Erasmus Journal for Philosophy and Economics,
Volume 18, Issue 2,
Winter 2025, pp. 1–37
https://doi.org/10.23941/ejpe.v18i2.1012

Do People See a Problem in the Non-Identity Problem? An Experimental Exploration Using the Dictator Game

ARNE NASGOWITZ
Centre for Applied Research at NHH (SNF), FAIR

PABLO SOTO-MOTA
Centro de Estudios Económicos, El Colegio de México

Abstract: Our actions may determine the potential well-being of future individuals, but also their number and identity. The Non-Identity Problem, a central topic in Population Ethics with broad implications for Economics and Policy, challenges the notion that we must act to prevent harm to future people. In this study, we use variations of the dictator game and vignettes to explore non-experts' intuitions about the problem. We find that participants' choices align with the No-Difference view, holding that moral responsibility is independent of who is affected.

Keywords: Non-Identity Problem, Population Ethics, Experimental Ethics, Dictator Game

JEL Classification: A13, C91, D63, D64

I. INTRODUCTION

"The eyes of all future generations are upon you. And if you choose to fail us, I say to you: We will never forgive you". With these words, Swedish activist Greta Thunberg (2019) warned participants at the United Nations Climate Action Summit about the consequences of inaction against climate change. Although at the beginning of her speech she referred mostly to the dire consequences of inaction for her own generation, her appeal ultimately included harm to *all future generations*, including individuals not yet born. This intuition—that failing to prevent problems like climate

AUTHORS' NOTE: We are grateful to two anonymous reviewers for their careful reading and constructive suggestions. We also thank Bertil Tungodden, Alexander Cappelen, Mathias Ekström, Akshay Moorthy, Ernesto Gavassa Pérez, and Jonas Pilgaard Kaiser for valuable comments on the early ideas that motivated this paper. Finally, we thank the participants at the 2022 Interdisciplinary Research Network for Economists and Philosophers for helpful feedback and discussion.

change harms future individuals—is widely shared. In a recent poll by Ipsos (2025), only 13% of respondents disagreed with the statement "If individuals like me do not act now to combat climate change, we will be failing future generations". Likewise, some philosophers have argued in favour of *longtermism*: the ethical claim that since future generations may span many millennia and their numbers are much larger than our generation, our moral responsibility toward them is significantly amplified (Greaves and MacAskill 2021). Recent research indicates that people who endorse longtermism express greater concern for climate change, greater support for pro-environmental policies, and greater engagement in pro-environmental behaviours (Syropoulos et al. 2024; Law et al. 2023). Related findings of Ogunbode et al. (2024), based on a global survey in eleven countries, show that most participants support the imperative to prevent harm to future generations, and that these beliefs strongly predict both environmental activism and everyday sustainable actions.

Despite its seeming universality, the idea that we have a moral obligation to avoid harming future individuals confronts a serious philosophical challenge known as the *Non-Identity Problem*. This problem arises from the observation that the consequences of our actions differ significantly for present and future individuals and from the widely held assumption that to harm someone is to make them worse off than they would otherwise have been. Our decisions can certainly affect some aspects of the well-being of people alive today, but future individuals depend entirely on us. Their environment, economic conditions, cultural context, and even the possibility of their existence depend on what we do today. Furthermore, because environmental and socioeconomic conditions influence reproductive choices, our actions can affect how many, when, and how future individuals will be born, directly determining their genetic identity. If the identities of future individuals are not fixed, to whom do we have moral responsibilities?

As an example of this, consider our decision on whether to deplete or conserve important natural resources. It is likely that depleting natural resources today would significantly decrease the quality of life of future people, so we would normally say that doing so is wrong. However, our choice whether to deplete or conserve the natural resources would lead to drastically different lifestyles, consumption patterns, and social relationships. Such differences would eventually lead to a change in reproductive decisions, inevitably resulting in different individuals coming into existence, each with distinct genetic identities, experiences, and

environments depending on our initial choice between conservation and depletion. Suppose that, in fact, we deplete those resources and, as a consequence, centuries from now, people face scarcity and live much worse than we do today. People in that future might echo Greta Thunberg's words, saying to us: "You have chosen to fail us. I say to you: We will never forgive you". Would they be justified in making this accusation? If the possibility of their lives depended on us depleting resources, how could they claim that we harmed them by doing so? If we had conserved the resources, they would not complain, though not because they would not face scarcity, but because they would not exist. Other individuals, with different genetic identities, experiences, and contexts, would be the ones living. Thus, to whom do we have moral obligations to protect the environment?

The Non-Identity Problem, first articulated by Derek Parfit in his seminal book *Reasons and Persons* (1984), is the observation that in identity-affecting situations—cases where actions determine the existence and identity of affected individuals—conventional notions of harm and seemingly reasonable utilitarian assumptions make it difficult to justify moral obligations. Since its inception, the problem has been central to Population Ethics, a field of philosophy that examines questions related to the existence, number, and welfare of present and future populations (Pivato 2023). The relevance of the problem is constantly discussed in debates on reproductive rights (Hall 2024; Magnusson 2020; Wrigley 2008), animal welfare (O'Brien 2021; Ferrari 2012; Palmer 2011), affirmative action (Lippert-Rasmussen 2020), environmental policy (Tank 2022), among many other topics.

Parfit's initial answer to address the Non-Identity Problem was what he called a *No-Difference view*. This is to directly deny that changes in the identity of those receiving the consequences of our actions are of moral significance. Returning to our initial example, the No-Difference view implies that the fact that those facing difficult conditions in future centuries would never have been born if we had acted differently does not make depleting resources any less wrong. Parfit and others after him have sought philosophical grounds to defend this intuition and reconcile it with broader moral theories, which typically rely on someone being negatively affected to judge an action as wrong. However, the No-Difference view remains a matter of ongoing debate among experts, and many philosophers reject it on theoretical grounds (for example, Roberts 2024; Boonin 2008, 2020). At the same time, although laypeople frequently

express moral judgments that align with the No-Difference view, it is unclear whether these judgments stem from an explicit commitment to it or from alternative moral intuitions. Clarifying this distinction is a central motivation of the present study.

In this paper, we investigate whether non-philosophers hold No-Difference views. We build upon the work of Kopec and Bruner (2022), who examined this question by designing an identity-affecting version of the Dictator Game. After partially replicating their experiment, we introduce a new comparison scenario to control for alternative mechanisms to their observations, such as compassion fade and the unidentifiable victim effect. Further, we explore participants' intuitions on identity-affecting scenarios using vignettes. Contrary to the findings of Kopec and Bruner (2022), in our experiment, participants did not behave more selfishly in the non-identity version of the Dictator Game compared to the standard presentation. We interpret these results as evidence of the prevalence of No-Difference views among participants. Moreover, our results indicate that participants generally reject practical arguments based explicitly on Non-Identity considerations. Although our findings diminish the relevance of the Non-Identity Problem for policy, we argue that they highlight the importance of future theoretical and experimental exploration, both in philosophy and economics, into the rationale behind No-Difference views and alternative conceptions of harm.

The rest of this paper proceeds as follows. In Section 2, we introduce the Non-Identity Problem and present No-Difference views as a potential solution. Section 3 focuses on the empirical literature exploring whether individuals hold No-Difference views. In Section 4, we describe our experiment, hypotheses, and sample. Section 5 describes our results, and finally we conclude in Section 6, discussing the implications of our findings.

II. WHAT IS THE NON-IDENTITY PROBLEM?

Since its introduction by Parfit (1984), the explanations of the Non-Identity Problem usually start with the contrast of an *identity-preserving situation*—a situation where no matter whether an action is made or not the same people exist—to an *identity-affecting situation*—a situation where an action either changes the number or identity of existing people. Then, a series of seemingly reasonable assumptions is presented only to notice that the assumptions are good enough to justify intuitive moral claims in identity-preserving situations but that they clash with our intuitions in

identity-affecting situations. We follow here the same way of presenting the problem.

Identity-preserving situation: Imagine that the government of a country decided last year to extract all the natural resources of a poor region, let us call it *Region A*, and to put all the pollutants of the country in that same region. As a consequence, people in *Region A* are now impoverished and live in a polluted environment.

Identity-affecting situation: this is the one we present in the Introduction to this paper. If we decide to deplete the resources of the planet and pollute it, in some centuries, people will have a low quality of life and will live in a polluted place. If we decide otherwise and preserve the planet's resources, future people will have a high quality of life and live in a non-polluted place. However, we observe that depleting or preserving the resources of the planet demands very different lifestyles, and thus, our decision would lead to different reproductive choices. Consequently, over time, different individuals would be born, leading to a situation where the whole population would be different if we deplete rather than preserve the resources of the planet.

We need a few assumptions to justify the claim that the government in the identity-preserving situation did something morally wrong. If last year the government had not decided to extract resources from *Region A* and pollute it, *Region A*'s population would have a better quality of life than the one they have today. They harmed individuals of *Region A* and harming someone is wrong. Harming many people is arguably worse.

We may want to say a similar argument about the depletion of resources and the polluting of the planet in the second scenario —the identity-affecting situation. However, we cannot easily follow the same argument as before. We cannot say that we harm future people by depleting their resources in the same way as we say that the government harms the population in *Region A*. The reason is that, contrary to people in *Region A*, the future people in the impoverished and polluted planet would not exist other than under those conditions. In other words, the difference is that in the identity-preserving situation by depleting resources, the government harms a specific group of people—the residents of *Region A* —, while in the identity-affecting situation, our decision changes who exists, so if we deplete the resources, future people cannot truly say *they* would be better off if we did the alternative. So, the problem emerges from the clash between our intuitive judgments in the cases and the following seemingly reasonable assumptions.

A: Counterfactual notion of harm: An act harms someone if it makes that person worse off than they would have been had the act not occurred.

