature, libertarian paternalists and welfarists, and by providing estimates of the prevalence of these paternalistic types in a general population sample.

More broadly, the paper contributes to the literature examining heterogeneity in people's social preferences (Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000; Konow, 2000; Andreoni and Miller, 2002; Charness and Rabin, 2002; Cappelen et al., 2007; Bellemare, Kröger, and van Soest, 2008; Cappelen, Konow, Sørensen, and Tungodden, 2013; Durante, Putterman, and van der Weele, 2014; Almås et al., 2020), by providing evidence of significant heterogeneity in people's paternalistic preferences. Our paper further contributes to the social preference literature by showing that the willingness to intervene does not depend on whether individuals are seen as responsible for making choices that are not aligned with their own preferences. Finally, the paper contributes to the literature on the intrinsic value of decision rights, power, and self-determination (Deci and Ryan, 1985; Fehr, Herz, and Wilkening, 2013; Bartling, Fehr, and Herz, 2014; Owens, Grossman, and Fackler, 2014; Sloof and von Siemens, 2017; Pikulina and Tergiman, 2020), by providing evidence that people do not only value their own autonomy but also respect other people's freedom to choose.

The paper is organized as follows. Section 2 presents Study 1. Section 3 presents the theoretical framework. Section 4 presents Study 2. Section 5 concludes. In the Appendix, we provide supplementary analysis, including adjustments for multiple hypothesis testing (Section A), the experimental procedures and instructions (Section B), and details about the pre-analysis plans (Section C).

2 Study 1

In this section, we present the experimental design, sample, empirical strategy, and main results of Study 1.

2.1 Experimental Design

The experiment has two types of participants: spectators and stakeholders. The spectators make intervention decisions that are consequential for the stakeholders. Our interest is in the spectators' intervention decisions, and the sole function of the stakeholders is to render the spectators' decisions consequential.

We first present the context of the intervention decision, before we introduce the treatment manipulations.

Context. Each spectator is matched with a stakeholder who will receive a bonus payment. There are two bonus options, a safe payment of USD 4 and a risky payment of USD 10 or USD 0 with equal probability. Absent an intervention, the stakeholder will make a choice between the two bonus options in a non-transparent choice environment. The non-transparent choice environment leads the stakeholder to be mistaken about the choice set. Specifically, the stakeholder is mistaken about the odds of the risky option and, as a consequence, the stakeholder prefers the risky option to the safe option. However, the stakeholder would prefer the safe option to the risky option if they were not mistaken about the odds of the risky option.

More formally, let s denote the option with the safe payment of USD 4 and let r denote the option with the risky payment of USD 10 or USD 0 with equal probability. In the non-transparent choice environment, the stakeholder mistakenly believes that the choice is not between s and r , but between s and a different risky option, which we refer to as $\tilde{r} \neq r$. The stakeholder's preference ranking is given by $\tilde{r} \succ s \succ r$.

The spectator is fully informed about the preference ranking of the stakeholder, and the experimental design builds on the assumption that the spectator believes that the stakeholder chooses according to their preferences. A preference-maximizing stakeholder will choose the risky option in the non-transparent choice environment because, in this case, the stakeholder mistakenly believes that the choice is between s and \tilde{r} , and it holds that $\tilde{r} \succ s$. Consequently, absent an intervention, the stakeholder ends up with the non-preferred option r . The spectator is given the opportunity to intervene to ensure that the stakeholder ends up with their preferred option s .

This context captures the key characteristics of situations that create opportunities for paternalistic interventions: a stakeholder is about to make a choice that is not aligned with their preferences, and a spectator can intervene to ensure that they receive their preferred option.

Nature of Intervention. To study the role of the nature of the intervention for the spectators' intervention decision, the spectators are randomly assigned either to treatments where they can implement a *hard* intervention or to treatments where they can implement a *soft* intervention. The hard intervention removes the stakeholder's freedom to choose. If the spectator implements the hard intervention, the stakeholder cannot make a choice but is allocated the safe option s . The soft intervention, in contrast, does not remove the stakeholder's freedom to choose. If the spectator implements the soft intervention, the stakeholder will make a choice in the transparent choice environment. The stakeholder will then know that the choice is between s and r . Taken together:

- The outcome of intervening, hard or soft, is that the stakeholder ends up with their preferred safe option s .
- The outcome of not intervening, hard or soft, is that the stakeholder ends up with their non-preferred risky option r .

It follows that if the spectators' willingness to intervene only depends on the outcome of an intervention, the share of spectators that implement the soft intervention would be equal to the share of spectators that implement the hard intervention. Hence, the experimental design allows us to identify whether the nature of the intervention matters for the spectators' willingness to intervene, and whether spectators prefer a hard intervention or a soft intervention.

Source of Mistake. To study the role of the source of mistake, the spectators are randomly assigned either to treatments in which the source of mistake is *internal* or to treatments in which the source of mistake is *external*. In treatments with the internal source of mistake, the spectator is informed that they are matched to a stakeholder who had to calculate the odds of the risky option and made a mistake in the calculations. In treatments with the external source of mistake, the spectator is informed that they are matched to a stakeholder who was unlucky and received incorrect information about the odds of the risky option. The source of mistake does not affect the choices of a preference maximizing stakeholder and, consequently, it does not affect the outcomes of intervening or not intervening.

It follows that if the spectators' willingness to intervene does not depend on the source of mistake, the share of spectators that intervene when the source of mistake is internal would be equal to the share of spectators that intervene when the source of mistake is external. Hence, the experimental design allows us to identify whether the source of mistake matters for the spectators' willingness to intervene, and whether spectators are more willing to intervene when the source of mistake is internal or external.

Treatment Design. We implemented a full factorial 2×2 between-subjects design. We refer to the four treatments as *Hard* \times *Internal*, *Hard* \times *External*, *Soft* \times *Internal*, and

Soft × *External*. Figure 1 summarizes the sequence of the main events in the experiment and shows at which stage the treatment manipulations come into play.

[Figure 1 about here]

2.2 Participants

In this section, we present the sample of spectators and explain how we recruited and matched stakeholders to spectators.

Spectators. The spectators were recruited from the general population in the United States through a professional data service company (Dynata). We sampled a total of 8,004 spectators in August 2019, based on quotas for gender, age, education, income, and region, to match a representative sample of the general population in the United States (aged 18 or older). The spectators had to pass an attention filter before being randomized with equal probability to one of the four treatments. Each spectator made a single intervention decision. The spectators were informed that one out of five spectator decisions would be randomly selected and implemented.

We elicited the spectators’ demographic background characteristics, including gender, age, education, and income. Further, since the intervention decision is made in a context where the stakeholder is about to make a choice in the domain of risk, we also measured the spectators’ own willingness to take risks by eliciting their self-assessment on an 11-point scale ranging from “Completely unwilling to take risks” (0) to “Very willing to take risks” (10). At the end of the experiment, spectators could self-identify as “Republican,” “Democrat,” or “Independent/Third Party.” Finally, we elicited the extent to which the spectators agree with the following statements: “People sometimes make choices that harm their own well-being” and “The government can sometimes improve its citizens’ well-being by restricting their freedom of choice.” The spectators could answer on a seven-point scale ranging from “fully disagree” (1) to “fully agree” (7).

