Environmental and Resource Economics (2024) 87:1045–1060
https://doi.org/10.1007/s10640-024-00854-1

Solar Geoengineering, Free-Driving and Conflict: An Experimental Investigation

Todd L. Cherry[1,2] · **Stephan Kroll**[3] · **David M. McEvoy**[4] · **David Campoverde**[1]

Accepted: 13 February 2024 / Published online: 27 March 2024

Abstract

As the international community continues to fall short on reducing emissions to avoid disastrous impacts of climate change, some scientists have called for more research into solar geoengineering (SGE) as a potential temporary fix. Others, however, have adamantly rejected the notion of considering SGE in climate policy discussions. One prominent concern with considering SGE technologies to help manage climate change is the so-called "free driver" conjecture. The prediction is that among countries with different preferences for the level of SGE, the country that prefers the most will deploy levels higher than the global optimum. This paper tests the free-driver hypothesis experimentally under different conditions and institutions. We find that aggregate deployment of SGE is inefficiently high in all settings, but slightly less so when players are heterogeneous in endowments or when aggregate deployment is determined by a best-shot technology. Despite persistent inefficiencies in SGE deployment, free-driver behavior, on average, is less extreme than the theoretical predictions.

Keywords Climate change · Solar geoengineering · Free driver · Experimental economics

1 Introduction

As global greenhouse gas (GHG) emissions continue to increase, many scientists are concerned that international efforts focused on mitigation will not be sufficient to avoid disastrous climate change. The possibility of missing emissions targets has triggered increased interest in technological interventions—broadly referred to as "geoengineering"—that may help avoid negative consequences from excessive emissions. Chief among these interventions is solar geoengineering (SGE), the method of cooling the planet by reflecting solar

✉ Todd L. Cherry
tcherry@uwyo.edu

1 Department of Economics, University of Wyoming, Laramie, WY, USA

2 CICERO Center for International Climate Research, Oslo, Norway

3 Department of Agricultural and Resource Economics, Colorado State University, Fort Collins, CO, USA

4 Department of Economics, Appalachian State University, Boone, NC, USA

radiation away from Earth. A recent National Academy of Sciences report called for increased research into SGE (NASEM 2021), while other studies have pointed out that there are many concerns with researching and deploying SGE technologies (Biermann et al. 2022), among them international governance.

SGE technologies are relatively cheap and could potentially be deployed by a single country. Weitzman (2015) hypothesized that the country with the highest preferred level of SGE would effectively determine the level of SGE, and in turn the climate, for all countries. He referred to this country as the "free driver," where a single actor, free from institutional or financial constraints, can alone drive the global climate by using cheap geoengineering technologies.[1] The fundamental issue is that countries likely have different preferences regarding the optimal geoengineered climate, and therefore SGE may be a good or a bad (referred to as a GoB) depending on the level of deployment (Ricke et al. 2013; Weitzman 2015; Wagner 2021). The prediction is that the free diver will deploy SGE levels that are inefficiently high relative to a social optimum. Given the important implications of this governance challenge, our study provides an empirical test of the free-driver hypothesis under different conditions using a set of controlled economic experiments.

This study adds to a new strand of literature that experimentally tests the strategic implications of the emerging climate engineering technologies (e.g., Cherry et al. 2022), and more specifically extends the initial experimental test of the free-driver hypothesis conducted by Abatayo et al. (2020). In that study, SGE leads to suboptimal outcomes for three reasons—free-driving behavior, lack of coordination and investment in counter-geoengineering. We extend Abatayo et al. (2020) along two important dimensions: First, while Abatayo et al. (2020) only examine countries with homogeneous endowments, we consider the more plausible scenario of countries having heterogeneous endowments from which they can choose SGE deployment levels. Second, in Abatayo et al. (2020) total SGE deployment by countries is determined by summation of individual SGE efforts. However, given an individual country can achieve its preferred SGE level unilaterally, some studies model SGE deployment as best-shot aggregation (e.g., Barrett 2007; Cherry et al. 2022). We consider both aggregation technologies. Combined our paper makes three contributions to the limited empirical evidence on the free-driver hypothesis: (i) we replicate Abatayo et al. findings under different conditions as a robustness check, (ii) we introduce heterogeneity to a public good-or-bad (GoB) game that characterizes SGE deployment, and (c) we consider the impact of characterizing overall SGE deployment level as determined by the best-shot aggregation technology instead of the summation technology.[2]

The extensions to Abatayo et al. (2020) test the free-driver hypothesis under different conditions and institutions that relate to the complexity of deploying SGE in the real world. In particular, we investigate the impact of heterogeneity in wealth among actors, which corresponds to countries having different endowments and thus different technological abilities to invest in SGE. Abatayo et al. (2020) reference the importance of this design choice (*p*. 13,397): "inequality in decision makers' endowments, which is a proxy for variable welfare of countries, might exacerbate free-driving, given that the regions that will suffer the most from global warming are also the poorest." The behavioral effects of endowment heterogeneity on the provision of public goods have been documented in previous studies

[1] The introduction and description of the term "free driver" initially appeared in a policy paper by Wagner and Weitzman (2012) and was later formalized in a theoretical model by Weitzman (2015).

[2] In a follow-up paper to Abatayo et al., Ghidoni et al. (2023) test in a similar set-up whether side-payments can decrease the extent of free driving behavior.

(e.g., Chan et al. 1999; Cherry et al. 2005), but the impact of endowment heterogeneity related to SGE (or GoB environments more generally) has not yet been researched. For that reason, we are less interested in why endowment heterogeneity might make a difference for behavior (because of, for example, different choice sets or inequality aversion), just in whether it makes a difference at all.