B: Person-Affecting view: "An act cannot be wrong if this act would be worse for no one" (Parfit 2017, 119).[1]

Notice that the counterfactual notion of harm implies that no one is harmed by being prevented from existing. The reason is that someone who does not exist is not someone who can be made better or worse off. Likewise, it implies that a person is not harmed by the fact that we make them exist, even if their well-being is low, provided that their life is worth living. Again, the reason is that it is not possible to say that such a person would have been better off by not existing. Finally, given the Person-Affecting view, since making a person exist with a low quality of life makes no one worse off, it follows that such an act does not harm anyone and therefore is not morally wrong. This conclusion contrasts with common moral intuitions. Indeed, several studies find that people judge actions such as environmental pollution to be morally wrong precisely because they believe these actions harm future individuals (Ipsos 2025; Syropoulos et al. 2024; Ogunbode et al. 2024). The Non-Identity Problem consists in determining if what is wrong is our moral assumptions or our intuitive judgements in the identity-affecting cases.

There are several possible answers to the Non-Identity Problem.[2] In general, there are those who either accept the apparent conclusion of the problem and say that we cannot say that harming future individuals is *wrong* and those who say that but deny the counterfactual notion of harm or Person-Affecting views. In the first group, we find, for example Boonin (2008) who says that the Non-Identity Problem is, in fact, a Non-Identity *Argument* and that we should agree with its conclusion. He argues that our resistance to accept the consequence that we do nothing wrong by knowingly creating people with a lower quality of life instead of other individuals with a higher quality of life is there only because "our moral sentiments are such that they produce misleading results" in identity-affecting cases.[3] Likewise, Weinberg (2008) attempts to dissolve the problem by arguing that *merely possible people* are morally irrelevant.

[1] Parfit refers to this principle as the *Narrow Deontic Principle*, one among several narrow person-affecting principles. We refer to it here simply as the *Person-Affecting view*, as this is how it is most commonly described in the literature (Roberts 2024).

[2] For in-depth discussions of the potential answers to the Non-Identity Problem see the entry by Roberts (2024) at the *Stanford Encyclopedia of Philosophy* or the book *The Non-Identity Problem and the Ethics of Future People* by Boonin (2020).

[3] By saying this Boonin is not claiming that we have no moral obligations to future generations such as taking care of the environment. He explains that in almost all practical

In the second group, some have discussed alternatives to the counterfactual notion of harm (Gardner 2015; Rivera-Lopez 2009; Harman 2004). Finally, others have argued to deny or qualify the Person-Affecting view (McMahan 2013; Steinbock 2009). In fact, since he presented the problem for the first time, Parfit (2017, 1984) defended the idea to avoid the non-identity problem by directly rejecting the Person-Affecting view and defended what he called:

No-Difference View: The morality of actions does not depend on the specific individuals who actually exist. The identity-affecting characteristics of a situation should not change the moral evaluation of our actions.

If we hold No-Difference views, the Non-Identity Problem dissolves. Moral evaluation then focuses solely on the overall consequences of whichever individuals exist independently of whether they would have existed if we had acted in a different way. Consequently, conserving resources becomes a clear moral obligation.

In a conference on Effective Altruism in 2016, Parfit lamented that his initial presentation of the problem motivated many philosophers to view it as an obstacle to justifying our environmental duties. Moreover, he commented that back when he wrote *Reasons and Persons* (1984) he wrongly expected that the No-Difference view would be accepted by most people, yet he observed, perhaps correctly, that most in the philosophical debates implicitly or explicitly rejected it. This article asks whether Parfit's prediction of the broad acceptance was right, if not among philosophers, among non-experts. That is, we ask whether individuals hold No-Difference views and, consequently, fail to see a problem in identity-affecting cases. In particular, we study individuals' distributive decisions in identity-preserving and identity-affecting situations using variations of the dictator game. We also examine how individuals evaluate arguments based on the Non-Identity Problem that are commonly used to reduce moral responsibility in applied contexts. Finally, we investigate whether behaviour in the experimental tasks is related to the perceived acceptability of such arguments.

As Kopec and Bruner (2022) emphasize, the moral intuitions of the general public speak directly to the practical relevance of the Non-Identity Problem, even if they do not contribute to its theoretical resolution. Many of the most pressing contexts in which the Non-Identity Problem arises,

situations, the assumptions of the problem are not met. For example, it is not often the case that making people with a low quality of life exist affects no one else other than them.

such as environmental degradation, climate policy, or reproductive health, require collective action, regulatory change, or long-term policy commitments. In democratic societies, such changes ultimately depend, directly or indirectly, on public acceptance.

While philosophers typically analyse the Non-Identity Problem on theoretical grounds rather than on intuitiveness, arguments invoking the problem may be used to discuss policy proposals. Whether these arguments are persuasive outside academic philosophy therefore depends on whether the underlying moral considerations resonate with nonexperts. If individuals do not find identity-affecting considerations morally salient, or if they instead hold No-Difference views, then the Non-Identity Problem may have limited relevance in real-world policy debates, regardless of its philosophical significance. From this perspective, investigating whether laypeople endorse No-Difference views is not meant to decide between competing philosophical positions, but rather to assess the extent to which the Non-Identity Problem constitutes a genuine obstacle to moral motivation and public justification in applied and policy-oriented contexts.

III. PREVIOUS EMPIRICAL LITERATURE AND CONTRIBUTION

To the best of our knowledge, only two studies have attempted to explore whether people hold No-Difference views. Our study was inspired by and builds directly on both of them.

In this section, we describe these studies in detail.

In the first study, Doolabh et al. (2019) conducted an online survey which investigated the perceived moral relevance of the Non-Identity Problem using a real-world example: the Zika virus epidemic, which primarily affects newborn babies. In their study, the sample was a group of educated individuals—50% of whom were philosophers—who self-selected to participate through an ad posted in an online philosophy magazine. Participants were asked to evaluate different public health interventions against Zika, some targeting newborns directly, others aimed at delaying pregnancies to reduce future infections. In addition to eliciting preferences over policy interventions, the authors asked participants how important the Non-Identity Problem was to their reasoning. They found evidence suggesting that most people do not find the Non-Identity Problem morally relevant, either because they do not fully understand it or because they rely on non-counterfactual notions of harm. In particular, although a majority of participants favoured interventions that benefit

existing newborns over those that benefit potential future individuals, they did so primarily because they viewed these interventions as more effective, not because they endorsed Person-Affecting views.

In the second related study, Kopec and Bruner (2022) designed an identity-affecting version of the Dictator Game to test whether non-experts make different distributive decisions compared to an identity-preserving situation, such as the traditional Dictator Game. In their novel setup, each dictator, as in the standard version, decided how much to share from a $1 endowment. However, the amount shared determined the identity of the recipient: each dictator was matched with eleven potential recipients, each associated with a specific amount between $0 and $1, in 10-cent increments. Only the recipient matched to the chosen amount received the money. Those not selected were not informed of the dictator's choice, thereby creating an identity-affecting condition. Moreover, the authors asked both dictators and recipients whether they believed the dictator's decision caused harm. Each treatment involved 177 participants, who received $0.50 as a participation fee.

Their findings suggest that participants do perceive a difference between identity-affecting and -preserving situations, thus rejecting No-Difference views. In the traditional Dictator Game, participants shared an average of 23.38%, a figure consistent with prior findings (Engel 2011). In contrast, in the identity-affecting version, participants shared only 15.5%, representing a 33.7% decrease (7.88 percentage points) in the amount shared. Furthermore, while 32% of participants in the traditional version kept the entire endowment, 45% did so in the identity-affecting condition. Most notably, participants in the identity-affecting treatment overwhelmingly reported that sharing nothing did not constitute harm, whereas the opposite view was more prevalent in the traditional setting.

Kopec and Bruner (2022) interpret their results as evidence that the Non-Identity Problem "likely does pose a potentially serious obstacle to moral motivations on problems like climate change" (186). However, they note that their findings could also be explained by *compassion fade*, due to the increased number of potential recipients in the identity-affecting version of their experiment. Compassion fade is a psychological phenomenon in which individuals act less altruistically—and thus more selfishly—as the number of people in need increases (Grodeck et al. 2022). Proposed mechanisms for compassion fade include the idea that, as the number of victims rises, individuals pay less attention to each one, reducing feelings of personal guilt for not helping (Tomm et al. 2019). Another

related mechanism is the identifiable victim effect: the observation that altruism tends to be stronger when the recipient is a clearly identifiable individual rather than an anonymous one, which becomes less likely as the size of the group increases (Lee and Feeley 2016; Jenni and Loewenstein 1997).

We observe, however, another possible mechanism. In the traditional Dictator Game, the dictator can directly eliminate inequality between themselves and the recipient by sharing half of the initial endowment. However, in the identity-affecting scenario, the distributive decision of the dictator has implications for the implemented inequality not only between her and the actual receiver but also within the potential pool of receivers or the group as a whole. Since individuals may care about distributive fairness (Cappelen et al. 2007), these considerations might confound the effect of introducing an identity-affecting element into the decision-making process. In fact, models that rationalize altruistic preferences would predict, without referring to the Non-Identity Problem, that individuals would share less money in the identity-affecting version of the Dictator Game (Andreoni and Miller 2002).

Our study contributes to the literature in two main ways. First, it revises the findings of Kopec and Bruner (2022) by increasing the sample size and disentangling the effects of compassion fade and group-level inequality. We implement a novel design in which dictators face a similar group of eleven potential recipients, but in a context that does not involve an identity-affecting situation. Second, following Doolabh et al. (2019), we directly ask participants about their moral judgments in identity-affecting scenarios and their acceptance of arguments based on the Non-Identity Problem. However, unlike their study, we use a sample of non-experts who have likely never encountered the Non-Identity Problem before. Thus, our second contribution is to provide a test of whether distributive decisions in a lab setting correlate with participants' underlying moral intuitions.

It is important to clarify that, like the studies discussed in this section, ours should be understood as an experimental exploration of moral intuitions among non-philosophers. That is, it is descriptive in nature—we investigate what people actually believe—and we fully abstain from making normative claims or engaging in philosophical analysis of potential solutions to the problem.

IV. THE EXPERIMENT

Our study consists of two parts, guided by two main research questions. In the first part, we investigate whether people make different distributive decisions in identity-affecting situations compared to an identity-preserving situation using variations of the Dictator Game. In our identity-affecting situations, we replicate the study by Kopec and Bruner (2022) and introduce a modified version of the Dictator Game, with which we aim to isolate the effects of the non-identity scenario from alternative mechanisms such as compassion fade and the identifiable victim effect. In the second part, we elicit individuals' moral intuitions in response to some real-world scenarios and relate them to their behaviour in the lab to observe whether their moral intuitions correlate with their distributive behaviour. Our study received approval from the IRB at our institution. All materials, including instructions and data, are publicly available in the online appendix to this paper.