Table 1 reports the main demographic characteristics of the sample in Study 1 and compares it to the population in the United States. From the leftmost column, we observe that the sample is gender-balanced and exhibits significant heterogeneity in age, education, and income. The median age in our sample is 46 years. In terms of education, 31% of the sample do not have any college education, while 15% have a Master’s degree or an even higher educational attainment. The median educational category is “some college.” About one quarter of the sample has a yearly income of less than USD 30,000, while slightly more than 10% have an income exceeding USD 150,000. The median income category is USD 30,000 to 60,000. Regarding political affiliation, 29% self-identify as Republicans, 33% as Democrats, and 28% as Independents/Third

Party, while 10% did not report a political affiliation. We also observe substantial heterogeneity in the spectators’ risk preferences: 23% of the spectators indicate that their willingness to take risks is below the midpoint of 5 on the 11-point scale from 0 to 10, and 64% indicate that their willingness to take risks is above the midpoint; the median response is 6. Comparison of the sample with the population averages in the United States (rightmost column) reveals that the sample closely mirrors the population statistics with respect to gender and age, but contains a slightly lower share with low education and with a household income of at least USD 150,000. Table A1 in the Online Appendix shows that the sample is balanced across treatments.

[Table 1 about here]

Stakeholders. We recruited the stakeholders through an online labor market platform (Amazon Mechanical Turk). The stakeholders could receive a bonus payment. We elicited the stakeholders’ preferences over the safe and the risky bonus option in both the transparent choice environment and in the non-transparent choice environment. Only stakeholders who prefer the safe option in the transparent choice environment and the risky option in the non-transparent choice environment were matched to a spectator.¹

2.3 Empirical Strategy

To examine how the nature of the intervention and the source of mistake causally affect the spectators’ willingness to intervene, we use the following empirical specification:

$$I_i = \beta_0 + \beta_1 S_i + \beta_2 E_i + \beta_3 S_i E_i + \gamma X_i + \varepsilon_i \quad (1)$$

The dependent variable I_i is an indicator for whether spectator i intervenes. Treatment *Hard*×*Internal* is the omitted category. S_i is an indicator for spectator i being in a treatment with a soft intervention, E_i is an indicator for spectator i being in a treatment with the external source of mistake, $S_i E_i$ is the interaction between S_i and E_i , X_i is a vector of background characteristics, and ε_i is an idiosyncratic error term. X_i includes political orientation, willingness to take risks, education, income, age, and gender. In the analysis, the background characteristics are defined by the following indicator variables: Republican indicates whether a spectator identifies as Republican or non-Republican. High Risk Taking, High Education, High Income, and High Age indicate whether a spectator is above or below the median of the respective characteristic in

¹See Section B.2 of the Appendix for further details about the preference elicitation and the matching protocol.

the sample. Female indicates whether a spectator is female or male. We estimate the models with and without the vector of background characteristics.

The coefficient β_1 provides an estimate of the causal effect of the nature of the intervention on the spectators' willingness to intervene. The coefficient β_2 provides an estimate of the causal effect of the source of mistake on the spectators' willingness to intervene. The coefficient β_3 provides an estimate of the interaction effect between the nature of the intervention and the source of mistake on the spectators' willingness to intervene. We also estimate the causal effect of the nature of the intervention when pooling the treatments with the hard intervention (*Hard*×*Internal* and *Hard*×*External*) and the treatments with the soft intervention (*Soft*×*Internal* and *Soft*×*External*), and we estimate the causal effect of the source of mistake when pooling the treatments with the internal source of mistake (*Hard*×*Internal* and *Soft*×*Internal*) and the treatments with the external source of mistake (*Hard*×*External* and *Soft*×*External*).

In the heterogeneity analysis, we study the average causal effect of the nature of the intervention and the source of mistake in different subgroups when pooling the respective treatments. In this analysis, we use the following specification:

$$I_i = \beta_0 + \beta_1 T_i + \beta_2 x_i + \beta_3 T_i x_i + \varepsilon_i \quad (2)$$

where x_i indicates a single background characteristic of spectator i , and $T_i x_i$ is the interaction between x_i and the treatment indicator $T_i = S_i, E_i$. Equation (2) is estimated separately for x_i indicating political orientation, willingness to take risks, education, income, gender, and age (single interaction model). In the analysis, we also estimate a model that jointly includes x_i and the interaction term $T_i x_i$ for each background characteristic (joint interaction model).

2.4 Results

We start by providing an overview of the spectators' intervention decisions. Figure 2 shows the share of spectators that intervene by treatment. The left panel shows the share of spectators that intervene in treatments *Hard*×*Internal* and *Soft*×*Internal*. The right panel shows the share of spectators that intervene in treatments *Hard*×*External* and *Soft*×*External*.

[Figure 2 about here]

We observe that only about a third of the spectators implement the hard intervention, both when the source of mistake is internal (33.5 percent) and when the source of mistake is external (30.0 percent). Hence, the large majority of spectators decide not to

restrict the stakeholder’s freedom to choose, even though the hard intervention would ensure that the stakeholder is allocated their preferred safe option. In contrast, a large majority of the spectators implement the soft intervention (internal: 85.9 percent, external: 87.5 percent), which preserves the stakeholder’s freedom to choose and ensures that the stakeholder can make a choice in the transparent choice environment.

[Table 2 about here]

Table 2 reports the regression analysis. In column (1), pooling the treatments with the hard intervention and the treatments with the soft intervention, we estimate the average causal effect of the nature of the intervention on the spectators’ willingness to intervene: the share of spectators that implement the soft intervention is 55.0 percentage points higher than the share of spectators that implement the hard intervention ($p < 0.01$). Column (2) shows that the estimated causal effect is virtually unaffected by the inclusion of the spectators’ background characteristics. We further note that the estimated coefficients for the background characteristics are small and in most cases not significant.

Columns (3) and (4) estimate the model with an interaction variable between the nature of the intervention and the source of mistake. We find only a small interaction effect: the estimated difference between the share of spectators that implement the hard intervention and the share of spectators that implement the soft intervention is 5.1 percentage points larger when the source of mistake is external rather than internal ($p < 0.01$). Consequently, the treatment effect of manipulating the nature of the intervention is large both when the source of mistake is internal (52.4 percent, $p < 0.01$) and when it is external (57.5 percent, $p < 0.01$).