Our experiments also consider different aggregation technologies for determining the level of SGE the group ultimately experiences. While some studies have modeled SGE using a summation technology (e.g., Abatayo et al. 2020; Weitzman 2015; Moreno-Cruz 2015), others have characterized aggregation as "best shot"—the group level of SGE is equal to the highest level deployed (e.g., Barrett 2007; Cherry et al. 2022).[3] The motivation for using a model with best-shot aggregation technology rests on SGE being a low-cost public good-or-bad (GoB) that enables a single country, the free driver, to unilaterally achieve its preferred SGE level. The conjecture is the country with the highest ideal SGE level responds to the actions of others to achieve its desired level of SGE, either by deploying alone or adding to others' deployment. There is little strategic uncertainty that might result in overshooting SGE deployment. Abatayo et al. (2020) point out explicitly that this strategic uncertainty and "the resulting variability in outcomes ends up being more detrimental to the economy in terms of economic surplus than the losses that would be incurred from effort escalation." (*p*. 13,397). To reduce strategic ambiguity in the lab, we consider the best-shot aggregation technology. While the free-driver prediction is fundamentally the same under both summation and best-shot aggregation technologies—in equilibrium the player with highest preferred level is the sole free driver—it is an interesting empirical question whether the different institutions impact actual behavior. Others have demonstrated the importance of different aggregation technologies in a summation and best-shot public goods environment (Cherry et al. 2005; Kroll et al. 2007), but not in a public good-or-bad (GoB) environment like SGE. By considering the behavioral implications of endowment heterogeneity and aggregation technologies, results will provide insights for the direction of future experimental studies in this emerging literature.

Our experimental design varies whether endowments are homogeneous or heterogeneous, the size of the endowments, and whether group-level SGE is determined by summation or best-shot technologies. Like Abatayo et al. in a subset of their treatments we also consider the effects of allowing counter-SGE. In the treatments with homogeneous endowments and summation technologies, our results are broadly consistent with the overall findings in Abatayo et al. (2020): we see group-level SGE deployment close to levels predicted by the free-driver hypothesis, and the availability of counter-SGE brings groups closer to the socially optimal level but at very high costs due to wasted resources. A more detailed individual-level analysis, however, reveals that, while the group-level data matches predictions quite well, individual behavior deviates from the theory—the predicted free driver rarely chooses their monetarily ideal level of SGE while others rarely choose zero. Rather, on average, the predicted free driver chooses a lower amount of SGE and is able to save on the cost of deployment, while the other players invest smaller but positive amounts.

We find that introducing heterogeneity in endowments reduces the average group deployment of SGE, moving slightly toward the social optimum. The best-shot technology causes a decrease in average net SGE deployment relative to summation, but significant resources are wasted in a zero-sum conflict as players other than the "free driver" invest

[3] Barrett (2007, p.38) points out that "geoengineering essentially constitutes a large project, a single best effort".

in SGE. Although introducing counter-SGE causes significant inefficiencies, overall we observe less conflict (and wasted resources) than predicted.

2 Geoengineering Game

The experimental design is a variant of the geoengineering game developed by Abatayo et al. (2020). Participants are randomly assigned to groups of three with each group member A, B or C having different single-peaked preferences about the level of SGE. Participants are aware of their own and others' ideal levels of SGE. Each group member simultaneously decides their contribution to SGE, and these are aggregated into a group-level SGE.[4] Members incur increasing costs as the actual level of SGE deviates from their ideal level. As is common in economic lab experiments, all decisions are financially consequential—subjects earn tokens, which at the end of the experiment are converted to a cash payment.

Specifically, group member i has endowments of $E_i = 100, 150$ or 200 tokens (i = player A, B or C). They choose their "production levels" g_i, whereby each unit of production costs four tokens, $c = 4$. The sum of production levels determines the group level of SGE. Individual payoffs are determined by the difference between the group level of SGE, G, and individually optimal levels of g_i^*:

$$\pi_i = E_i - a * \left|G - g_i^*\right| - c * |g_i|$$

where the "penalty" a is a group member's cost for that difference. As in Abatayo et al. we set the penalty cost equal to $a = 10$ for each token difference, independent of whether the group level is too high or too low compared to the individually optimal level. Individually optimal levels for players A, B and C are always $g_A^* = 2$, $g_B^* = 6$ and $g_C^* = 10$. Following Abatayo et al. (2020), we allow subjects to produce up to 10% of their endowment: $g_i \in \left[0, \frac{1}{10}E_i\right]$ and $g_i \in \left[-\frac{1}{10}E_i, \frac{1}{10}E_i\right]$ for the geoengineering-only and counter-geoengineering treatments, respectively.

The experimental design varies four treatment variables: (i) the level of endowment, (ii) the equality of endowments, (iii) the aggregation technology of SGE, and (iv) the ability to counter-engineer. For (i) and (ii), the level of endowment can take one of three values (100, 150 or 200), and participants having either homogenous endowments ($E_1 = E_2 = E_3 = 100$ or $= 150$ or $= 200$) or heterogeneous endowments ($E_i = 200, E_j = 150, E_k = 100$, with $i,j,j \in \{A, B, C\}$). For (iii), the group SGE level is either defined by the sum of individual contributions (i.e., summation: $G = g_1 + g_2 + g_3$) or by the highest individual contribution among the members (i.e., best-shot: $G = max\{g_1, g_2, g_3\}$). And for (iv), participants either do not have the ability to offset others' contributions to SGE ($g_i \in \left[0, \frac{1}{10}E_i\right]$) or they do have the ability to counter others' SGE ($g_i \in \left[-\frac{1}{10}E_i, \frac{1}{10}E_i\right]$).

Table 1 summarizes the experimental design. Three no-counter-geoengineering treatments vary endowment levels, endowment equity, and aggregation technology. Two counter-geoengineering treatments vary endowment inequity. Two items are worth highlighting. First, within the homogeneous endowment treatment, the level of endowment varies

[4] As is standard practice in economic experiments, the instructions use neutral language and avoid terms like solar geoengineering, climate change and public goods.

Table 1 Summary of treatments

	Counter-geoengineering	Endowment distribution	Endowment level	Aggregation technology
Treatment 1 (n = 840)	No	Homogeneous	100, 150 or 200	Summation
Treatment 2 (n = 600)	No	Heterogeneous	100, 150 and 200	Summation
Treatment 3 (n = 780)	No	Homogeneous	100, 150 or 200	Best-shot
Treatment 4 (n = 900)	Yes	Homogeneous	100, 150 or 200	Summation
Treatment 5 (n = 720)	Yes	Heterogeneous	100, 150 and 200	Summation

n indicates the number of individual production decisions

across three levels ($E_1 = E_2 = E_3 = 100$ or $= 150$ or $= 200$). And second, given the nature of the best-shot feature, the comparison of aggregation technologies is limited to the case of homogeneous endowments and no counter-geoengineering.