IV. I. EXPERIMENTAL DESIGN

In our primary experimental task, participants made three distributive choices within the framework of the Dictator Game. We randomized the order of the three different distributive decisions for each participant. Participants were informed that one of their distributive choices would be implemented at random. The instructions for each task were accompanied by a graphical explanation that can be consulted in the Appendix to this paper. The details of each variation of the Dictator Game are as follows:

Traditional DG: Each participant, acting as a *dictator*, receives an endowment of 1 USD and is asked to share any amount between 0 and 1 USD with an anonymous individual who did not receive such an endowment.

Identity-Affecting DG: This variation is directly inspired by Kopec and Bruner (2022). In it, each participant, acting as a *dictator*, receives an endowment of 1 USD and decides how much to share with an anonymous individual. However, we introduce 11 potential recipients, each randomly matched with a possible transfer amount going from \$0 to \$1 in 10-cent increments. For example, one potential recipient may be matched to \$0.10, another to \$0.80. If the dictator decides to share \$0.10, only the recipient matched to \$0.10 receives the money; the others receive nothing. Crucially, dictators know that recipients who receive nothing will never learn that they were part of the study. Thus, in this variation,

dictators are in an identity-affecting situation: their decision not only determines the allocation of resources but also who is effectively part of the experiment and experiences its consequences.

Random-Recipient DG: In this variation, the dictator chooses how to allocate their endowment of 1 USD between themselves and a recipient who initially received nothing. However, the recipient is randomly selected from a pool of eleven potential recipients *after* the participant makes their decision.

Kopec and Bruner (2022) argue that their Identity-Affecting Dictator Game places participants in an identity-affecting situation, insofar as the dictator's decision determines not only the allocation of resources but also which individual becomes part of the experiment and experiences its consequences. On this basis, they interpret a comparison between the Identity-Affecting DG and the Traditional DG as a test of whether participants regard identity-affecting features as morally relevant.

However, as discussed in the previous section, there are alternative mechanisms that could plausibly explain their findings, including compassion fade, the identifiable victim effect, and concerns about inequality within a larger pool of potential recipients. These alternative mechanisms arise from a structural difference between the experimental implementation of an identity-affecting situation and the paradigmatic identity-affecting cases discussed in the philosophical literature. In the experiment, all potential recipients already exist, and dictators are aware of this. As a result, participants may plausibly hold altruistic, fairness-based, or inequality-averse preferences over the entire pool of recipients. By contrast, in paradigmatic identity-affecting cases, such as those involving future generations, the pool of potential beneficiaries does not yet exist, and we, existing individuals, are aware that their actions determine not only outcomes, but who will exist at all.

For this reason, while the Identity-Affecting DG shares an important structural feature with philosophical identity-affecting cases, namely, that the identity of the beneficiary is contingent on the agent's decision, it should be understood as an experimental device rather than as a literal instantiation of non-existence. Its purpose is to isolate and operationalize a specific decision-theoretic feature of identity-affecting situations under controlled conditions. We introduce the Random-Recipient DG precisely to disentangle the identity-contingency of the decision from the presence of multiple potential recipients. Although the Random-Recipient DG involves the same number of potential recipients as the Identity-Affecting

DG, the identity of the recipient is determined independently of the dictator's choice. As such, it does not constitute an experimental analogue of an identity-affecting situation, but it allows us to isolate the behavioural effects of a larger recipient pool.

The Random-Recipient DG should therefore be understood, like the Identity-Affecting DG, as an experimental device rather than an attempt to more closely approximate real-world identity-affecting scenarios. Its value lies in separating potential mechanisms. Comparing the Random-Recipient DG to the Traditional DG identifies the effect of introducing a pool of potential recipients. Comparing the Identity-Affecting DG to the Random-Recipient DG isolates the residual effect of making the recipient's identity contingent on the dictator's decision.

The second part of the study consisted of survey questions on four different real-life scenarios involving various aspects of identity-affecting intuitions related to the non-identity problem. Our goal with these questions was not to test participants' grasp of the concepts around the Non-Identity Problem as a philosophical problem. Rather, we used them to elicit moral intuitions in situations that, though not necessarily identity-affecting in a strict sense, reflect aspects of the type of thinking in the Non-Identity Problem. In particular, the idea that some actions might change those affected.

In all scenarios, we ask the participants how much they agree or disagree with a given argument on a scale of 0 to 10, where 0 is *not at all* and 10 is *totally*. In all cases, participants who reject the No-Difference view would express a higher level of agreement with each of the arguments presented. The exact descriptions of these situations can be found in the online repository. The cases are as follows.

The case of depletion: We ask participants how much they agree with the claim that we do not harm future generations by depleting resources because doing otherwise would have resulted in different people being born. This is a classic example of a non-identity situation and it is the case we use in our main analysis.

The case of school reforms: We present a situation in which a reform improves aggregate outcomes but alters who receives the benefits. Participants are asked to evaluate the policy, even though the beneficiaries are not the same individuals as under the status quo. Although this is not strictly an identity-affecting scenario-the new beneficiaries already exist-the case captures the intuition that changes in identity might not matter.

Historical discrimination: This case explores retrospective compensation when descendants of historically discriminated groups would not exist had the original injustices not occurred.

Animal consumption: This case investigates whether reducing the reproduction of animals raised for food harms these animals, since they would not exist otherwise because these animals are raised exclusively to be consumed. This case, though not exactly an identity-affecting situation, relates to the Non-Identity Problem by using the same intuition that we may not affect someone by making him exist when the alternative is not existing at all.

After responding to these cases, participants answered questions about their sociodemographic background. Finally, they were presented with a summary of the non-identity problem and asked how much, after learning about it, they agreed with the conclusion that we can cause harm to future generations on a scale of 0 to 10, where 0 is *not at all* and 10 is *totally*. In this final question, participants who reject the No-Difference view are expected to disagree with the statement.

IV. II. HYPOTHESES

We preregistered the following hypotheses and methods:[4]

H1: Participants will share less money in the situation where each amount is linked to a different person (Identity-Affecting DG), compared to the Traditional Dictator Game.

H2: Participants will share similar amounts in the Random-Recipient DG and the Identity-Affecting DG.

H3: Participants' behaviour in the Dictator Game will correlate with their responses to survey questions about real-world identity-affecting scenarios.

Our first hypothesis predicted similar results to those of Kopec and Bruner (2022) when comparing the traditional version of the Dictator Game with the identity-affecting version. However, our second hypothesis posits that this difference is fully explained by factors other than the identity-affecting nature of the decision. In other words, if behaviour in the Random-Recipient DG mirrors that in the Identity-Affecting DG, this would support the view that individuals hold No-Difference intuitions. Finally, our third hypothesis is motivated by evidence from previous surveys suggesting that people believe it is possible to harm future

[4] The pre-registration can be accessed here (Anonymous link for peer review) and we discuss deviations from the pre-registration in Appendix C.

individuals regardless of identity considerations. We thus predicted similar patterns of moral judgment in our vignettes and in the experimental task.[5]

IV. III. SAMPLE

We conducted our experiment on Prolific, an online platform designed to recruit research participants for academic and commercial studies. Previous evidence shows that Prolific participants are more attentive, engaged, honest, and diverse than those on comparable platforms such as MTurk, CloudResearch, and Dinata (Douglas et al. 2023; Albert and Smilek 2023; Peer et al. 2022). We recruited 841 participants for our experiment in late August 2022. The sample was restricted to individuals residing in the United States with an approval rate of 95% or higher in previous tasks. After excluding participants who did not complete the experiment (12) or failed our attention checks (6), we retained a final sample of 823 participants. On average, participants spent six minutes on the survey and received a payment of $2.27, which included a fixed participation fee and a bonus based on their choices in the experiment.

In addition, we contacted 549 individuals to act as recipients of the amounts shared by our participants—the dictators. These subjects were drawn from a pool of 3,235 participants in a previous study who had been informed they might receive a *surprise bonus* without further explanation. To avoid deception in our Identity-Affecting DG and Random-Recipient DG conditions, we matched each participant in our study with either one or eleven individuals from that pool, depending on which decision was randomly selected for payment. After participants made their sharing decisions, we sent the corresponding payment to the matched individual without informing the others of either the decision or the existence of the new experiment. Therefore, as our instructions indicated, only recipients who received a positive amount ever became aware of the experiment and the dictator's decision, effectively creating an experimental identity-affecting situation.

Table 1 presents the main characteristics of our sample. By design, we recruited an equal number of men and women. Their age ranges from 18 to 93 years, with the average between 37 and 38 years. Most subjects reside in urban areas (75%) and have at most a high school degree (59%). We

[5] In the study, we focus only on the Depletion vignette, although we present correlations and analysis in the Appendix. The reason is that only this case corresponds uncontroversially to a non-identity situation.