We summarize the analysis of how the nature of the intervention affects the spectators’ willingness to intervene as follows:

Result 1: *The nature of the intervention has a substantial causal effect on the spectators’ willingness to intervene, both when the source of the stakeholder’s mistake is internal and when it is external.*

We now turn to an analysis of the causal effect of the source of mistake on the spectators’ willingness to intervene. Columns (3) and (4) show that the estimated share of spectators that implement the hard intervention is 3.5 percentage points lower when the source of mistake is external rather than internal ($p < 0.05$), while the estimated share of spectators that implement the soft intervention is 1.6 percentage points higher, but the difference is not statistically significant ($p = 0.15$ and $p = 0.13$, respectively). In columns (5) and (6), pooling the treatments with the internal source of mistake and the treatments with the external source of mistake, we find that there is no significant

average effect of the source of mistake on the spectators' willingness to intervene. Taken together, we conclude:

Result 2: *The source of the stakeholder's mistake does not have a substantial causal effect on the spectators' willingness to intervene, irrespective of the nature of the intervention.*

The fact that we have a large-scale general population sample allows us to study whether different subgroups of the population differ in their intervention decisions. In Figure 3, we report the heterogeneity analysis for subgroups defined by political orientation, willingness to take risks, education, income, age, and gender. The left panel shows the estimated interaction effect between an indicator for the spectator being in one of the treatments with the soft intervention and an indicator for the respective background characteristic. The right panel shows the estimated interaction effect between an indicator for the spectator being in one of the treatments with external source of mistake and an indicator for the respective background characteristic.

[Figure 3 about here]

In the left panel of Figure 3, we observe that the estimated average causal effect of the nature of the intervention on the spectators' willingness to intervene, both for the single and the joint interaction model, is not significantly different across subgroups, with the exception of High Education and High Income. The difference between above-median and below-median educated spectators is significant in both the single interaction model and the joint interaction model and robust to multiple hypothesis testing (see Tables A4 and A7 in the Appendix). In the single interaction model, the estimated average causal effect is 6.0 percentage points larger for above-median educated spectators than for below-median educated spectators ($p < 0.01$). Above-median educated spectators are slightly less likely to implement the hard intervention (30.4 vs. 32.9 percent) and slightly more likely to implement the soft intervention (88.5 vs. 85.0 percent) than below-median educated spectators. The difference between spectators with above-median and below-median income is significant only in the single interaction model but not in the joint interaction model, and it is not robust to multiple hypothesis testing. The right panel shows that the estimated average causal effect of the source of mistake on the spectators' willingness to intervene is not different across subgroups (see Tables A5 and A8 in the Appendix).

There is no significant interaction effect with respect to political orientation, both in terms of the nature of the intervention and the source of mistake. Further, taking into account that we do not find a significant level effect of being Republican (see Table 2),

it follows that Republicans and non-Republicans make very similar intervention decisions in the experiment.² This finding may suggest that political disagreements about paternalistic policies are more related to disagreements about the consequences of paternalistic interventions, which are controlled for in our experiment, than to fundamental differences in paternalistic preferences.

The heterogeneity analysis also shows that the estimated average causal effect of the nature of the intervention is large and highly significant in all subgroups ($p < 0.01$ in all tests, see Table A4). In contrast, the estimated average causal effect of the source of mistake is small and not significant in all subgroups ($p > 0.14$ in all tests, see Table A5).

We sum up the heterogeneity analysis in the following result:

Result 3: *There are only small differences in intervention decisions across subgroups. In all subgroups, the nature of the intervention has a substantial average causal effect on the spectators' willingness to intervene, while the source of the stakeholder's mistake does not.*

3 Theoretical Framework

Let θ^H denote the share of spectators that implement the hard intervention and θ^S the share of spectators that implement the soft intervention. Study 1 establishes, both for the internal and the external source of mistake, the following empirical pattern:

$$0 < \hat{\theta}^H < \hat{\theta}^S < 1, \quad (3)$$

where $\hat{\theta}^H$ and $\hat{\theta}^S$ denote the estimated shares. This raises three questions about the spectators' intervention decisions:

- Why do some spectators implement the hard intervention?
- Why do more spectators implement the soft intervention than the hard intervention?
- Why do some spectators not implement the soft intervention?

We introduce a theoretical framework to shed light on these questions.

²Table A6 shows that this finding is robust to focusing on the sub-sample of Republicans and Democrats and to different model specifications.

3.1 Spectators' Preferences

We assume that a spectator's preferences are defined over the stakeholder's welfare and the stakeholder's freedom to choose.

Let $W(b)$ denote a spectator's evaluation of the stakeholder's welfare, which is determined by the bonus option, $b \in \{s, r\}$, that the stakeholder ends up with. We assume that either $W(s) > W(r)$ or $W(s) < W(r)$. Let $U(b)$ represent the preference ranking of the stakeholder, with $U(s) > U(r)$ for all stakeholders in the experiment. A spectator's welfare evaluation is aligned with the stakeholder's preferences if and only if $W(b) = U(b)$, which in the experiment would entail that $W(s) > W(r)$. Let θ^A denote the share of spectators whose welfare evaluations are aligned with the stakeholder's preferences.

Let $F(c)$ denote a spectator's evaluation of the stakeholder's freedom to choose, which is determined by the stakeholder's choice environment, $c \in \{c^{+t}, c^{+nt}, c^{-}\}$, where c^{+t} denotes a transparent choice environment, c^{+nt} a non-transparent choice environment, and c^{-} an environment in which the stakeholder has no choice. We assume that $F(c^{+t}) \geq F(c^{+nt}) > F(c^{-})$.

3.2 Spectators' Intervention Decisions

A spectator can choose to intervene, i , or not to intervene, ni . In the experiment, $b(i) = s$ and $b(ni) = r$, both in treatment *Hard* and in treatment *Soft*. Hence, the welfare consequences of intervening and of not intervening are the same in the two treatments. Further, $c = c^{-}$ if a spectator intervenes in treatment *Hard* and $c = c^{+t}$ if a spectator intervenes in treatment *Soft*. In both treatments, $c(ni) = c^{+nt}$. Hence, the stakeholder's freedom to choose is strictly reduced by intervening in treatment *Hard* and is weakly increased by intervening in treatment *Soft*.

We make the following two minimal assumptions about a spectator's intervention decisions:

A1. A spectator intervenes if $W(b(i)) > W(b(ni))$ and $F(c(i)) \geq F(c(ni))$.

A2. A spectator does not intervene if $W(b(i)) < W(b(ni))$ and $F(c(i)) \leq F(c(ni))$.

The two assumptions imply that a spectator does not make dominated intervention decisions: a spectator who considers that an intervention strictly increases the stakeholder's welfare and at least weakly increases the stakeholder's freedom to choose will intervene, and a spectator who considers that an intervention strictly decreases the stakeholder's welfare and at least weakly decreases the stakeholder's freedom to choose will not intervene.

We can now make the following observation:

Observation 1: Assumptions A1 and A2 imply that $\theta^H \leq \theta^A \leq \theta^S$.