One hundred ninety-two subjects, recruited from the student body at the University of Wyoming through ORSEE (Greiner 2015), participated in one of 10 sessions (two sessions for each treatment) in the spring of 2022. The experiment was programmed and conducted using oTree (Chen et al. 2016). Each session consisted of 20 repeated but independent decision periods. We observe 3840 individual-level decisions (1280 unique groups). To minimize dependence across periods, we use a stranger design to control for reputation effects—i.e., groups are reshuffled with different participants to form a new group each period, with subjects not knowing the identity of the other two group members. We also avoid income effects and risk smoothing across periods by randomly drawing one period to determine the payoffs for participants. Note that participants are assigned a new player type (A, B or C) each period, so they likely have different monetarily ideal SGE levels and different endowment levels from one period to the next.

3 Social Optimum and Equilibrium Predictions

Table 2 summarizes treatment-specific equilibrium predictions for individual and group production in SGE. Treatments vary with the level and equity of endowments, the aggregation technology, and the ability to invest in counter-geoengineering, and the game-theoretic predictions lead to two aggregate outcomes—a free-driver result and a geoengineering conflict result. We organize our testable hypotheses around these two predicted outcomes. Similar to Abatayo et al. (2020), we also explain why economic surplus in the equilibria is smaller compared to surplus at the social optimum, with economic surplus measured as sum of individual payoffs within a group.

Free-driver hypothesis. We start with Treatment 1 (the baseline), which features homogeneous endowments, summation aggregation, and no counter SGE. In this treatment, theory predicts the free-driver result: the group experiences excessive SGE because the group member that prefers the highest SGE level (Player C) will impose their ideal level of SGE on the group. Specifically, theory predicts the Player C will act as free driver and produce their monetarily preferred level of SGE ($g_C = 10$), and other group members, anticipating the free-driver's actions, produce nothing ($g_B = g_A = 0$). This yields excessive SGE for the group ($G = 10 > G^* = 6$), which leads to excessive cooling and lower surplus (between

Table 2 Social optimum, Nash equilibrium predictions and total surplus

Treatment	Endowment			Predicted SGE				Surplus at Nash Eq	Surplus at optimum
	E_C	E_B	E_A	g_C	g_B	g_A	G		
1	100	100	100	10	0	0	10	170 (86.6%)	196
	150	150	150	10	0	0	10	320 (92.5%)	346
	200	200	200	10	0	0	10	470 (94.8%)	496
2	100	200	150	10	0	0	10	320 (92.5%)	346
	100	150	200	10	0	0	10	320 (92.5%)	346
	150	200	100	10	0	0	10	320 (92.5%)	346
	150	100	200	10	0	0	10	320 (92.5%)	346
	200	150	100	10	0	0	10	320 (92.5%)	346
	200	100	150	10	0	0	10	320 (92.5%)	346
3	100	100	100	10	0	0	10	170 (86.6%)	196
	150	150	150	10	0	0	10	320 (92.5%)	346
	200	200	200	10	0	0	10	470 (94.8%)	496
4	100	100	100	10	6	−10	6	116 (59.2%)	196
	150	150	150	15	6	−15	6	226 (65.3%)	346
	200	200	200	20	6	−20	6	336 (67.7%)	496
5	100	200	150	10	11	−15	6	226 (65.3%)	346
	100	150	200	10	15	−20	5	180 (52.0%)	346
	150	200	100	15	−4	−10	6	254 (73.4%)	346
	150	100	200	15	10	−20	5	180 (52.0%)	346
	200	150	100	20	−4	−10	6	234 (67.6%)	346
	200	100	150	20	1	−15	6	226 (65.3%)	346

"Surplus at Nash Eq" is the predicted economic surplus at the Nash equilibrium and, in parentheses, as a percentage of total economic surplus at the social optimum. "Surplus at optimum" is the total economic surplus at the social optimum

86.6% and 94.8% of surplus at the social optimum, depending on the endowment level). The first three rows in Table 2 summarize the predicted free-driver result in the baseline treatment, which serves also as replication of Abatayo et al. (2020) and provides a basis for the subsequent hypotheses.

While baseline Treatment 1 (and also Treatment 4 with counter-engineering) is a replication of the main treatments in Abatayo et al. (2020), it is worthwhile to point out the key differences in the experimental designs: (i) we did not allow communication between subjects, (ii) decision-makers are, in line with standard procedures in lab experiments, individuals instead of teams of two, (iii) in order to mimic the one-shot nature of GE decision-making and to avoid unwarranted reputational effects, we use a stranger design instead of a partner design, and (iv) we have three decision-makers in one group instead of two or six. Thus, our paper provides a robustness check because these differences could lead to differences in behavior between our subjects and the subjects in Abatayo et al.[5]

[5] Ghidoni et al. (2023), the extension of Abatayo et al. (2020) with side-payments, does also not allow communication between subjects, only within a decision-making unit.

We extend the baseline treatment to consider how the free-driver result is affected by the level of endowments, equity of endowments, and the aggregation technology. Table 2 details the game-theoretic predictions across endowment levels, endowment homogeneity and heterogeneity (Treatments 1 and 2), and summation and best-shot aggregation technologies (Treatments 1 and 3). In all cases in Treatments 1—3, the predicted individual and group SGE levels remain the same ($g_C = 10; g_B = g_A = 0; G = 10 > G^* = 6$).[6]This leads to three additional hypotheses: we expect the free-driver result will be equivalent across endowment levels in the baseline treatment (*endowment level hypothesis*), across homogeneous and heterogeneous endowments treatments (*endowment equity hypothesis),* and across summation and best-shot aggregation treatments (*aggregation technology hypothesis*), with similar effects on economic surplus as in the baseline treatment translate.