	MEAN				T-TEST DIFFERENCE		
	TDG (1)	RR DG (2)	IA DG (3)	Total (4)	(1)–(2)	(1)–(3)	(2)–(3)
Female	0.50	0.50	0.49	0.50	0.00	0.01	0.01
	(0.03)	(0.03)	(0.03)	(0.02)			
Age	36.81	38.36	37.93	37.70	-1.55	-1.13	0.42
	(0.82)	(0.80)	(0.84)	(0.47)			
High education	0.51	0.66	0.59	0.59	-0.15***	-0.08*	0.06
	(0.03)	(0.03)	(0.03)	(0.02)			
Urban	0.74	0.76	0.76	0.75	-0.03	-0.02	0.01
	(0.03)	(0.03)	(0.03)	(0.02)			
Ethnicity	0.72	0.68	0.72	0.71	0.04	0.00	-0.04
White	(0.03)	(0.03)	(0.03)	(0.02)			
Black	0.09	0.10	0.09	0.09	-0.01	0.00	0.01
	(0.02)	(0.02)	(0.02)	(0.01)			
Asian	0.07	0.10	0.06	0.08	-0.03	0.01	0.04
	(0.02)	(0.02)	(0.01)	(0.01)			
Other	0.09	0.09	0.12	0.10	0.00	-0.02	-0.02
	(0.02)	(0.02)	(0.02)	(0.01)			
Political conservative	0.47	0.50	0.50	0.49	-0.02	-0.02	0.00
	(0.03)	(0.03)	(0.03)	(0.02)			
Environm. concerns (0–10)	5.22	5.01	5.12	5.12	0.22	0.11	-0.11
	(0.14)	(0.14)	(0.14)	(0.08)			
N	274	272	268	814			

Table 1: Sample descriptives and balance

Notes: Information on ethnicity is missing for 16 participants. Political conservative if the answer is above 5 on a scale from 0 (left) to 10 (right).

also collected data on participants' ethnicity, political ideology, and environmental concerns. Compared to previous studies, our sample allows for better heterogeneity analyses due to its size and improves the representativeness of our results with respect to the US general population, as our respondents are gender balanced and tend to be less educated than participants in previous studies.

Table 1 also reports the balance between groups who start with different experimental tasks. We find that the sample overall is balanced in terms of gender, age, ethnicity and political views. However, participants who started the experiment with the traditional dictator game were, on average, significantly less likely to have an above-high-school education. To control for this imbalance, we control for education in our between-subject analysis.

V. RESULTS

V. I. EMPIRICAL STRATEGIES

Following our pre-registration, we test our hypotheses using two types of analysis. First, we leverage the within-subject design to (i) test whether individuals differ in sharing behaviour in the Random-Recipient DG and the Identity-Affecting DG conditions (Hypothesis 2) and (ii) to compare individual behaviour in the experimental tasks with responses to the vignette scenarios (Hypothesis 3). For the latter analysis, we compute the difference between each participant's response in the Random-Recipient DG and the Identity-Affecting DG, and assess its correlation with responses in the Depletion case, which represents a purely identity-affecting scenario. Second, we test our first hypothesis by restricting the analysis to the distributive choice presented first to each participant. This approach is equivalent to a between-subjects design and enables a close replication of Kopec and Bruner (2022).[6]

These analytical approaches support different interpretations. The between-subjects design cleanly isolates the effect of each condition on sharing behaviour. The within-subject design also captures behavioural differences across conditions but considering the influence of participants' prior exposure to other scenarios.[7] Accordingly, within-subject results should be interpreted as reflecting the net effect of each condition, acknowledging that participants may have learned key distinctions across conditions based on prior exposure. Consequently, we interpret discrepancies between the two analyses as evidence of an effect of reflecting on the problem on behaviour.

Consequently, we proceed by first reporting results from the within-subject design, both from comparing decisions and relating them to answers in the Vignette, before turning to the between-subject design.

V. II. MAIN RESULTS

WITHIN SUBJECT ANALYSIS

Figure 1 shows how participants distribute the bonus between themselves and a recipient. Starting with the traditional DG, we observe that roughly

[6] The main differences between our experiment and theirs are: (1) our sample is larger, allowing for heterogeneity analyses; (2) we implemented our experiment on Prolific, whereas they used MTurk; (3) participants in our sample are more representative of the general population, being on average less educated and with an equal gender split.

[7] Following standard practice in within-subject designs, our within-analysis uses a random effects model and includes fixed effects for the order tasks were displayed to participants.

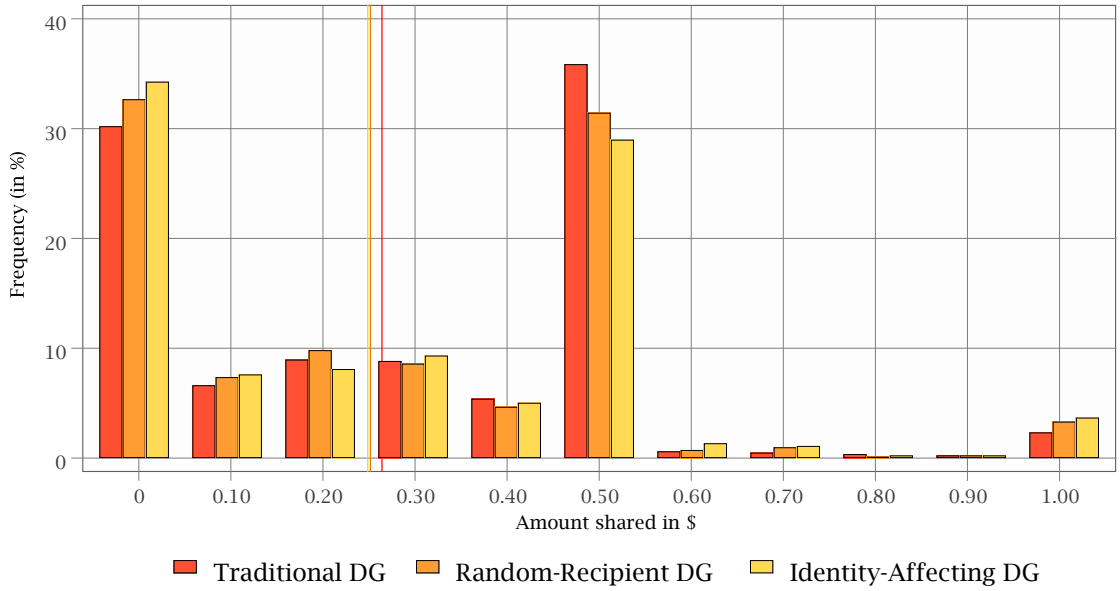

Figure 1: Amount shared in Dictator Game (within-subject)

Notes: The Figure shows the participant's allocation to the recipient across the 3 different tasks (within-subject design). The vertical lines indicate the average amount shared in each of the three tasks.

one out of three participants do not share any of their bonus with the recipient they are matched to. Two-thirds share some amount of the money, with most deciding to split it equally, that is, share $0.50. Turning to the variations of the Dictator Game, we observe a drop in the number of participants who share half of their bonus with the recipient who is either chosen by them (in the Identity-Affecting DG) or at random (in the Random-Recipient DG). Instead, participants seem more likely to share smaller amounts. This is consistent with participants who prefer equality, but take into account the whole pool of potential recipients. In the traditional Dictator Game, an equal split leads to perfect equality. With the larger pool of potential recipients, when perfect equality between the sender and all potential recipients is no longer possible, sharing nothing is a possibility to implement equality between all recipients.

While the overall distributions for the traditional and the Identity-Affecting DG are not statistically different from one another (K-S p-value=0.127), the average amount redistributed is. Table 2 shows that participants share significantly less in the Identity-Affecting DG compared to the traditional DG. Even though the differences are not big, this is in line with our first hypothesis and the findings by Kopec and Bruner (2022). As explained earlier, this analysis includes the effect of being exposed to

all three conditions. Thus, it also accounts for the effect of reflecting on the problem after observing the other conditions. We therefore include

	AVERAGE	ZERO	EQUAL SHARE
	(1)	(2)	(3)
Random-Recipient DG	-0.011**	0.025***	-0.044***
	(0.005)	(0.009)	(0.011)
Identity-Affecting DG	-0.014***	0.041***	-0.069***
	(0.005)	(0.009)	(0.011)
Observations	2442	2442	2442
Mean	0.279	0.279	0.279
RR = IA (*p-value*)	0.607	0.083	0.029

Table 2: Effect of task on amount shared (within-subject)

Notes: The dependent variables are the amount shared in the traditional Dictator Game (column 1), a dummy equals 1 if the amount shared is zero (column 2) and a dummy equals 1 if the participant shares half of the bonus (column 3). Estimated using a random-effects model. Controls include question order, gender, age, age squared, education, urban, environmental attitudes, conservative, and ethnicity. Robust standard errors. Statistical significance is indicated by $*p < 0.1$, $**p < 0.05$, $***p < 0.01$.

fixed effects for the order in which the tasks were displayed to the participant. The within-subject design allows us to go beyond the average sharing shown in Figure 1. We classify respondents based on their responses across the three tasks into the following groups: Those who (i) share the same amount in all three tasks, (ii) share the same amount in the IA DG and the RR DG (but less than in the TDG), and (iii) share less in the IA DG compared to the RR DG. We find that 67.44% of the sample do share the same amount throughout all tasks. This can also be seen in Figure B.4, where we plot the amount shared in each task against each other. Most individuals share either nothing or $0.50 throughout the three tasks. However, 6.63% of respondents share the same amount in the RR DG and the IA DG and less than in the TDG. We interpret this as a sign of inequality concerns among the group of (potential) recipients which is bigger in the RR DG and the IA DG. Finally, 12.16% of respondents share less in the IA
DG compared to the RR DG, suggesting identity-affecting views.

Our second hypothesis concerns the comparison between the Random-Recipient DG and the Identity-Affecting DG, which in our view is the appropriate comparison to identify the effect of an identity-affecting situation. In Table 2, we report the within-subject estimation results. As shown in column 1, average sharing in both the Random-Recipient DG and the Identity-Affecting DG is significantly below the traditional DG. To test whether this difference is due to the identity-affecting character of the decision, we compare the coefficients for the Random-Recipient DG and the Identity-Affecting DG. As indicated by the p-value of 0.604, the coefficients are not statistically different from each other. In line with this, the distributions are not statistically different from each other (K-S p-value=0.99).

If there is a difference between the traditional DG and the Identity-Affecting DG, as indicated by the within-subject estimates, then this difference is likely driven by the larger pool of potential recipients and *not* by the identity-affecting character of the decision. This is because the difference relative to the traditional DG emerges when introducing more recipients (Random-Recipient DG) and not when letting the selection directly affect the identity of the recipient (Identity-Affecting DG). This confirms our second hypothesis of no difference between the Random-Recipient DG and the Identity-Affecting DG. We also show in Table 2 that this is partially driven by an increase in participants who keep the whole bonus to themselves (column 2) and a reduction in participants who share half of their bonus (column 3).