Proof. (i) Consider a spectator who implements the hard intervention. By A2, $W(s) > W(r)$. It follows that $\theta^H \leq \theta^A$. (ii) Consider a spectator who does not implement the soft intervention. By A1, $W(s) < W(r)$. It follows that $\theta^A \leq \theta^S$. \square

The empirical pattern observed in Study 1, $0 < \hat{\theta}^H < \hat{\theta}^S < 1$, is consistent with Observation 1. However, Study 1 does not provide us with a measure of θ^A , the share of spectators whose welfare evaluations are aligned with the stakeholder's preferences. To provide a stricter test of the theoretical framework, we implement Study 2, which comprises a treatment that provides us with an estimate of θ^A , along with estimates of θ^H and θ^S . Study 2 also allows us to study the prevalence of the main paternalistic types in the normative literature, which we now turn to.

3.3 Paternalistic Types

Within the theoretical framework, libertarian paternalism and welfarism can be formalized as follows:

Libertarian Paternalist. A libertarian paternalist intervenes if and only if $F(c(i)) \geq F(c(ni))$ and $W(b(i)) > W(b(ni))$, with $W(b) = U(b)$.

Welfarist. A welfarist intervenes if and only if $W(b(i)) > W(b(ni))$.

The two paternalistic types satisfy assumptions A1 and A2.

A libertarian paternalist intervenes if and only if the intervention preserves the stakeholder's freedom to choose and strictly increases the stakeholder's welfare, with the welfare evaluation being aligned with the stakeholder's preferences. Hence, a libertarian paternalist does not implement any intervention that reduces the stakeholder's freedom to choose, irrespective of how the intervention affects the stakeholder's welfare. A welfarist intervenes if and only if the intervention strictly increases the stakeholder's welfare, irrespective of how the intervention affects the stakeholder's freedom to choose. The welfare evaluation of a welfarist may or may not be aligned with the stakeholder's preferences.

It follows that:

- Welfarists whose welfare evaluations are aligned with the stakeholder's preferences implement both the hard and the soft intervention.
- Libertarian paternalists implement the soft but not the hard intervention.

- Welfarists whose welfare evaluations are not aligned with the stakeholder's preferences neither implement the hard nor the soft intervention.

We now consider how the theoretical framework can be used to study the prevalence of the two paternalistic types.

3.4 Prevalence of Paternalistic Types

Let σ^{LP} denote the share of spectators that are libertarian paternalists, σ^{W_a} the share of spectators that are welfarists whose welfare evaluations are aligned with the stakeholder's preferences, and $\sigma^{W_{na}}$ the share of spectators that are welfarists whose welfare evaluations are not aligned with the stakeholder's preferences.

We establish the following observation:

Observation 2: Assumptions A1 and A2 imply that (i) the share of welfarists whose welfare evaluations are aligned with the stakeholder's preferences is given by $\sigma^{W_a} = \theta^H$, (ii) the share of welfarists whose welfare evaluations are not aligned with the stakeholder's preferences is given by $\sigma^{W_{na}} = 1 - \theta^S$, (iii) the share of libertarian paternalists is given by $\sigma^{LP} = \theta^A - \theta^H$, and (iv) $\sigma^{W_a} + \sigma^{W_{na}} + \sigma^{LP} \leq 1$.

Proof. (i) Consider a spectator who implements the hard intervention. By A2, $W(s) > W(r)$. By A1, the spectator implements the soft intervention. Hence, the spectator is a welfarist whose welfare evaluation is aligned with the stakeholder's preferences. It follows that $\sigma^{W_a} = \theta^H$. (ii) Consider a spectator who does not implement the soft intervention. By A1, $W(r) > W(s)$. By A2, the spectator does not implement the hard intervention. Hence, the spectator is a welfarist whose welfare evaluation is not aligned with the stakeholder's preferences. It follows that $\sigma^{W_{na}} = 1 - \theta^S$. (iii) Consider a spectator for whom it holds that $W(s) > W(r)$. By A1, the spectator implements the soft intervention. Hence, the spectator is either a welfarist (if the spectator implements the hard intervention) or a libertarian paternalist (if the spectator does not implement the hard intervention). It follows that $\sigma^{W_a} + \sigma^{LP} = \theta^A$. Taking into account (i), it follows that $\sigma^{LP} = \theta^A - \theta^H$. (iv) It follows from (i), (ii), and (iii) that $\sigma^{W_a} + \sigma^{W_{na}} + \sigma^{LP} = \theta^H + (1 - \theta^S) + (\theta^A - \theta^H) = 1 - \theta^S + \theta^A$. By A1, $\theta^A \leq \theta^S$. It follows that $\sigma^{W_a} + \sigma^{W_{na}} + \sigma^{LP} \leq 1$. \square

It follows from Observation 2 that if *all* spectators are either welfarists or libertarian paternalists, then $\theta^A = \theta^S$. However, A1 and A2 allow for behavior that cannot be explained by either welfarists or libertarian paternalists: spectators who implement the soft intervention, even though their welfare evaluations are not aligned with the stakeholder's preferences. These spectators must (i) evaluate the stakeholder's freedom to choose to be strictly greater in a transparent choice environment than in a

non-transparent choice environment (otherwise, they would violate A2 by implementing the soft intervention) and (ii) consider the increase in the stakeholder’s freedom to choose from implementing the soft intervention to outweigh what they evaluate to be a loss in the stakeholder’s welfare. Taken together, these spectators cannot be welfarists because they care about the stakeholder’s freedom to choose and they cannot be libertarian paternalists because their welfare evaluations are not aligned with the stakeholder’s preferences. The share of such spectators is given by $\theta^S - \theta^A$. In principle, all spectators could agree with (i) and (ii), which would be the case if $0 = \theta^H = \theta^A < \theta^S = 1$.

We now turn to Study 2, which tests the theoretical framework (Observation 1) and estimates the shares of the paternalistic preference types (Observation 2) based on both the spectators’ intervention decisions and their welfare evaluations.

4 Study 2

We first describe the experimental design, sample, and empirical strategy, before we present and discuss the main results of Study 2.

4.1 Experimental Design and Participants

Experimental Design. Study 2 replicates treatments *Hard*×*Internal* and *Soft*×*Internal* from Study 1, which, for short, we refer to as treatments *Hard* and *Soft* in the following. Study 2 adds a treatment, labeled *Welfare*. The context of treatment *Welfare* is identical to the context of treatments *Hard* and *Soft*. In contrast to treatments *Hard* and *Soft*, a spectator in treatment *Welfare* does not have the option to give the stakeholder the freedom to choose, but must allocate either the preferred safe option or the non-preferred risky option to the stakeholder. We use treatment *Welfare* to directly elicit whether a spectator considers the safe option or the risky option to promote the welfare of the stakeholder.

Participants. The spectators and stakeholders were recruited from the same populations as in Study 1, using the same procedures. Subjects who participated in Study 1 could not participate in Study 2. We sampled a total of 6,033 spectators in January 2020. The middle column of Table 1 reports the characteristics of the sample in Study 2, which is very similar to the sample in Study 1. Table A9 shows that the sample is largely balanced across treatments, with slightly fewer Republicans and slightly more spectators with a higher education and a higher income in treatment *Welfare*.