Geoengineering conflict hypothesis. Introducing counter-geoengineering changes game-theoretic predictions; instead of a free-driver result, we expect a *geoengineering conflict result*. Generally, in both treatments 4 and 5, Player C maximizes their production of SGE ($g_C = 10, 15$ or 20) even if that maximum level is beyond their monetarily preferred level. This is because Player A, the member with the lowest monetarily preferred level ($g_A^* = 2$), anticipates the excessive SGE production by Player C and, in absolute terms, chooses the maximum level of counter SGE ($g_A = -10, -15$ or -20). Player B, the member with the median monetarily preferred level, which also matches the socially optimal level, chooses a SGE level (g_B) so the group SGE is optimal ($g_B^* = G^* = 6$). Thus, in each equilibrium of the treatments with counter-geoengineering, the group achieves (or nearly achieves) the socially optimal level of SGE, but only after costly conflict that substantially reduces surplus (59.2% to 67.7% of the surplus at the social optimum in Treatment 4; 52.0% to 73.4% in Treatment 5). Table 2 summarizes the predicted *geoengineering conflict result* in Treatments 4 and 5.

We also consider how the geoengineering conflict result is affected by the level and equity of endowments. Again, Table 2 provides the game-theoretic predictions across endowment levels and endowment equity (homogeneous vs. heterogeneous). First, for endowment level, recall that Players A and C maximize production of offsetting counter SGE and SGE. To the extent this holds across homogeneous endowment levels (Treatment 4), larger homogeneous endowments will translate to greater symmetric offsetting production of counter SGE and SGE, while still leading to the socially optimal level of SGE. Thus, in Treatment 4, we expect higher homogeneous endowments will increase the cost of the geoengineering conflict (surplus, however, will still be higher because of the larger endowments).

Things are a bit different with heterogeneous endowments. Players A and C still maximize production of counter SGE and SGE, but heterogeneous endowments lead to *asymmetric* levels of SGE and counter SGE. Player B still responds by choosing a level of SGE (g_B) so the group SGE is at, or close to, their monetarily preferred (and socially optimal) level ($g_B^* = G^* = 6$). Comparing Treatments 4 and 5, we expect to see a similar geoengineering conflict result—a costly (symmetric or asymmetric) investment in competing counter SGE and SGE that reduces surplus even when group SGE is at, or close to, the social optimal level.

[6] Given SGE is relatively inexpensive, we parameterize endowment heterogeneity to focus on income effects rather than capacity effects. However, introducing counter-engineering creates a conflict in which endowment heterogeneity entails capacity effects.

Table 3 Mean group and individual SGE in treatment 1

	Combined	E = 100	E = 150	E = 200	$\Delta_{200-100}$
Group SGE					
Socially optimal SGE	6.00	6.00	6.00	6.00	0.00
Predicted SGE	10.00	10.00	10.00	10.00	0.00
Observed SGE	10.01 (3.77)	9.35 (3.99)	10.11 (3.33)	10.59 (3.94)	1.24 (13.3%)
Individual SGE					
Player A ($g_A^* = 2$)	1.21 (1.99)	1.15 (1.88)	1.14 (1.80)	1.34 (2.32)	0.19 (16.5%)
Player B ($g_B^* = 6$)	2.41 (2.32)	2.18 (2.09)	2.47 (2.27)	2.58 (2.58)	0.40 (18.3%)
Player C ($g_C^* = 10$)	6.40 (2.36)	6.01 (2.56)	6.50 (2.39)	6.67 (2.06)	0.66 (11.0%)

Unconditional means in table; *p*-values in text based on conditional tests; standard deviations in parentheses in columns 2-5; percentage increases from E=100 in parentheses in column 6

4 Results

4.1 Free-Driver Hypotheses

We begin by considering the free-driver hypotheses with a review of the individual and group-level SGE production in the baseline treatment. Table 3 reports the optimal SGE, the predicted and observed group SGE levels, and the observed individual SGE levels in Treatment 1. The first column in Table 3 includes the combined numbers, the next three columns segment by endowment level, and the last column shows the difference between average contributions by subjects with high and low endowments.

The average group production of SGE (10.01) exceeds the socially optimal level ($G^* = 6$) and is statistically equivalent to the predicted level of 10.00 ($p = 0.167$) when controlling for round effects and clustering standard errors.[7] Thus, on average, group SGE outcomes are consistent with the free-driver hypothesis. The individual SGE production choices by player types, however, depart from predicted behavior. Recall, we expect players A and B to produce zero SGE and Player C to produce 10. But we observe in Treatment 1 that Players A and B produce 1.21 and 2.41 on average, while Player C produces 6.40. This suggests that Player C produces less than their monetarily ideal level while Players A and B produce levels above equilibrium predictions. Ultimately, the group invests close to the predicted 10, but Player C does not bear the entire burden and, in a game-theoretic sense, plays on average their best response to the other players' choices (but not vice versa). Thus, results support the group-level free-driver result (excessive SGE and loss of surplus) but with some shifting of behavior at the individual level.

Table 3 also sheds light on the *endowment level hypothesis* by reporting the mean group and individual SGE results from Treatment 1 by endowment level. Results reveal a small but consistent positive relationship between endowment level and SGE production. At the group level, mean SGE levels are 1.24 (13.3%) higher with 200 endowments compared to 100 endowments (10.59 vs. 9.35) but the difference is insignificant when controlling for

[7] We report p-values from conditional t-tests from least squares estimates that take advantage of the panel nature of the data to control for round effects, while also accounting for observational dependence with robust standard errors clustered at the session level.