Our results are also in line with the finding in Kopec and Bruner (2022) that both senders and receivers state that receivers are harmed less when receiving lower amounts in the identity-affecting scenario. While they interpret it as confirming their hypothesis of participants holding Person-Affecting views, this is also consistent with recipients who are less likely to associate lower amounts with harm when there are ten other recipients who (also) receive no bonus due to compassion fade. Consequently, we interpret these results as suggestive that individuals' behaviours are consistent with No-Difference views.

In line with our pre-registration, we analyse whether the findings differ between different subsamples. We do not find significantly different results in any sub-sample. However, it is worth noting that the highly educated individuals (secondary school or above) are those who are most likely to differentiate between the Random-Recipient and the Identity-

Affecting DG, which is in line with Kopec and Bruner (2022), who sample from a relatively higher education background.

V. II. I CORRELATION WITH VIGNETTES

A common point of criticism of laboratory experiments is the lack of correspondence to real-world scenarios. In this section, we therefore test how participants' behaviour in the lab corresponds to real-world scenarios that follow a similar identity-affecting intuition. We focus on the case of depletion of natural resources since, among the cases we presented to participants, it is the only one that reflects a purely Identity-Affecting situation, and it is a very common example in the literature. In this case, we explained to participants that if we decide to deplete natural resources, future generations will live with a lower stock of natural resources. However, these people will also be different from those who would live in a world where we would have decided to not deplete natural resources. We ask participants to consider this argument and to state whether depleting natural resources causes harm to future generations.

We compare these answers to behaviour in the incentivized experiment in two ways. First, we link the answers to the difference in sharing between the Random-Recipient DG and the Identity-Affecting DG. Participants who answer differently in these two variations of the Dictator Game are more likely to hold Person-Affecting views and we therefore expect them to agree to a lower degree with the statement that we can cause harm to future generations. In contrast, participants who do not differentiate in these two scenarios might hold No-Difference views and therefore agree more with the statement.

We do not observe that participants who differentiate more between the Random-Recipient and the Identity-Affecting DG in the incentivized experiment agree less that the depletion of resources harms future generations. In fact, contrary to our third preregistered hypothesis, we did not find a significant correlation between the behaviour in the experiment and agreement with any of the other vignettes either. As Table A.3 in the Appendix shows, this is the case both if we consider the difference in behaviour directly or if we consider a binary variable that takes the value of one when individuals shared more in the Random-Recipient DG compared to the Identity-Affecting DG.

These observations do not necessarily show that the experiment design fails to capture participants' intuitions about the Non-Identity Problem because of several reasons. First, the experiment was incentivized,

while the vignettes, by nature, are not. Second, since we do not find big differences in sharing behaviour in the experiment in the first place, the statistical power to detect any correlation with the vignettes was low. Third, the experiment and the depletion vignette, though reflecting identity-affecting situations, relate to different domains: the experiment is about the distribution of resources while the vignette is about environmental damage. So, it is possible, though future research should investigate this in depth, that individuals hold different moral perspectives in different domains, explaining the lack of correlation that we observed.

We therefore move one step back and look at the distributions of answers in the depletion vignette and compare them to the distribution of answers to the question on whether we can harm future generations after learning about the Non-Identity Problem. Notice that these questions are essentially the same; the only difference being that participants answered the depletion vignette before learning about the Non-Identity Problem. So, changes in the distribution of answers would suggest that learning about the problem convinces people about the idea of Person-Affecting views.

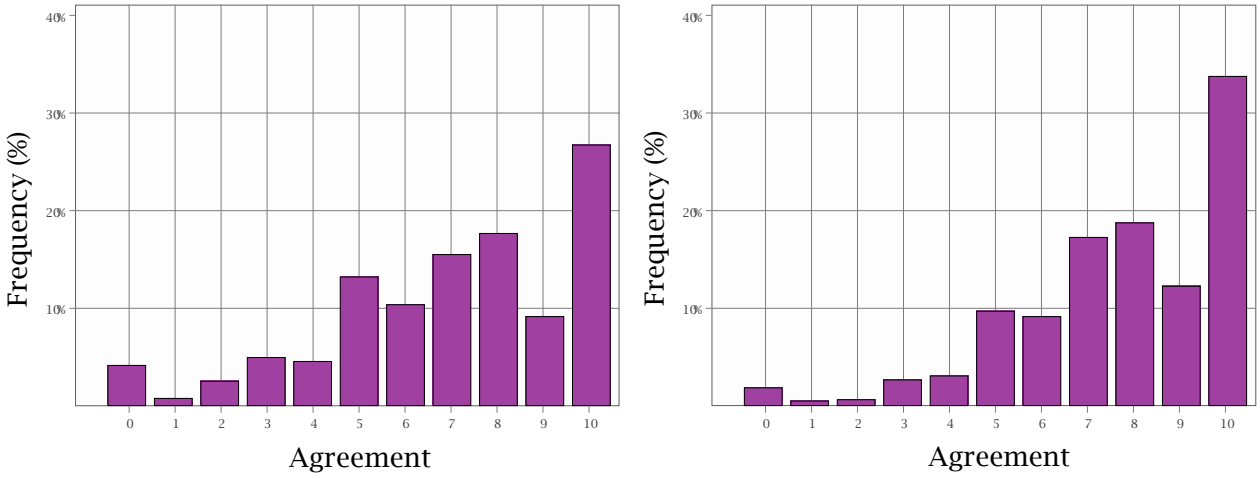

(a) Depletion

(b) Philosophical argument

Figure 2: Vignette

Notes: The figures show the agreement on a scale from 0 (not at all) to 10 (very much) with the statement that (a) we harm future generations by depleting resources and (b) with the philosophical argument on person-affecting views.

Figure 2 shows this comparison. First, we observe that answers to both questions follow a very similar distribution (r = 0.61).[8] Second, we observe that the majority of participants agreed, answered more than 5 on a scale from 0 to 10, that we can harm future individuals both in the depletion vignette and after reading the argument. In fact, they agreed more with the statement than in the vignette (McNemar's test, χ^2=43.23, p-value<0.001, N=814).

Since participants consistently responded that it is possible to harm future individuals, we interpret the vignette results as consistent with the hypothesis that a majority of non-experts hold No-Difference views. This is in line with the findings by Doolabh et al. (2019), who found that participants, including professional philosophers, do not consider the Non-Identity Problem a challenge to their intuition in real-life identity-affecting situations. Moreover, although our observation that participants agreed more often with the idea that we can harm future individuals after direct exposure to the problem should not be interpreted as causal, it suggests that encountering a philosophical challenge to moral intuitions in vignettes does not necessarily alter them. This interpretation aligns with findings from Grodeck and Schoenegger (2023), who challenged participants' intuitions on concrete moral cases, also related to population ethics, with widely accepted philosophical principles. They found that when conflict arose between participants' philosophical beliefs and their intuitions about concrete cases, participants were more likely to revise their endorsement of the former rather than the latter.

V. II. II. BETWEEN-SUBJECT ANALYSIS

Figure 3 presents the distributions of the amount shared by participants in the between-subject analysis, that is, considering only the first decision participants made. Across all conditions we observe that roughly one out of three participants do not share any of their bonus with the recipient they are matched to (31.82%) and we do not observe significant differences across conditions on this. Also, we observe that across all conditions, two-thirds share some amount of the money, with most deciding to split it evenly between themselves and the recipient (33.17%). Yet, we observe a difference in this proportion across conditions. More participants chose to share half of their initial endowment in the Traditional DG condition (38.7%) than in the Identity-Affecting DG condition (30.6%). A two

[8] The agreement with the other three scenarios can be seen in Figure B.5.

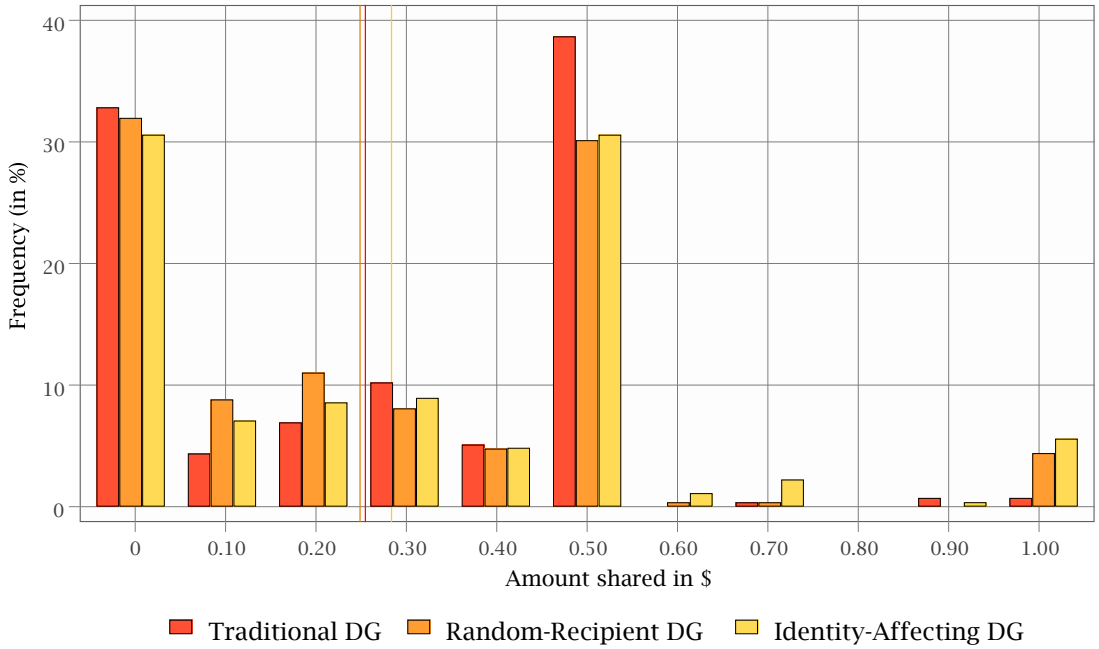

Figure 3: Amount shared in Dictator Game (between subject)

Notes: The Figure shows the participant's allocation to the recipient in the first experimental task (between-subject design), across the 3 different tasks. The vertical lines indicate the average amount shared in each of the three tasks.

sample test of proportions reveals that this difference is statistically significant (diff.=8.1 percentage points, z=1.98, p-value=0.048). However, this difference does not translate to the total amount shared by participants. Overall, participants shared 28.6% of their initial endowment.[9] Moreover, contrary to our first hypothesis, we do not observe a difference in the money participants shared in the Traditional Dictator Game when compared to the Identity-Affecting Dictator Game. Finally, in a Kolmogorov-Smirnov test, we do not observe significant differences in the distributions of shared amounts across these conditions (p-value=0.43).