4.2 Empirical Strategy

To analyze the spectators' decisions in Study 2, we use the following empirical specification:

$$D_i = \beta_0 + \beta_1 H_i + \beta_2 S_i + \gamma X_i + \varepsilon_i \quad (4)$$

The dependent variable D_i is an indicator for whether spectator i intervenes in treatments *Hard* or *Soft*, or allocates the safe option in treatment *Welfare*. Treatment *Welfare* is the omitted category in the regression model. H_i is an indicator for spectator i being in treatment *Hard*, S_i is an indicator for spectator i being in treatment *Soft*, X_i is a vector of background characteristics (political orientation, willingness to take risks, education, income, age, and gender), and ε_i is an idiosyncratic error term. We estimate the model with and without the vector of background characteristics.

The regression model can be used to test Observation 1 ($\theta^H \leq \theta^A \leq \theta^S$). It follows from the regression model that $\hat{\theta}^H = \hat{\beta}_0 + \hat{\beta}_1$, $\hat{\theta}^S = \hat{\beta}_0 + \hat{\beta}_2$, and $\hat{\theta}^A = \hat{\beta}_0$. Hence, Observation 1 is rejected in the data if $\hat{\beta}_1 > 0$ or $\hat{\beta}_2 < 0$.

It follows from Observation 2 and the regression model that the estimated shares of libertarian paternalists and welfarists in our sample are given by: $\hat{\sigma}^{Wa} = \hat{\theta}^H = \hat{\beta}_0 + \hat{\beta}_1$, $\hat{\sigma}^{Wna} = 1 - \hat{\theta}^S = 1 - \hat{\beta}_0 - \hat{\beta}_2$, and $\hat{\sigma}^{LP} = \hat{\theta}^A - \hat{\theta}^H = -\hat{\beta}_1$. If all spectators are either libertarian paternalists or welfarists, then $\hat{\beta}_2 = 0$ and $\hat{\sigma}^{Wa} + \hat{\sigma}^{Wna} + \hat{\sigma}^{LP} = 1$.

4.3 Results

We start by providing an overview of the spectators' decisions. Figure 4 shows the share of spectators that intervene in treatments *Hard* and *Soft*, respectively, and the share that allocates the preferred safe option to the stakeholder in treatment *Welfare*. We observe from treatments *Hard* and *Soft* that the empirical pattern from Study 1 replicates in Study 2. About a third of the spectators (35.4 percent) implement the hard intervention and the large majority of the spectators (82.5 percent) implement the soft intervention. Study 2 thus replicates Result 1 from Study 1: the nature of an intervention has a strong causal impact on the spectators' willingness to intervene. Moreover, Figure 4 shows that 69.8 percent of the spectators allocate the safe option to the stakeholder in treatment *Welfare*, while 30.2 percent of the spectators allocate the risky option to the stakeholder.

[Figure 4 about here]

Table 3 reports the regression analysis. In column (1), we observe that the estimated share of spectators that intervene in treatment *Soft* is 12.7 percentage points higher

than the estimated share of spectators that allocate the safe option to the stakeholder in treatment *Welfare* ($p < 0.01$). We further estimate that the share of spectators that intervene in treatment *Hard* is 34.4 percentage points lower than the share of spectators that allocate the safe option in treatment *Welfare* ($p < 0.01$). Column (2) shows that the estimated treatment differences are virtually unaffected by the inclusion of the spectators' background characteristics.

[Table 3 about here]

The estimates in Table 3 are in line with Observation 1: (i) the estimated share of spectators that implement the hard intervention is strictly smaller than the estimated share of spectators whose welfare evaluations are aligned with the stakeholder's preferences ($\hat{\theta}^H < \hat{\theta}^A$), and (ii) the estimated share of spectators that implement the soft intervention is strictly larger than the estimated share of spectators whose welfare evaluations are aligned with the stakeholder's preferences ($\hat{\theta}^A < \hat{\theta}^S$).

Figure 5 shows that Observation 1 holds across subgroups of the general population. In each subgroup, we observe that the estimated share of spectators that implement the hard intervention is strictly smaller and the estimated share of spectators that implement the soft intervention is strictly larger than the estimated share of spectators whose welfare evaluations are aligned with the stakeholder's preferences (see Table A12).³

[Figure 5 about here]

We can now state the following result:

Result 4: *The spectators' intervention decisions are consistent with the theoretical framework, which provides evidence that the willingness to intervene is determined by how an intervention affects the stakeholder's welfare and freedom to choose.*

Given that the spectators' intervention decision are consistent with Observation 1, we can use the estimates in Table 3 to study the prevalence of the two main paternalistic types in our sample. Based on column (1), the estimated share of libertarian paternalists, $\hat{\sigma}^{LP}$, is 34.4 percent and the estimated share of welfarists, $\hat{\sigma}^{W_a} + \hat{\sigma}^{W_{na}}$, is 52.9 percent. The estimated share of welfarists whose welfare evaluations are aligned with the stakeholder's preferences, $\hat{\sigma}^{W_a}$, is 35.4 percent and the estimated share of welfarists

³Figure 5 also shows that the share of spectators allocating the preferred safe option to the stakeholder is largest among below-median risk takers and smallest among above-median risk-takers. This suggests that some stakeholders rely on their own preferences in their welfare evaluations (Ambuehl et al., 2021).

whose welfare evaluations are not aligned with the stakeholder’s preferences, $\hat{\sigma}^{Wna}$, is 17.5 percent. We summarize the analysis of the prevalence of the main paternalistic types as follows:

Result 5: *The large majority of the spectators are either libertarian paternalists or welfarists: about a third of the spectators are estimated to be libertarian paternalists and about half of the spectators are estimated to be welfarists.*

The estimation results imply that only the decisions of 12.7 percent of the spectators (given by $\hat{\theta}^S - \hat{\theta}^A$) cannot be explained by the two paternalistic types. These spectators cannot be welfarists because they care about the stakeholder’s freedom to choose and they cannot be libertarian paternalists because their welfare evaluations are not aligned with the stakeholder’s preferences. Overall, given that both these spectators and the libertarian paternalists value the stakeholder’s freedom to choose, we estimate that for almost half of the sample, 47.1 percent, the willingness to intervene is determined not only by how an intervention affects the stakeholder’s welfare but also by how an intervention affects the stakeholder’s freedom to choose.

Finally, Figure 5 shows that the prevalence of the two paternalistic types is quite similar across subgroups. In each subgroup, we find that about half of the spectators are welfarists, but also that a significant share of spectators are libertarian paternalists.