Table 4 Mean group and individual SGE by treatments w/o counter

	Treatment 1	Treatment 2	Treatment 3
Endowments	Equal	Unequal	Equal
Counter-geoengineering	No Counter	No Counter	No Counter
Aggregation tech	Summation	Summation	Best-shot
Group SGE			
Socially optimal SGE	6.00	6.00	6.00
Predicted SGE	10.00	10.00	10.00
Observed (Net) SGE	10.01 (3.77)	9.15 (2.97)	9.48 (2.01)
Total SGE	10.01 (3.77)	9.15 (2.97)	15.50 (6.04)
Individual SGE			
Player A ($g_A^* = 2$)	1.21 (1.99)	0.47 (1.34)	2.62 (3.61)
Player B ($g_B^* = 6$)	2.41 (2.32)	1.89 (1.67)	3.76 (3.77)
Player C ($g_C^* = 10$)	6.40 (2.36)	6.79 (2.25)	9.11 (2.00)

Unconditional means in table; *p*-values in text based on conditional tests; standard deviations in parentheses

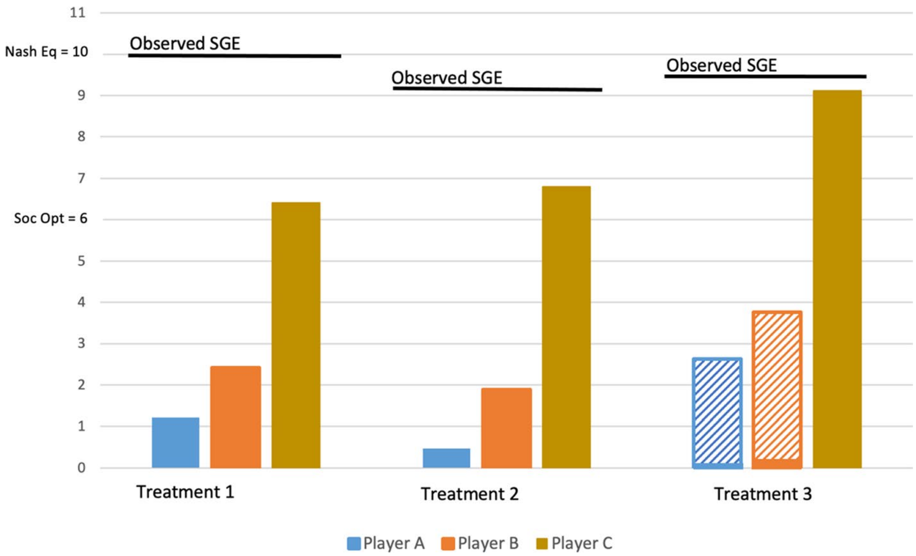

Fig. 1 Contribution to observed SGE by player type and treatment. Note: For players A and B in treatment 3, striped and solid bars represent mean inconsequential and consequential best-shot production

round effects and clustering standard errors ($p = 0.252$). This general relationship emerges across player types, but all player-level differences between endowment levels are statistically insignificant ($p > 0.15$). In summary, our data support the hypothesis that different endowment levels do not impact behavior when groups are homogeneous.

To examine the *endowment equity hypothesis*, we turn to Table 4 and Fig. 1 for a comparison between the homogeneous and heterogeneous endowments treatments (Treatments 1 and 2). The comparison reveals that average group SGE production is lower when endowments are heterogeneous (9.15 vs. 10.01; $p = 0.021$). The level of group SGE, however, remains well above the socially optimal level ($G^* = 6$). A review of individual player-type SGE production reveals the group outcome resulted from competing responses, with the introduction of heterogeneous endowments appearing to lower production by players A and B ($p=0.051$ and $p=0.105$) while Player C production is statistically unchanged ($p=0.551$). Overall, contrary to the endowment equity hypothesis, results suggest that endowment equity matters. Group level production of SGE is reduced by introducing heterogeneity in endowments (via non-free drivers), and this in turn increases efficiency—total surplus for heterogenous groups in Treatment 2 is 82% of possible surplus and therefore higher than in homogeneous groups where total surplus is below 80% for all three endowment levels.

We can also use Table 4 and Fig. 1 to examine the *aggregation technology hypothesis* by comparing the treatments with summation and best-shot aggregation (Treatments 1 and 3). Two results are noteworthy. First, in the baseline treatment with summation aggregation, average group SGE production was 10.01, which is significantly higher than the 9.48 observed in the best-shot treatment ($p=0.025$). In both treatments, group SGE levels are considerably higher than the socially optimal level, which results in a reduction of economic surplus. However, with the best-shot aggregation (Treatment 3), the final average group-level SGE (9.48) disregards the unproductive individual production (6.02) that is less than the "best-shot" level.[8]

Second, the individual SGE levels in Table 4 reveal that the best-shot technology causes an increase in SGE production for all three player types relative to the summation (Treatment 1), but only the difference in Player C's decisions (6.40 vs. 9.11) is significant ($p=0.008$). This finding is consistent with the conjecture that the best-shot aggregation mitigates strategic uncertainty. Figure 2, which shows the best-response functions for all three players, provides insights why one can expect higher individual production levels in the summation treatment, particularly from Player C. In this treatment, C's best response to any positive deployment by Players A and B is to decrease their own deployment away from 10. Thus, if a C player expects (correctly, as it turns out) any "trembling hand" by the other players, they should decrease their own deployment. In contrast, C's best response to positive deployment by A or B in the best-shot treatment is, as long as their deployment is less than 10, fixed at 10 (see top part of Fig. 2)—uncertainty is basically eliminated for Players C, which of course was the motivation for implementing the best-shot treatment. With the best-shot technology, Player C's behavior is more aligned with the free-driver hypothesis of 10 compared to the other treatments.

The same relative incentives exist for Players A and B when responding to positive deployment by others. The bottom portion of Fig. 2 illustrates that the best response for both Players A and B, respectively, is decreasing in the deployment of the two other players until turning to zero at 2 (6) for Player A (B). The implication is that if the players expect positive deployment by others for any reason, the best response is weakly higher SGE for all players in the best shot treatment.[9]

[8] This unproductive production is akin to investments in SGE capacity, costly negotiations, lobbying, etc. Future studies may consider experimental designs that rebate unproductive production.

[9] We thank an anonymous reviewer for pointing out the differences in best responses between the summation and best shot treatments and suggesting a "trembling hand" argument as potential explanation of the observed difference between summation and best shot. Note that in a summation treatment with more "non-C players", strategic uncertainty and therefore the expectations of trembling hands increase even more in the summation treatment and therefore we might expect the free-driver's production to decline even more relative to its average production given a best-shot technology.