In Table 3, we take advantage of our larger sample and extend this analysis to different subsamples and use Ordinary Least Squares analysis to identify the effect of the treatment condition on the amount shared by participants. We observe that participants in the Identity-Affecting DG did not share a different amount than those in the Traditional Dictator Game under any subsample. Also, we do not observe a statistically significant difference between sharing behaviour in the Random-Recipient DG and the Identity-Affecting DG.

[9] This amount is similar to that found in a meta-analysis of the Dictator Game. Engel (2011) report that dictator share 28.35% of their endowment across different conditions and considering both field and lab experiments. In a more recent meta-analysis, focusing only on lab experiments as ours, Umer et al. (2022) reported that dictators share 29.6% of their initial endowment.

	BASE	GENDER		EDUCATION		AGE	
		Male	Female	Low	High	Low	High
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Random-Recipient DG	-0.002	-0.016	0.013	-0.014	0.014	0.011	-0.012
	(0.021)	(0.031)	(0.029)	(0.035)	(0.027)	(0.029)	(0.031)
Identity-Affecting DG	0.027	0.029	0.026	0.008	0.045	0.031	0.023
	(0.022)	(0.033)	(0.031)	(0.035)	(0.030)	(0.030)	(0.034)
Observations	814	408	406	336	478	471	343
Mean	0.286	0.273	0.299	0.302	0.274	0.282	0.292
RR = IA (*p-value*)	0.211	0.195	0.681	0.585	0.272	0.537	0.285

Table 3: Dictator game by sub-sample (between-subject)

Notes: The dependent variables measure the amount redistributed. Same controls as in Table 2. Robust standard errors. Statistical significance is indicated by *p < 0.1, **p < 0.05, ***p < 0.01.

We interpret our findings as a failure to replicate the results reported by Kopec and Bruner (2022), who did find a difference between the Traditional Dictator Game and the Identity-Affecting Dictator Game using a design similar to our between-subject analysis. Several factors could explain this discrepancy: our sample was larger and more diverse in terms of gender and education, we used a different platform to distribute the experiment, and there were differences in how instructions were presented. Despite this, our findings are consistent with the idea that their initial results could be explained by mechanisms such as compassion fade, preferences for fairness, or the unidentifiable-victim effect and not by Person-Affecting views. Therefore, we also interpret our results as supporting the hypothesis that non-experts generally hold No-Difference views.

VI. CONCLUSIONS

In this paper, we use a method from Experimental Economics, the dictator game, to explore a key question in Population Ethics: the Non-Identity Problem. A key challenge in examining behaviour in identity-affecting situations is that allowing a decision to influence a person's identity often introduces a larger pool of potential recipients, which may independently

affect behaviour. Building on and extending Kopec and Bruner (2022), we designed an experiment to disentangle the effect of identity-affecting situations from the effect of a broader recipient pool. To achieve this, participants made choices in three variants of the Dictator Game: the traditional DG, the Random-Recipient DG, and the Identity-Affecting DG.

When focusing solely on participants' first decision (between-subject analysis), we found no significant differences in choices across the three variants. However, when analysing all decisions (within-subject analysis), we observed that the average amount shared was slightly lower in both the Random-Recipient and Identity-Affecting DGs. While this difference between the identity-affecting and traditional DG replicates the findings of Kopec and Bruner (2022), our results indicate that the effect is driven by the larger recipient pool rather than the identity-affecting nature of the task. This result is consistent with the hypothesis that participants hold No-Difference views, meaning that they do not consider identity-affecting situations morally different than identity-preserving ones. Moreover, the finding that in the within-subject analysis we observe a difference between the Traditional DG and the other conditions, but we do not see such difference in the between subject analysis suggests that participants' responses evolved as they encountered new conditions, possibly due to learning or heightened moral reflection. Furthermore, in a study using vignettes, participants considered resource depletion to be harmful to future generations, even when it might change their identity. More broadly, the majority agreed that we can both benefit and harm future generations.

Taken together, our findings support Parfit's prediction that No-Difference views are widely accepted as a solution to the Non-Identity Problem, at least among non-philosophers (Parfit 2016). While philosophers continue to debate this issue, the general public seems to regard it as unproblematic. This does not mean that it is of no importance or morally irrelevant, but it suggests that it may have little weight in public discussions on topics such as climate change.

Finally, our study highlights the value of experimental research for understanding philosophical intuitions among non-experts. Although empirical findings cannot directly settle normative debates, knowing which principles people endorse helps philosophers anticipate which issues are likely to resonate in practical applications and democratic discussions. This approach can improve the way philosophy engages with the moral challenges we face today and those that future generations will inherit.

REFERENCES

Albert, Derek A., and Daniel Smilek. 2023. "Comparing Attentional Disengagement between Prolific and MTurk Samples." *Scientific Reports* 13 (1): 20574. https://doi.org/10.1038/s41598-023-46048-5.

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. https://doi.org/10.1111/1468-0262.00302.

Boonin, David. 2008. "How to Solve the Non-Identity Problem." *Public Affairs Quarterly* 22 (2): 129–159.

Boonin, David. 2014. *The Non-Identity Problem and the Ethics of Future People*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199682935.001.0001.

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. https://doi.org/10.1257/aer.97.3.818.

DiSilvestro, Russell. 2009. "Reproductive Autonomy, the Non-Identity Problem, and the Non-Person Problem." *Bioethics* 23 (1): 59–67. https://doi.org/10.1111/j.1467-8519.2008.00681.x.

Doolabh, Keyur, Lucius Caviola, Julian Savulescu, Michael J. Selgelid, and Dominic Wilkinson. 2019. "Is the Non-Identity Problem Relevant to Public Health and Policy? An Online Survey." *BMC Medical Ethics* 20 (1): 46. https://doi.org/10.1186/s12910-019-0379-5.

Douglas, Benjamin D., Patrick J. Ewell, and Markus Brauer. 2023. "Data Quality in Online Human-Subjects Research: Comparisons between MTurk, Prolific, CloudResearch, Qualtrics, and SONA." *PLOS ONE* 18 (3): e0279720. https://doi.org/10.1371/journal.pone.0279720.

Engel, Christoph. 2011. "Dictator Games: A Meta Study." *Experimental Economics* 14 (4): 583–610. https://doi.org/10.1007/s10683-011-9283-7.

Ferrari, Arianna. 2012. "Animal Disenhancement for Animal Welfare: The Apparent Philosophical Conundrums and the Real Exploitation of Animals. A Response to Thompson and Palmer." *NanoEthics* 6 (1): 65–76. https://doi.org/10.1007/s11569-012-0139-1.

Gardner, Molly. 2015. "A Harm Based Solution to the Non-Identity Problem." *Ergo* 2. https://doi.org/10.3998/ergo.12405314.0002.017.

Greaves, Hilary, and William MacAskill. 2021. *The Case for Strong Longtermism*. GPI Working Paper No. 5-2021. Global Priorities Institute, University of Oxford.

Grodeck, Ben, Toby Handfield, and Matthew Kopec. 2022. "The Effect of Compassion Fade on Altruistic Behavior: Experimental Evidence for a Guilt Mitigation Account." SSRN Working Paper. https://doi.org/10.2139/ssrn.4301600.

Hall, Georgina. 2025. "Why Nonidentity Is Not a Problem: Parfitian Defence of Clinicians Refusing to Provide Assisted Reproductive Technologies." *The Journal of Medicine and Philosophy: A Forum for Bioethics and Philosophy of Medicine* 50 (3): 168–180. https://doi.org/10.1093/jmp/jhaf010.

Harman, Elizabeth. 2004. "Can We Harm and Benefit in Creating?" *Philosophical Perspectives* 18: 89–113. https://doi.org/10.1111/j.1520-8583.2004.00022.x.

Ipsos. 2025. *People & Climate Change 2025: Public Attitudes to the Climate Crisis and the Transition to Net Zero*. Technical report. Ipsos.

Jenni, Karen, and George Loewenstein. 1997. "Explaining the Identifiable Victim Effect." *Journal of Risk and Uncertainty* 14 (3): 235–257.

Kopec, Matthew, and Justin Bruner. 2022. "No Harm Done? An Experimental Approach to the Nonidentity Problem." *Journal of the American Philosophical Association* 8 (1): 169–189. https://doi.org/10.1017/apa.2021.1.

Law, Kyle Fiore, Stylianos Syropoulos, and Liane Young. 2023. "Why Do Longtermists Care about Protecting the Environment? An Investigation on the Underlying Mechanisms of Pro-Climate Policy Support." *Sustainability* 15 (24): 16732. https://doi.org/10.3390/su152416732.

Lee, Seyoung, and Thomas Hugh Feeley. 2016. "The Identifiable Victim Effect: A Meta-Analytic Review." *Social Influence* 11 (3): 199–215. https://doi.org/10.1080/15534510.2016.1216891.

Lippert-Rasmussen, Kasper. 2020. *Making Sense of Affirmative Action*. Oxford University Press.

Magnusson, Erik. 2019. "Children's Rights and the Non-Identity Problem." *Canadian Journal of Philosophy* 49 (5): 580–605. https://doi.org/10.1080/00455091.2018.1463798.

McMahan, Jeff. 2013. "Causing People to Exist and Saving People's Lives." *The Journal of Ethics* 17 (1): 5–35. https://doi.org/10.1007/s10892-012-9139-1.

O'Brien, Gary David. 2022. "Beneficence, Non-Identity, and Responsibility: How Identity-Affecting Interventions in Nature Can Generate Secondary Moral Duties." *Philosophia* 50 (3): 887–898. https://doi.org/10.1007/s11406-021-00419-2.