5 Conclusions

The paper studies paternalistic preferences in two large-scale experiments with participants from the general population in the United States. The experimental context captures the key characteristics of situations that create opportunities for paternalistic interventions: a stakeholder is about to make a choice that is not aligned with their own preferences, and a spectator can intervene to ensure that they receive their preferred option. Our experimental study provides, in two independent samples, causal evidence for the nature of an intervention being of great importance for the willingness to intervene. Only about a third of the spectators implement a hard intervention that removes the stakeholder’s freedom to choose, while a large majority implement a soft intervention that provides information without restricting the choice set. We find that this result holds regardless of the stakeholder’s responsibility for being mistaken about the choice set – whether the source of mistake is internal or external – and in different subgroups of the population.

We introduce a theoretical framework with two paternalistic types – libertarian paternalists and welfarists – and find that the behavior of the large majority of spectators can be rationalized by the two paternalistic preference types: about a third of the spec-

tators are estimated to be libertarian paternalists and about half of the spectators are estimated to be welfarists. The estimated share of libertarian paternalists reveals that a significant part of the population in the United States is reluctant to implementing hard interventions: they prefer leaving people “free to fail,” rather than enhancing their welfare by restricting their choice options.

To study whether people consider that this experimental context is relevant for the policy debate on paternalism in the United States, we also asked two general questions. First, we asked the spectators whether they agree that people sometimes make choices that are harmful to themselves. We find that the large majority of the spectators agree with the statement (see upper panel of Figure A1). Second, we asked the spectators whether they agree that the government can sometimes improve people’s lives by restricting their freedom to choose. We find sizable agreement with this view of the government but also that almost half of the spectators express some skepticism regarding the government’s ability to improve people’s lives with hard paternalistic policies (see lower panel of Figure A1). This skepticism could reflect general distrust in the government (e.g., Kuziemko, Norton, Saez, and Stantcheva, 2015) but also that people consider the freedom to choose to be an integral component of a good life. Interestingly, we find that the share of people disagreeing that the government can improve people’s lives with hard paternalism is very close to the estimated share of spectators that value the stakeholder’s freedom to choose in our experiment, 47.0 percent vs. 47.1 percent. Further, we find a strong positive association between abstaining from implementing the hard intervention in our experiment and disagreement with the view that the government can improve people’s lives by means of hard paternalism ($p < 0.001$). Taken together, we find evidence that the spectators perceive the experimental context of the present study to be of relevance and that the willingness to intervene in the experiment is predictive of people’s attitudes toward governmental paternalistic policies.

The experimental paradigm developed in this paper allows addressing a multitude of intriguing questions on paternalistic preferences. First, it is important to study the extent to which paternalistic preferences are domain-specific. Are paternalistic preferences different when stakeholders make choices over time than when they make choices in the domain of risk? Second, what is the role of the relation between the spectator and the stakeholder? A growing literature has focused on the economics of parenting (Francesconi and Heckman, 2016; Doepke and Zilibotti, 2017; Cobb-Clark, Salamanca, and Zhu, 2019; Agostinelli, Doepke, Sorrenti, and Zilibotti, 2020), which raises the question of whether people make different paternalistic considerations when acting as parents. Finally, it would be interesting to study cultural variation in paternalistic preferences, and the extent to which this variation can contribute to explain cross-country differences in people’s attitudes to paternalistic policies (Sunstein et al., 2018). Paternalistic interventions are prevalent across societies, and it is of great importance to understand how they are justified and relate to people’s own paternalistic preferences.

References


- Agostinelli, Francesco, Matthias Doepke, Giuseppe Sorrenti, and Fabrizio Zilibotti (2020). “It takes a village: The economics of parenting with neighborhood and peer effects,” Working Paper 27050, National Bureau of Economic Research.
- Alesina, Alberto, Stefanie Stantcheva, and Edoardo Teso (2018). “Intergenerational mobility and preferences for redistribution,” *American Economic Review*, 108(2): 521–554.
- Almås, Ingvild, Alexander W. Cappelen, and Bertil Tungodden (2020). “Cutthroat capitalism versus cuddly socialism: Are Americans more meritocratic and efficiency-seeking than Scandinavians?” *Journal of Political Economy*, 128(5): 1753–1788.
- Ambuehl, Sandro, Douglas B. Bernheim, and Axel Ockenfels (2021). “What motivates paternalism? An experimental study,” *American Economic Review*, 111(3): 787–830.
- Andreoni, James and John Miller (2002). “Giving according to GARP: An experimental test of the consistency of preferences for altruism,” *Econometrica*, 70(2): 737–753.
- Arad, Ayala and Ariel Rubinstein (2018). “The people’s perspective on libertarian-paternalistic policies,” *The Journal of Law and Economics*, 61(2): 311–333.
- Arneson, Richard J (1980). “Mill versus paternalism,” *Ethics*, 90(4): 470–489.
- Bartling, Björn, Ernst Fehr, and Holger Herz (2014). “The intrinsic value of decision rights,” *Econometrica*, 82(6): 2005–2039.
- Bellemare, Charles, Sabine Kröger, and Arthur van Soest (2008). “Measuring inequity aversion in a heterogeneous population using experimental decisions and subjective probabilities,” *Econometrica*, 76(4): 815–839.
- Beshears, John, James J. Choi, David Laibson, Brigitte C. Madrian, and William L. Skimmyhorn (2022). “Borrowing to save? The impact of automatic enrollment on debt,” *The Journal of Finance*, 77(1): 403–447.
- Bolton, Gary E. and Axel Ockenfels (2000). “ERC: A theory of equity, reciprocity, and competition,” *American Economic Review*, 90(1): 166–193.
- Camerer, Colin F., Samuel Issacharoff, George Loewenstein, Ted O’Donoghue, and Matthew Rabin (2003). “Regulation for conservatives: Behavioral economics and the case for asymmetric paternalism,” *University of Pennsylvania Law Review*, 151(3): 1211–1254.