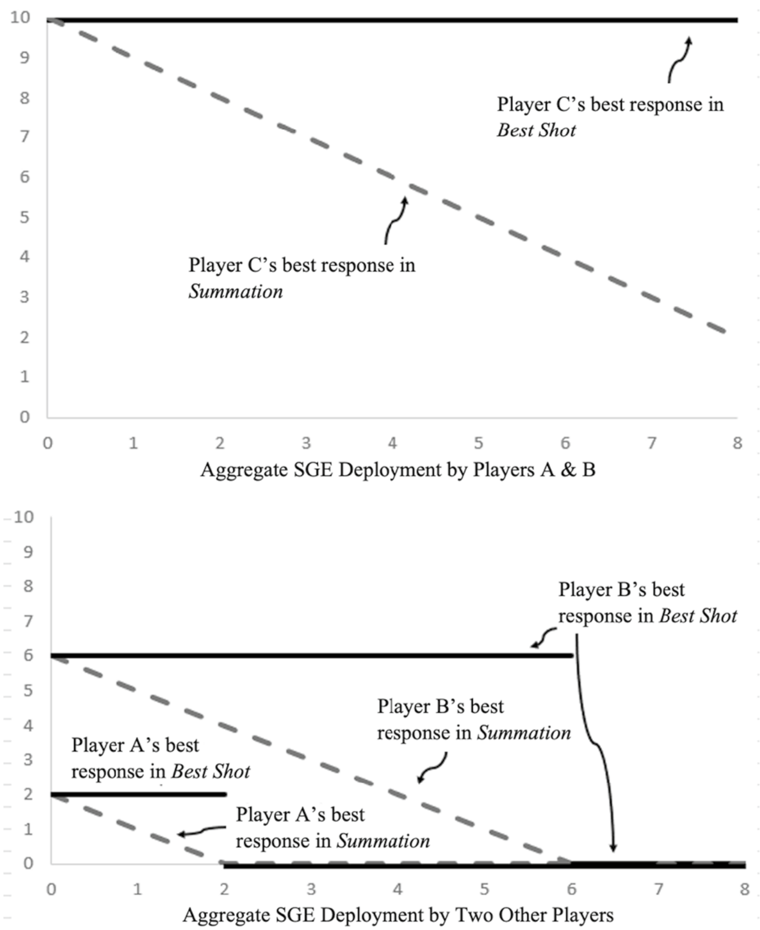

Fig. 2 Best response functions in summation vs. best shot

Overall, the best-shot technology leads to lower final group level SGE relative to summation. This result, however, comes at an efficiency cost: Player C increases production but players A and B invest in unproductive SGE (on average 2.62 and 3.76, respectively), which reduces surplus.[10]

4.2 Geoengineering Conflict Hypotheses

We now turn to the geoengineering conflict hypotheses. Table 5 provides the average group and individual SGE production levels in the two treatments with opportunities for counter-geoengineering.

[10] Note that Player C's average production in Treatment 3, 9.11, is not equal, contrary to what theory would predict, to Net SGE, 9.48, because there were 39 cases, mainly in the first half of the experiment, where Player C's production was actually not the highest one and therefore did not provide the best shot.

Table 5 Mean group and individual SGE by counter-geoengineering treatments

	Treatment 4	Treatment 5
Endowments Counter-geoengineering Aggregation Tech	Equal Counter Summation	Unequal Counter Summation
Group net SGE (net of conflict)		
Social optimum	6.00	6.00
Predicted	6.00	6.00
Observed	6.06 (8.78)	6.45 (7.93)
Group total SGE (SGE+|counter SGE|)		
Predicted	36.00^	37.50^
Observed	19.32 (6.74)	24.01 (6.56)
Individual SGE		
Player A ($g_A^* = 2$)	−5.36 (5.71)	−7.52 (5.90)
Player B ($g_B^* = 6$)	3.22 (3.60)	3.97 (4.49)
Player C ($g_C^* = 10$)	8.20 (5.21)	10.00 (4.08)

Unconditional means in table; p-values in text based on conditional tests; standard deviations in parentheses; ^indicates an average because predicted levels with heterogeneous endowments depend on endowment levels; Treatment 2 predicted levels average 36.0 and range from 26.0 to 46.0 and Treatment 4 predicted levels average 37.5 and range from 29.0 to 45.0

Table 6 Mean group and individual SGE in treatment 4 by endowment

	E = 100	E = 150	E = 200	$\Delta_{200-100}$
Group outcomes				
Net SGE	3.64 (6.72)	6.46 (7.98)	7.75 (10.70)	4.11 (112.9%)
Total SGE	16.50 (4.93)	19.13 (5.93)	22.08 (7.89)	5.58 (33.8%)
Player outcomes				
Player A ($g_A^* = 2$)	−5.30 (3.89)	−5.10 (5.16)	−5.72 (7.48)	−0.42 (−7.9%)
Player B ($g_B^* = 6$)	2.88 (3.88)	2.86 (3.41)	3.97 (3.47)	1.09 (37.8%)
Player C ($g_C^* = 10$)	6.06 (4.25)	8.70 (4.97)	9.50 (5.72)	3.44 (56.8%)

Unconditional means in table; p-values in text based on conditional tests; standard deviations in parentheses in columns 2-4; percentage increases from E=100 in column 5

From Table 5, the observed average net production of SGE is close to the hypothesized value of 6 in each treatment (6.06 and 6.45), and our tests generally support this hypothesis ($p = 0.344$ and $p = 0.086$ for Treatments 4 and 5, respectively). Moreover, the observed net production of SGE is statistically equivalent between Treatments 4 and 5 ($p = 0.581$), suggesting that heterogeneity in endowments does not impact these decisions given counter-geoengineering opportunities.