Ogunbode, Charles A., Rouven Doran, Arin H. Ayanian, Joonha Park, Akira Utsugi, Karlijn L. van den Broek, Jihane Ghorayeb, et al. 2024. "Climate Justice Beliefs Related to Climate Action and Policy Support around the World." *Nature Climate Change* 14 (11): 1144–1150. https://doi.org/10.1038/s41558-024-02168-y.

Palmer, Clare. 2011. "Animal Disenhancement and the Non-Identity Problem: A Response to Thompson." *NanoEthics* 5 (1): 43–48. https://doi.org/10.1007/s11569-011-0115-1.

Parfit, Derek. 2016. "The Non-Identity Problem." Talk presented at EAGxOxford 2016, Examination Schools, University of Oxford, June 2016. Transcript. Effective Altruism Forum. https://forum.effectivealtruism.org/posts/YpjXuZkwzSnJxWE5h/transcript-of-a-talk-on-the-non-identity-problem-by-derek.

Parfit, Derek. 2017. "Future People, the Non-Identity Problem, and Person-Affecting Principles." *Philosophy & Public Affairs* 45 (2): 118–157. https://doi.org/10.1111/papa.12088.

Peer, Eyal, David Rothschild, Andrew Gordon, Zak Evernden, and Ekaterina Damer. 2022. "Data Quality of Platforms and Panels for Online Behavioral Research." *Behavior Research Methods* 54 (4): 1643–1662. https://doi.org/10.3758/s13428-021-01694-3.

Pivato, Marcus. 2023. "Population Ethics in an Infinite Universe." *Philosophical Studies* 180 (12): 3383–3414. https://doi.org/10.1007/s11098-023-02014-5.

Rivera-López, Eduardo. 2009. "Individual Procreative Responsibility and the Non-Identity Problem." *Pacific Philosophical Quarterly* 90 (3): 336–363.

Roberts, Melinda A. 2024. "The Nonidentity Problem." In *The Stanford Encyclopedia of Philosophy*, Summer 2024 ed., edited by Edward N. Zalta. Metaphysics Research Lab, Stanford University.

Schoenegger, Philipp, and Ben Grodeck. 2023. "Concrete over Abstract: Experimental Evidence of Reflective Equilibrium in Population Ethics." In *Experiments in Moral and Political Philosophy*, 43–61. Routledge.

Steinbock, Bonnie. 2009. "Wrongful Life and Procreative Decisions." In *Harming Future Persons: Ethics, Genetics and the Nonidentity Problem*, 155–178. Springer Netherlands. https://doi.org/10.1007/978-1-4020-5697-0_8.

Syropoulos, Stylianos, Kyle Fiore Law, and Liane Young. 2024. "The Case for Longtermism: Concern for the Far Future as a Catalyst for Pro-Climate Action." *npj Climate Action* 3 (1): 27. https://doi.org/10.1038/s44168-024-00110-8.

Tank, Lukas. 2022. "Climate Change and Non-Identity." *Utilitas* 34 (1): 84–96.

Thunberg, Greta. 2019. "Greta Thunberg's Speech at the UN Climate Action Summit." Speech transcript, NPR, September 23, 2019.

Tomm, Brandon M., Paul Slovic, and Jiaying Zhao. 2019. "The Number of Visible Victims Shapes Visual Attention and Compassion." *Journal of Vision* 19 (10): 105c. https://doi.org/10.1167/19.10.105c.

Umer, Hamza, Takashi Kurosaki, and Ichiro Iwasaki. 2022. "Unearned Endowment and Charity Recipient Lead to Higher Donations: A Meta-Analysis of the Dictator Game Lab Experiments." *Journal of Behavioral and Experimental Economics* 97: 101827. https://doi.org/10.1016/j.socec.2022.101827.

Weinberg, Rivka. 2008. "Identifying and Dissolving the Non-Identity Problem." *Philosophical Studies* 137 (1): 3–18. https://doi.org/10.1007/s11098-007-9168-y.

Pablo Soto-Mota is an Assistant Professor at the Center for Economic Studies at El Colegio de México. His research lies at the intersection of behavioral economics and ethics, with a focus on moral judgment, fairness, and institutional design. He uses experimental and survey methods.
Contact e-mail:<pisoto@colmex.mx>

Arne Nasgowitz is a researcher at the Centre for Applied Research (SNF) at the Norwegian School of Economics. He is a development economist and his research interests include education, labor markets, beliefs and norms. His work emphasizes experimental evidence from low-income contexts, particularly in East Africa.
Contact e-mail:<arne.nasgowitz@snf.no>

APPENDIX
A. TABLES

	AMOUNT SHARED		
	(1)	(2)	(3)
Random-Recipient DG	-0.005	-0.001	-0.002
	(0.021)	(0.022)	(0.021)
Identity-Affecting DG	0.026	0.028	0.027
	(0.022)	(0.022)	(0.022)
Observations	814	814	814
Mean	0.286	0.286	0.286
Base Controls		✓	✓
Add. Controls			✓
RR = IA (*p-value*)	0.178	0.212	0.211

Table A.1: Effect of task on amount shared (between-subject)

Notes: Baseline controls include gender, age, age squared, education and urban. Additional controls are environmental attitudes, conservative and ethnicity. Robust standard errors. Statistical significance is indicated by $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	DEPLETION	HISTORICAL DISCRIM.	EDUCATION	ANIMALS
	(1)	(2)	(3)	(4)
	Panel A: Full sample			
Δ Lab	0.047	-0.013	0.247	0.155
	(0.166)	(0.185)	(0.166)	(0.193)
Observations	2442	2442	2442	2442
	Panel B: > 0			
Δ Lab	-0.084	-0.290	0.134	0.150
	(0.170)	(0.193)	(0.169)	(0.199)
Observations	1644	1644	1644	1644

Table A.2: Vignette studies

Notes: The dependent variables measures the agreement on a scale from 0 to 10 with "we cause harm to future societies by depleting natural resources" (column 1), "we ought to make compensations on the ground of historical discrimination" (column 2), "an educational program increasing the average amount of numbers progressing to university" (column 3), "we would harm animals by stop consuming animal products and hence reduce their reproduction rates" (column 4). Experiment is a dummy equal to one when the respondent gives less in the Identity-Affecting DG compared to the Random-Recipient DG. Panel B restricts the sample to participants who shared a positive amount in the Random-Recipient DG. Same controls as in Table 2. Robust standard errors. Statistical significance is indicated by $*p < 0.1$, $**p < 0.05$, $***p < 0.01$.

	Base	Gender		Education		Age	
		MALE	FEMALE	LOW	HIGH	LOW	HIGH
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Random-Recipient DG	-0.044***	-0.010	-0.013*	-0.016*	-0.008	-0.006	-0.018**
	(0.011)	(0.007)	(0.007)	(0.008)	(0.006)	(0.007)	(0.007)
Identity-Affecting DG	-0.069***	-0.019**	-0.009	-0.008	-0.018***	-0.007	-0.023***
	(0.011)	(0.007)	(0.007)	(0.008)	(0.006)	(0.007)	(0.007)
Observations	2442	1224	1218	1008	1434	1413	1029
Mean	0.279	0.263	0.296	0.296	0.268	0.275	0.286
RR = IA (*p-value*)	0.029	0.248	0.609	0.331	0.108	0.927	0.461

Table A.3: Dictator Game by sub-sample (within-subject)

Notes: The dependent variables measure the amount redistributed. Same controls as in Table 2. Robust standard errors. Statistical significance is indicated by $*p < 0.1$, $**p < 0.05$, $***p < 0.01$.