- Cappelen, Alexander W., Astri Drange Hole, Erik Ø. Sørensen, and Bertil Tungodden (2007). “The pluralism of fairness ideals: An experimental approach,” *American Economic Review*, 97(3): 818–827.
- Cappelen, Alexander W., James Konow, Erik Ø. Sørensen, and Bertil Tungodden (2013). “Just luck: An experimental study of risk taking and fairness,” *American Economic Review*, 103(3): 1398–1413.
- Charness, Gary and Matthew Rabin (2002). “Understanding social preferences with simple tests,” *Quarterly Journal of Economics*, 117(3): 817–869.
- Cobb-Clark, Deborah A., Nicolas Salamanca, and Anna Zhu (2019). “Parenting style as an investment in human development,” *Journal of Population Economics*, 32(4): 1315–1352.
- Deci, Edward L. and Richard M. Ryan (1985). *Intrinsic Motivation and Self-Determination in Human Behavior*, New York: Springer.
- DellaVigna, Stefano and Elizabeth Linos (2022). “RCTs to scale: Comprehensive evidence from two nudge units,” *Econometrica*, 90(1): 81–116.
- Diepeveen, Stephanie, Tom Ling, Marc Suhrcke, Martin Roland, and Theresa M. Marteau (2013). “Public acceptability of government intervention to change health-related behaviours: A systematic review and narrative synthesis,” *BMC Public Health*, 13(756).
- Doepke, Matthias and Fabrizio Zilibotti (2017). “Parenting with style: Altruism and paternalism in intergenerational preference transmission,” *Econometrica*, 85(5): 1331–1371.
- Durante, Ruben, Louis Putterman, and Joël J. van der Weele (2014). “Preferences for redistribution and perception of fairness: An experimental study,” *Journal of the European Economic Association*, 12(4): 1059–1086.
- Dworkin, Gerald (1972). “Paternalism,” *The Monist*: 64–84.
- Evers, C., D.R. Marchiori, A. F. Junghans, J. Cremers, and D. T. D. De Ridder (2018). “Citizen approval of nudging interventions promoting healthy eating: The role of intrusiveness and trustworthiness,” *BMC Public Health*, 18(1): 1–10.
- Fehr, Ernst, Holger Herz, and Tom Wilkening (2013). “The lure of authority: Motivation and incentive effects of power,” *American Economic Review*, 103(4): 1325–59.
- Fehr, Ernst and Klaus M. Schmidt (1999). “A theory of fairness, competition and cooperation,” *Quarterly Journal of Economics*, 114(3): 817–868.

- Fong, Christina (2001). “Social preferences, self-interest, and the demand for redistribution,” *Journal of Public Economics*, 82(2): 225–246.
- Francesconi, Marco and James J. Heckman (2016). “Child Development and Parental Investment: Introduction,” *The Economic Journal*, 126(596): F1–F27.
- Gangadharan, Lata, Philip J. Grossman, Kristy Jones, and C. Matthew Leister (2018). “Paternalistic giving: Restricting recipient choice,” *Journal of Economic Behavior and Organization*, 151: 143–170.
- Hausman, Daniel and Michael McPherson (2009). “Preference satisfaction and welfare economics,” *Economics and Philosophy*, 25(1): 1–25.
- Jacobsson, Fredric, Magnus Johannesson, and Lars Borgquist (2007). “Is altruism paternalistic?” *Economic Journal*, 117(520): 761–781.
- Jung, Janice Y. and Barbara A. Mellers (2016). “American attitudes toward nudges.” *Judgment & Decision Making*, 11(1): 62–74.
- Kaplow, Louis and Steven Shavell (2000). “Fairness versus welfare,” *Harvard Law Review*, 114(961).
- Kaplow, Louis and Steven Shavell (2001). “Any non-welfarist method of policy assessment violates the pareto principle,” *Journal of Political Economy*, 109(2): 281–286.
- Kaplow, Louis and Steven Shavell (2006). *Fairness versus welfare*, Cambridge, Massachusetts: Harvard University Press.
- Kassirer, Samantha, Emma E Levine, and Celia Gaertig (2020). “Decisional autonomy undermines advisees’ judgments of experts in medicine and in life,” *Proceedings of the National Academy of Sciences*, 117(21): 11368–11378.
- Kiessling, Lucas, Shyamal Chowdhury, Hannah Schildberg-Hörisch, and Matthias Sutter (2021). “Parental paternalism and patience,” IZA DP No. 14030.
- Konow, James (2000). “Fair shares: Accountability and cognitive dissonance in allocation decisions,” *American Economic Review*, 90(4): 1072–1091.
- Kuziemko, Ilyana, Michael I. Norton, Emmanuel Saez, and Stefanie Stantcheva (2015). “How elastic are preferences for redistribution? Evidence from randomized survey experiments,” *American Economic Review*, 105(4): 1478–1508.
- Le Grand, Julian and Bill New (2015). *Government Paternalism: Nanny State Or Helpful Friend?*, Princeton, NJ: Princeton University Press.
- List, John A., Matthias Rodemeier, Sutanuka Roy, and Gregory K. Sun (2023). “Judging nudging: Understanding the welfare effects of nudges versus taxes,” Working Paper 311520, National Bureau of Economic Research.


- Nussbaum, Martha (2001). “Symposium on amartya sen’s philosophy: Adaptive preferences and women’s options,” *Economics and Philosophy*, 17(1): 67–88.
- OECD (2017). “Behavioural insights and public policy: Lessons from around the world,” Technical report, OECD Publishing.
- Owens, David, Zachary Grossman, and Ryan Fackler (2014). “The control premium: A preference for payoff autonomy,” *American Economic Journal: Microeconomics*, 6(4): 138–161.
- Pikulina, Elena S. and Chloe Tergiman (2020). “Preferences for power,” *Journal of Public Economics*, 185: 104173.
- Reisch, Lucia A. and Cass R. Sunstein (2016). “Do Europeans like nudges?” *Judgment and Decision Making*, 11(4): 310–325.
- Romano, Joseph P. and Michael Wolf (2005). “Stepwise multiple testing as formalized data snooping,” *Econometrica*, 73(4): 1237–1282.
- Romano, Joseph P. and Michael Wolf (2016). “Efficient computation of adjusted p -values for resampling-based stepdown multiple testing,” *Statistics & Probability Letters*, 113(1): 38–40.
- Sen, Amartya (2004). *Rationality and Freedom*, Harvard University Press.
- Sloof, Randolph and Ferdinand A. von Siemens (2017). “Illusion of control and the pursuit of authority,” *Experimental Economics*, 20: 556–573.
- Sunstein, Cass R., Lucia A. Reisch, and Julius Rauber (2018). “A worldwide consensus on nudging? Not quite, but almost,” *Regulation & Governance*, 12(1): 3–22.
- Tannenbaum, David, Craig R. Fox, and Todd Rogers (2017). “On the misplaced politics of behavioral policy interventions,” *Nature Human Behaviour*, 1: 0130.
- Thaler, Richard H. and Cass R. Sunstein (2003). “Libertarian paternalism,” *American Economic Review*, 93(2): 175–179.
- Thaler, Richard H. and Cass R. Sunstein (2021). *Nudge: The Final Edition: Improving Decisions About Money, Health, and the Environment*, New Haven: Yale University Press.
- US Census Bureau (2018). “Data from www.census.gov, 2018–2020,” Library Catalog: www.census.gov Section: Government.

Figure 1: Sequence of Events in Study 1




Notes: The figure shows the events that take place in each stage of Study 1.

Figure 2: Spectator Decisions by Treatment — Study 1




Note: The figure shows the share of spectators that intervene by treatment. The left panel shows the share of spectators intervening in treatments *Hard*×*Internal* and *Soft*×*Internal*. The right panel shows the share of spectators intervening in treatments *Hard*×*External* and *Soft*×*External*. The black bars indicate standard errors.