Net production of SGE when counter-geoengineering is possible can be deceiving. Average total production in SGE, defined as the sum of the absolute values of positive and negative production of SGE, is, of course, much higher than net production (19.3 and 24.0 for Treatments 4 and 5, respectively) but lower than the predicted totals ($p = 0.066$ and $p = 0.038$, respectively). The implication is that although investments are made in both SGE and counter-SGE, behavior is less extreme in both directions than the predictions in

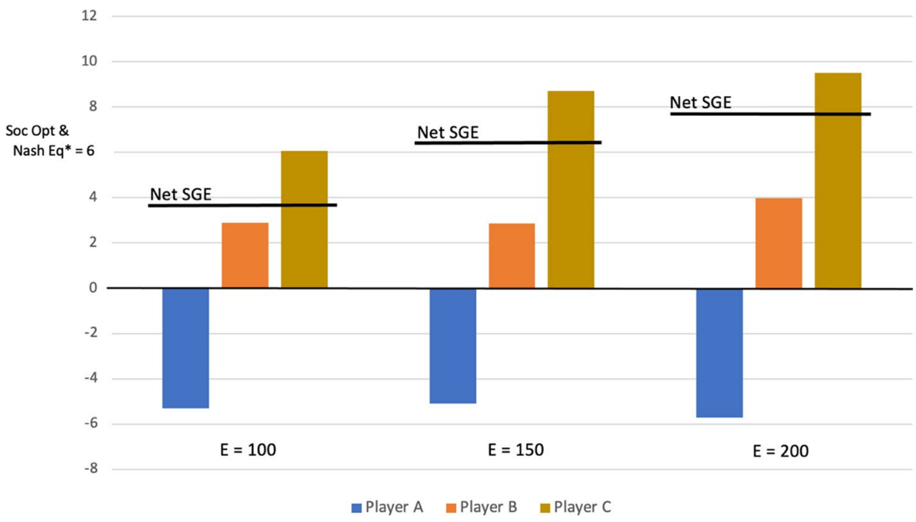

Fig. 3 Contribution to net SGE in treatment 4 by player type and endowment

Table 2. Players engage on average in less conflict and waste fewer resources than theory predicts, an observation similar to Abatayo et al. (2020). When comparing total SGE between both treatments, as predicted, total SGE production is higher when endowments are heterogeneous ($p=0.082$).

Table 6 and Fig. 3 reports the results from Treatment 4 by the level of the endowment. Recall (from Table 2) the prediction that net SGE production is unaffected by endowment levels, but total SGE production should increase with endowments. From Table 6 and Fig. 3, we observe an increase in both net and total SGE production as endowment levels increase. In contrast to the theoretical predictions, the difference in production is significant for net SGE (4.11, $p=0.016$) but insignificant for total SGE (5.58, $p=0.107$). Therefore, we find some evidence that endowment levels matter, but only for net SGE production.

Now consider the individual SGE choices. Recall from Table 2, Player B's predicted production of SGE is unaffected by endowment levels, and the absolute production of SGE is predicted to increase in endowment levels for Players A and C. Our data show that Player A and B's average investment in SGE was not significantly impacted by the different endowment levels ($p=0.574$ and $p=0.434$ for Player A and B, respectively). The change in Player C's average SGE production was positive and marginally significant ($p=0.058$).

This is only part of the story, though. As in Abatayo et al. (2020), efficiency in the counter-engineering treatments (Treatments 4 and 5) are impacted by two factors—the counter-engineering conflict (C chooses a high SGE level, expecting that A chooses a high negative level, and vice versa) and the lack of coordination. The latter factor can be illustrated with a numerical example—if C chooses 10 in two periods and A chooses −10 in two periods, then obviously the averages are 10 and −10 over those two periods, and both times B's choice will determine the group level of SGE. If C chooses 5 and 15 in two periods while A chooses −15 and −5, then average choices are still 10 and −10. But group-level SGE in the first period equals −10 plus B's choice and in the second period 10 plus B's choice. The averages hide a coordination problem that has huge efficiency effects. For this reason, economic surplus in Treatments 4 and 5 is lower than predicted in the equilibrium,[11] since in equilibrium surplus losses are only due to conflict but not due to coordination.

5 Conclusion

The challenge of managing climate change is typically thought of as a problem of providing a global public good. Countries would be collectively better off reducing greenhouse gas emissions below dangerous thresholds, but each has an incentive to free ride off others' efforts while maintaining business as usual. However, when countries have opportunities to address the climate problem complementarily by investing in SGE technologies, a different set of challenges arises. Given the low cost of deployment, it is possible for a single wealthy nation to quickly impact the global climate. The capable country with the highest preferred level of SGE could act as a "free driver," deploying a level of SGE that is inefficiently high relative to the global optimum (Wagner and Weitzman 2012; Weitzman 2015).

The governance challenges posed by potential free drivers has motivated social science research to consider the problem (Weitzman 2015; Abatayo et al. 2020; Aldy et al. 2021; Ghidoni et al. 2023). Our study provides an empirical test of the free-driver hypothesis using a set of controlled economic experiments. Our experimental design replicates and extends the basic setup from the seminal experimental work of Abatayo et al. (2020) by varying and interacting whether endowments are homogeneous or heterogeneous, the size of the endowments, and whether players can invest in counter-SGE. In addition, we test the effect of whether group-level SGE is determined by summation or best-shot technologies.

Consistent with Abatayo et al. (2020), in our baseline treatment we find clear support for the free-driver hypothesis at the group-level. However, we also observe that individual behavior is less consistent with predictions. Although the predicted "free driver" deploys the highest level of SGE, the level is less than their monetarily individual optimum, and we observe positive levels of deployment by other members. In effect, the free driver achieves their desired level without unilaterally bearing the entire cost of deployment. The implication for the real world is that even if a country deploys less than their theoretically ideal level, overall deployment might still be inefficiently high. Abatayo et al. (2020) also observe positive deployment by members other than the predicted free driver, but at a noticeably smaller percentage. In their study, the predicted free driver, on average, contributes 86% of total deployment while we observe average contributions of 64%. This inconsistency may be due to important differences between the two experimental designs, which we detail in Sect. 3. In particular, unlike Abatayo et al. (2020) we do not allow communication between subjects and we use a strangers matching protocol.

Additionally, we observe that changing endowment levels, as predicted, has no effect on behavior, but introducing heterogeneity in endowments reduces the average production of SGE (moving slightly toward the social optimum). This observation offers a silver lining—countries in the real world are highly endogenous along several dimensions (vulnerability to climate change, capacity for SGE deployment, aversion to global inequality, etc.), and our findings indicate that endowment heterogeneity might have a positive effect on SGE deployment, consistent with what various papers on heterogeneity in standard linear public good experiments report.