B. Figures

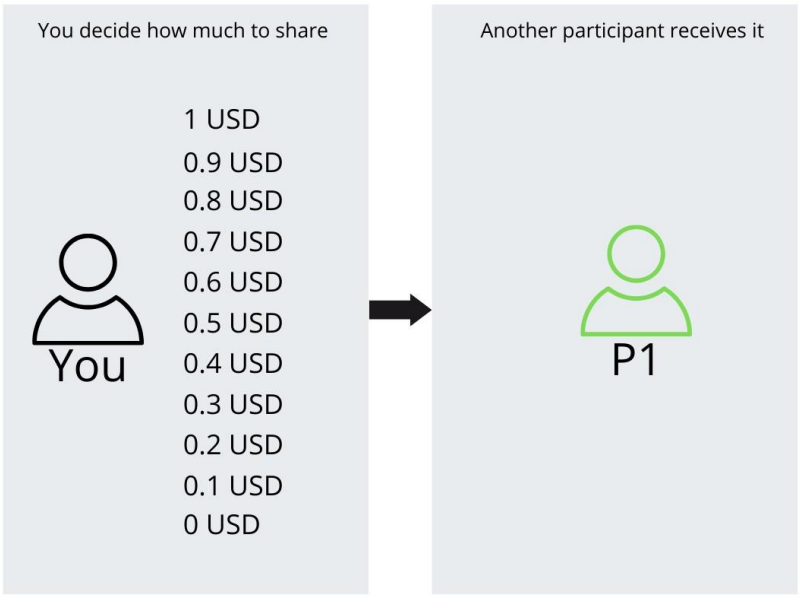

Figure B.1: Traditional Dictator Game

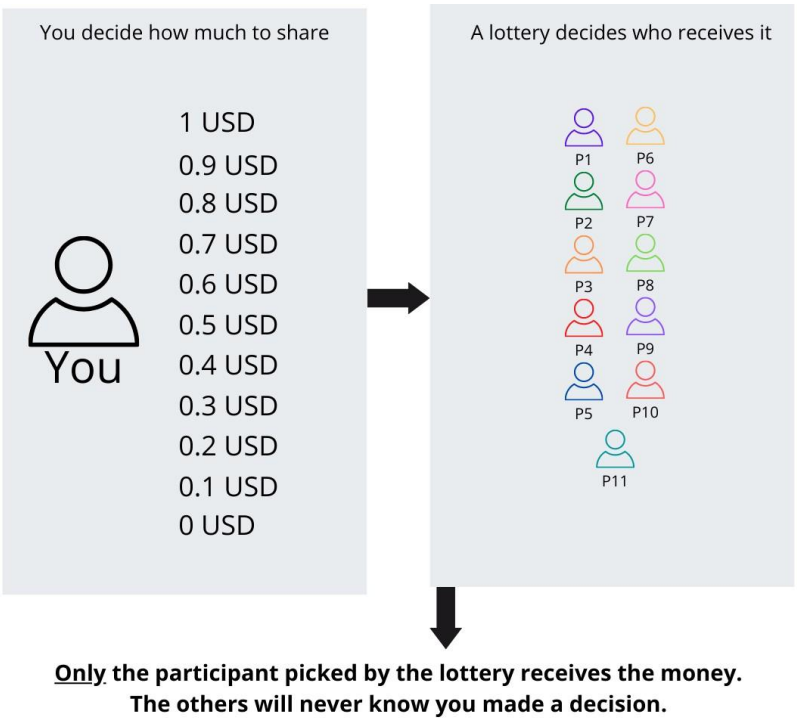

Figure B.2: Random-Recipient Dictator Game

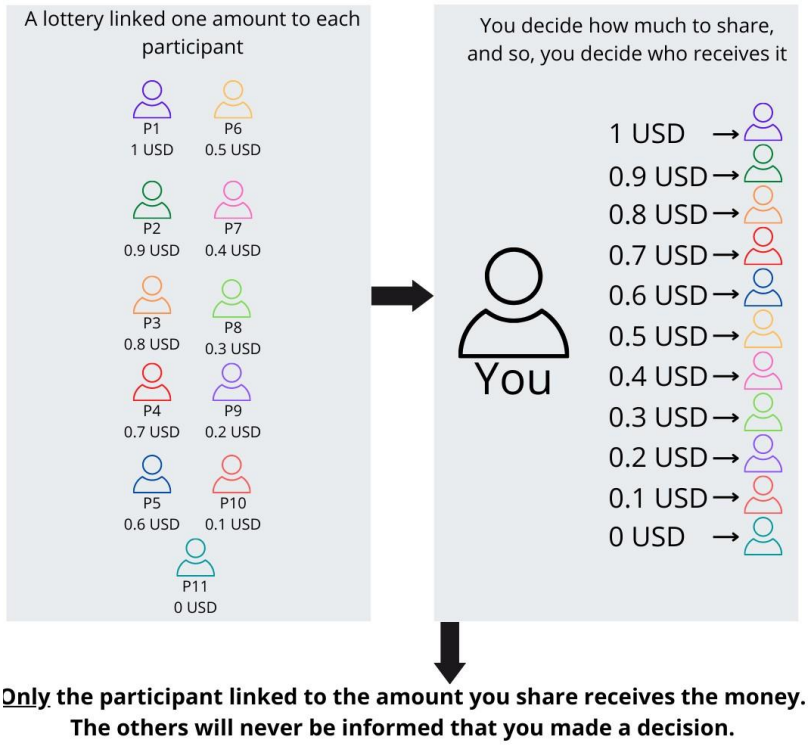

Figure B.3: Identity-Affecting Dictator Game

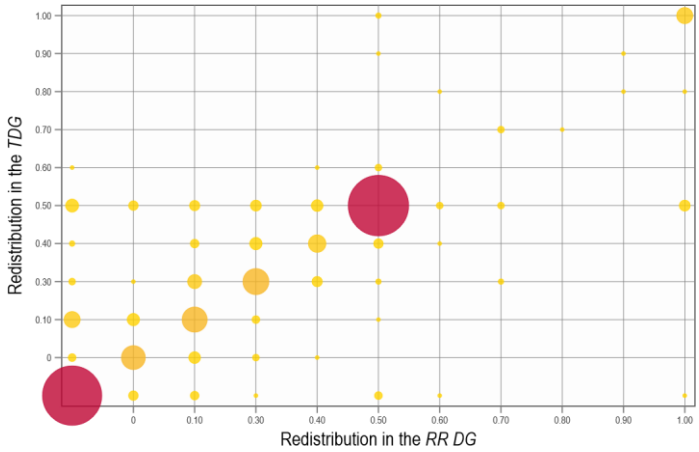

(a) TDG vs RR DG

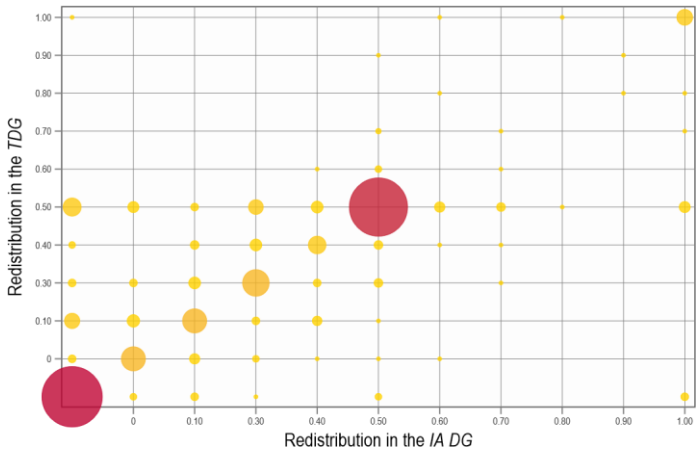

(b) TDG vs IA DG

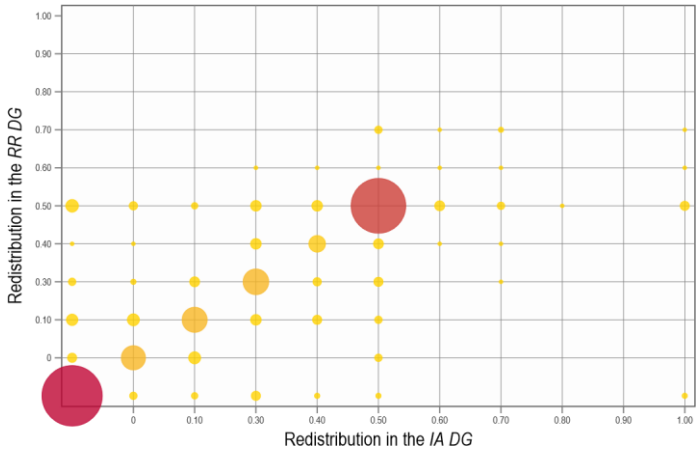

(c) RR DG vs IA DG

Figure B.4: Scatterplot within-subject

Notes: The figures plot the amounts shared in the two versions of the dictator game indicated below the plot.

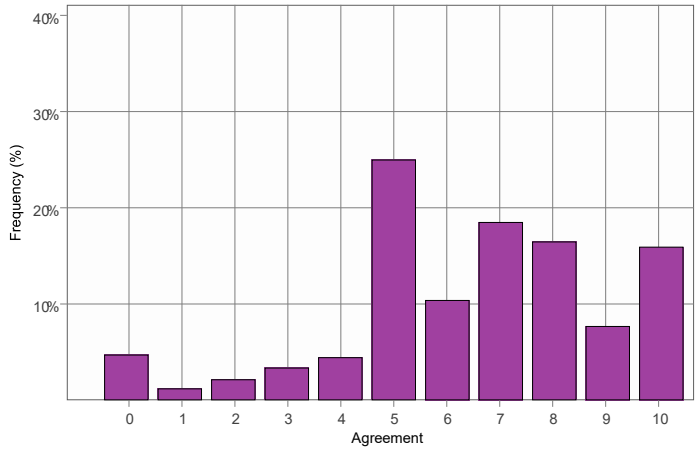

(a) School reform

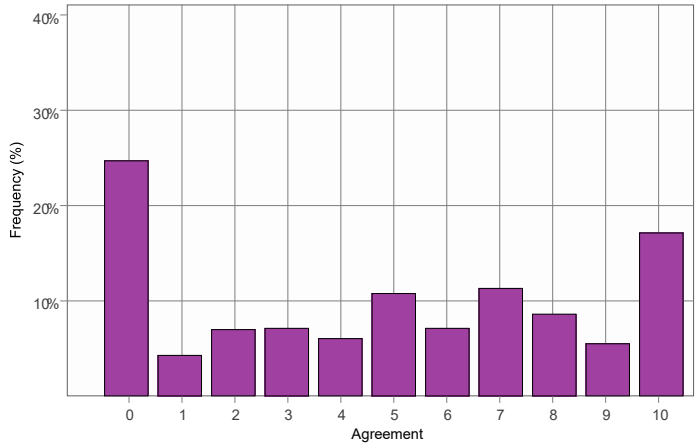

(b) Historical discrimination

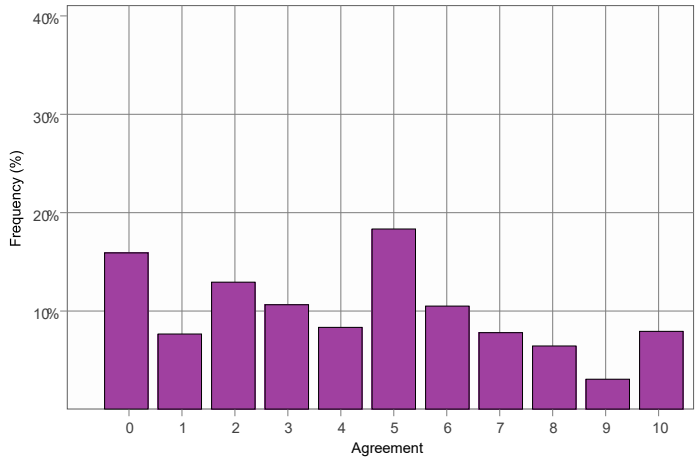

(c) Animal consumption

Figure B.5: Other Vignettes

Notes: The figures show the agreement on a scale from 0 (not at all) to 10 (very much) (a) with a proposed school reform that would make 4 children better and 2 worse off (out of 10), (b) that we ought to compensate historically discriminated groups and (c) with that we would harm animals by reducing their reproduction.

C. Pre-registration

The study was pre-registered with as predicted under the pre-registration number 104751 and can be accessed here.

We follow the pre-registration with one exception. We pre-registered to examine the difference between the traditional DG and the identity-affecting DG using the between-subject design, and the difference between the Random recipient DG and the identity-affecting DG using the within-subject design. In the paper, we use both approaches, within- and between-subject design, throughout the whole analysis.