Figure 3: Heterogeneous Treatment Effects — Study 1




Note: The figure shows the estimated coefficients and 95% confidence intervals for the interaction effects, both in the single interaction models and the joint interaction model. The left panel shows the estimated interaction effect between an indicator for the spectator being in treatment $Soft \times Internal$ or $Soft \times External$ and an indicator for the respective background characteristic. The right panel shows the estimated interaction effect between an indicator for the spectator being in treatment $Soft \times External$ or $Hard \times External$ and an indicator for the respective background characteristic. See Tables A4 and A5 in Appendix A for the underlying regression results.

Figure 4: Spectator Decisions by Treatment — Study 2



Note: The left bar shows the share of spectators intervening in treatment *Hard*. The middle bar shows the share of spectators allocating the preferred safe option to the stakeholder in treatment *Welfare*. The right bar shows the share of spectators intervening in treatment *Soft*. The black bars indicate standard errors.

Figure 5: Spectator Decisions by Subgroup — Study 2



Note: The figure shows, for different subgroups in the general population, the share of spectators that intervene in treatment *Hard*, the share of spectators allocating the preferred safe option to the stakeholder in treatment *Welfare*, and the share of spectators that intervene in treatment *Soft*. See Table A12 for the underlying regression results.

Table 1: Sample Descriptives

	Study 1		Study 2		U.S. Population
	Mean	SD	Mean	SD	Mean
Female	0.51	(0.50)	0.51	(0.50)	0.51
Age 18–34	0.30	(0.46)	0.30	(0.46)	0.31
Age 35–44	0.18	(0.39)	0.17	(0.38)	0.18
Age 45–54	0.18	(0.38)	0.19	(0.39)	0.19
Age 55–64	0.16	(0.37)	0.16	(0.37)	0.16
Age 65–	0.18	(0.38)	0.18	(0.38)	0.17
Edu: Highschool or less	0.31	(0.46)	0.27	(0.44)	0.37
Edu: Some College	0.21	(0.41)	0.23	(0.42)	0.20
Edu: Bachelor or Associate	0.33	(0.47)	0.39	(0.49)	0.30
Edu: Master or above	0.15	(0.36)	0.12	(0.33)	0.14
Income < 30,000	0.26	(0.44)	0.26	(0.44)	0.25
Income 30–60,000	0.26	(0.44)	0.30	(0.46)	0.25
Income 60–100,000	0.23	(0.42)	0.24	(0.43)	0.22
Income 100–150,000	0.14	(0.35)	0.13	(0.34)	0.14
Income > 150,000	0.11	(0.32)	0.08	(0.27)	0.14
Republican	0.29	(0.45)	0.31	(0.46)	
Democrat	0.33	(0.47)	0.35	(0.48)	
Independent/Third Party	0.28	(0.45)	0.27	(0.44)	
Risk Taking: low (0–4)	0.23	(0.42)	0.22	(0.41)	
Risk Taking: median (5)	0.13	(0.34)	0.12	(0.33)	
Risk Taking: high (6–10)	0.64	(0.48)	0.66	(0.47)	
Observations	8,004		6,033		

Notes: Sample descriptives for spectators in Study 1 and Study 2. We asked the spectators to identify as either male or female, and we elicited the exact year of age. Education was elicited using the categories Less than High School, High School/GED, Some College, Associate’s Degree, Bachelor’s Degree, Master’s Degree, Professional Degree (JD, MD), and Doctoral Degree. Income was elicited using the income brackets as shown in the table. The spectators were asked to identify as either “Republican,” “Democrat,” or “Independent/Third Party,” and they had the option not to answer this question. “Risk Taking” was measured on a scale from 0 to 10, with 0 indicating “Completely unwilling to take risks” and 10 indicating “Very willing to take risks.” We benchmark our sample composition against values for the population in the United States taken from the U.S. Census Bureau (2018).

Table 2: Regression Results — Study 1

	(1)	(2)	(3)	(4)	(5)	(6)
Soft Intervention	0.550*** (0.009)	0.550*** (0.009)	0.524*** (0.013)	0.524*** (0.013)		
External Source			-0.035** (0.015)	-0.035** (0.015)	-0.009 (0.011)	-0.008 (0.011)
Soft × External			0.051*** (0.018)	0.051*** (0.018)		
Republican		0.004 (0.010)		0.003 (0.010)		0.009 (0.012)
High Risk Taking		-0.031*** (0.009)		-0.031*** (0.009)		-0.036*** (0.011)
High Education		0.016 (0.010)		0.016 (0.010)		0.020* (0.012)
High Income		-0.026** (0.010)		-0.026*** (0.010)		-0.025** (0.012)
High Age		-0.003 (0.010)		-0.003 (0.010)		-0.008 (0.012)
Female		0.017* (0.010)		0.016* (0.010)		0.009 (0.011)
Constant	0.318*** (0.007)	0.329*** (0.013)	0.335*** (0.011)	0.347*** (0.015)	0.596*** (0.008)	0.612*** (0.015)
Observations	8,004	8,004	8,004	8,004	8,004	8,004
R^2	0.313	0.315	0.314	0.316	0.000	0.003

Notes: The table reports OLS regressions. The dependent variable is an indicator for whether the spectator chooses to intervene. The data from all four treatments is included. “Soft Intervention” is an indicator for the spectator being in a treatment with the soft intervention, “External Source” is an indicator for the spectator being in a treatment where the source of mistake is external. “Soft×External” is the interaction between these two variables. “Republican” is an indicator for identifying with the Republican party. “High Risk Taking,” “High Education,” “High Income,” and “High Age” are indicator variables for having above-median willingness to take risks, education, income, and age, respectively. “Female” is an indicator for being female. The results are robust to adjusting for multiple-hypothesis testing and to using Probit models (see Tables A2 and A3 in Appendix A). Robust standard errors in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$

Table 3: Regression Results — Study 2

	(1)	(2)
Soft Intervention	0.127*** (0.013)	0.128*** (0.013)
Hard Intervention	-0.344*** (0.015)	-0.344*** (0.015)
Republican		0.012 (0.012)
High Risk Taking		-0.094*** (0.012)
High Education		0.030** (0.012)
High Income		-0.009 (0.013)
High Age		-0.015 (0.012)
Female		0.050*** (0.012)
Constant	0.698*** (0.010)	0.711*** (0.017)
Observations	6033	6033
R^2	0.169	0.182

Notes: The table reports OLS regressions. The dependent variable is an indicator for whether the spectator chooses to intervene or select the safe option. Treatment *Welfare* serves as omitted category in models (1) and (2) that include the data from all three treatments. Model (3) comprises only the data from treatment *Welfare*. “Soft Intervention” and “Hard Intervention” are indicators for the spectator being in treatment *Soft* and *Hard*, respectively. “Republican” is an indicator for identifying with the Republican party. “High Risk Taking,” “High Education,” “High Income,” and “High Age” are indicator variables for having above-median willingness to take risks, education, income, and age, respectively. “Female” is an indicator for being female. The results are robust to adjusting for multiple-hypothesis testing and to using Probit models (see Tables A10 and A11 in Appendix A). Robust standard errors in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$