[11] Actual surplus range: 24.6%-50.1%. Equilibrium surplus range: 59.2%-67.7%.

Barrett (2007, p.38) states that "geoengineering essentially constitutes a large project, a single best effort." The idea is that the country that invests the most resources into its deployment capabilities will have the ability to set the global thermostat. In theory, that country is the predicted free driver and there is little reason for other countries to make positive investments. We implement a "best shot" treatment to test the implications relative to summation deployment. The best-shot feature also reduces the strategic ambiguity that exists in the simultaneous summation deployment used in Abatayo et al. (2020) and our Treatments 1, 2, 4 and 5. Here we observe lower levels of average net SGE compared to the summation game, even though the free driver deploys more SGE (closer aligned with the theoretical prediction), and significant resources are wasted in the process. Although we cannot say how SGE will be deployed in practice (if it ever is), our results suggest that if deployment is a single best effort then we may observe behavior that closer correlates to free driver predictions. However, we do observe a significant amount of wasted resources by the other players in the best shot formulation, suggesting that other behavioral motives are at play.

When players can invest in counter-SGE, average net SGE is consistent with the theoretical predictions—groups reach the socially optimal level, but at significant costs as players engage in zero-sum deployment of SGE and countermeasures. However, we observe less conflict and wasted resources than predicted. Thus, we find that, although investments are made in both SGE and counter-SGE, behavior is less extreme in both directions than theory predicts. In contrast with predictions, higher endowments cause an increase in average net SGE. Finally, introducing heterogeneity in endowments has no impact on average net SGE production when players can invest in counter-SGE. These results offer competing implications for SGE in the field: As in Abatayo et al. countermeasures might not be as inefficient as theory predicts, but richer countries might increase their geoengineering more than predicted.

Although many scientists are confident that SGE technologies could be effective (e.g., Keith 2001; Pope et al. 2012), these technologies have yet to be deployed. In the absence of field data, we turn to controlled experimentation to investigate how free-driver incentives impact behaviors in a simulated environment. Across treatments, we find strong evidence that aggregate deployment of SGE and counter-SGE is inefficiently high. However, on average, the supposed free drivers behave less extreme than predicted.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/s10640-024-00854-1.

Acknowledgements We thank Juan Moreno-Cruz, Mark Borsuk, Tyler Felgenhauer, Khara Grieger, Billy Pizer and Jonathan Wiener for valuable comments.

Funding This material is based upon work supported by the National Science Foundation under Grant No. 2033855.

Declarations

Competing interests The authors declare no competing interests.

References

Abatayo AL, Bosetti V, Casari M, Ghidoni R, Tavoni M (2020) Solar geoengineering may lead to excessive cooling and high strategic uncertainty. Proc Natl Acad Sci 117(24):13393–13398

Aldy J, Felgenhauer T, Pizer WA, Tavoni M, Belaia M, Borsuk ME, Ghosh A, Heutel G, Heyen D, Horton J, Keith D, Merk C, Cruz JM, Reynolds JL, Ricke K, Rickels W, Shayegh S, Smith W, Tilmes S, Wagner G, Wiener JB (2021) Social science research to inform solar geoengineering. Science 375(6569):815–818

Barrett S (2007) Single best efforts: global public goods that can be supplied unilaterally or minilaterally. Why cooperate? The incentive to supply global public goods. Oxford University Press, Oxford, pp 22–46

Biermann F, Oomen J, Gupta A, Ali SH, Conca K, Hajer MA, Kashwan P, Kotzé LJ, Leach M, Messner D, Okereke C, Persson Å, Potocnik J, Schlosberg D, Scobie M, Van Deveer SD (2022) Solar geoenegineering: the case for an international non-use agreement. Wires Clim Change 13(3):754. https://doi.org/10.1002/wcc.754

Chan KS, Mestelman S, Moir R, Muller RA (1999) Heterogeneity and the voluntary provision of public goods. Exp Econ 2(1):5–30

Chen DL, Schonger M, Wickens C (2016) oTree–an open-source platform for laboratory, online and field experiments. J Behav Exp Financ 9:88–97

Cherry TL, Kroll S, Shogren JF (2005) The impact of endowment heterogeneity and origin on public good contributions: evidence from the lab. J Econ Behav Organ 57(3):357–365

Cherry TL, Kroll S, McEvoy D, Campoverde D, Cruz JM (2022) Climate cooperation in the shadow of solar geoengineering: an experimental investigation of the moral hazard conjecture. Environ Polit. https://doi.org/10.1080/09644016.2022.2066285

Ghidoni R, Abatayo AL, Bosetti V, Casari M, Tavoni M (2023) Governing climate geoengineering: side-payments are not enough. J Assoc Environ Resour Econ 10(5):1149–1177

Greiner B (2015) “Subject pool recruitment procedures: organizing experiments with ORSEE. J Econ Sci Assoc 1(1):114–125

Keith DW (2001) Geoengineering. Nature 409(6818):420–420

Kroll S, Cherry TL, Shogren JF (2007) The impact of endowment heterogeneity and origin on contributions in best-shot public good games. Exp Econ 10(4):411–428

Moreno-Cruz JB (2015) Mitigation and the geoengineering threat. Resour Energy Econ 41:248–263

National Academies of Sciences, Engineering, and Medicine (NASEM) (2021) Reflecting sunlight: recommendations for solar geoengineering research and research governance. The National Academies Press, Washington, DC

Pope FD, Braesicke P, Grainger RG, Kalberer M, Watson IM, Davidson PJ, Cox RA (2012) Stratospheric aerosol particles and solar-radiation management. Nat Clim Chang 2(10):713–719

Ricke KL, Moreno-Cruz JB, Caldeira K (2013) Strategic incentives for climate geoengineering coalitions to exclude broad participation. Environ Res Lett 8(1):014021

Wagner G (2021) Geoengineering: the gamble. Polity Press, Cambridge

Wagner G, Weitzman M (2012) Playing god. Foreign Policy, Washington, DC

Weitzman M (2015) A voting architecture for the governance of free-driver externalities. Scand J Econ 117:1049–1068