Discussion papers
E-papers of the Department of Economics e Management – University di Pisa

Nicola Campigotto, Marco Catola,
Simone D'Alessandro, Pietro Guarnieri,
Lorenzo Spadoni

Curbing Energy Consumption through Voluntary Quotas: Experimental Evidence

Discussion paper n. 299

2023

Discussion paper n. 299, presented: **July 2023**

Authors' address/Indirizzo degli autori:

Nicola Campigotto — Department of Economics and Management, University of Trento, Via Vigilio Inama, 5, 38122 Trento TN, Italy. E-mail: nicola.campigotto@unitn.it

Marco Catola — School of Business and Economics, Maastricht University, Tongersestraat 53, 6211 LM Maastricht, Netherlands. E-mail: m.catola@maastrichtuniversity.nl

Simone D'Alessandro — Department of Economics and Management, University of Pisa, Via Cosimo Ridolfi, 10, 56124 Pisa PI, Italy. E-mail: simone.dalessandro@unipi.it

Pietro Guarnieri — Department of Economics and Management, University of Pisa, Via Cosimo Ridolfi, 10, 56124 Pisa PI, Italy. E-mail: pietro.guarnieri@unipi.it

Lorenzo Spadoni — Department of Economics and Law, University of Cassino and Southern Lazio, Via S. Angelo – Località Folcara 03043 Cassino (FR), Italy. E-mail: lorenzo.spadoni@unicas.it

Please cite as:/Si prega di citare come:

Nicola Campigotto, Marco Catola, Simone D'Alessandro, Pietro Guarnieri, Lorenzo Spadoni (2023), "Curbing Energy Consumption through Voluntary Quotas: Experimental Evidence", Discussion Papers, Department of Economics and Management – University of Pisa, n. 299 (http://www.ec.unipi.it/ricerca/discussion-papers).

Discussion Papers Series contact: pietro.battiston@unipi.it

Discussion paper

n. 299

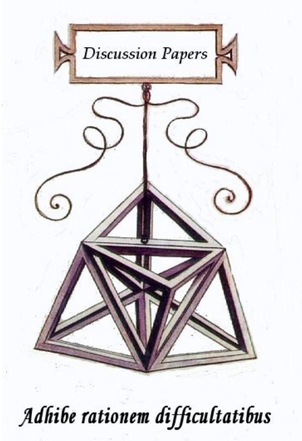

Nicola Campigotto, Marco Catola, Simone D'Alessandro, Pietro Guarnieri, Lorenzo Spadoni

Curbing Energy Consumption through Voluntary Quotas: Experimental Evidence

Abstract

This paper explores the potential of voluntary consumption quotas as a strategy to address resource supply shortages. The results of an incentivized online experiment are presented in which a Nash demand game was used to model an energy consumption problem. Participants had the option to join an energy conservation programme by accepting a consumption quota. Those who accepted the quota traded off their maximum demand for energy in exchange for the certainty that their demand would be met, while those who rejected the quota could demand and possibly earn more but risked suffering from a power outage, in which case they received nothing. Three different quota schemes are examined, and their policy implications are discussed. Our findings suggest that voluntary quotas may lead to a significant decrease in overall demand and contribute to enhancing consumption security.

Keywords: energy consumption; online experiment; Nash demand game; power outages; voluntary quotas.

JEL CLassification: C72; C99; Q48.

Curbing Energy Consumption through Voluntary Quotas: Experimental Evidence

Nicola Campigotto[1,*], Marco Catola[2], Simone D'Alessandro[3], Pietro Guarnieri[3], Lorenzo Spadoni[4]

[1]Department of Economics and Management, University of Trento
[2]School of Business and Economics, Maastricht University
[3]Department of Economics and Management, University of Pisa
[4]Department of Economics and Law, University of Cassino and Southern Lazio
[*]Corresponding author. Email: nicola.campigotto@unitn.it

This version: July 4, 2023

Abstract

This paper explores the potential of voluntary consumption quotas as a strategy to address resource supply shortages. The results of an incentivized online experiment are presented in which a Nash demand game was used to model an energy consumption problem. Participants had the option to join an energy conservation programme by accepting a consumption quota. Those who accepted the quota traded off their maximum demand for energy in exchange for the certainty that their demand would be met, while those who rejected the quota could demand and possibly earn more but risked suffering from a power outage, in which case they received nothing. Three different quota schemes are examined, and their policy implications are discussed. Our findings suggest that voluntary quotas may lead to a significant decrease in overall demand and contribute to enhancing consumption security.

Keywords: energy consumption, online experiment, Nash demand game, power outages, voluntary quotas.

JEL Classification: C72, C99, Q48.

1 Introduction

Global energy consumption has been increasing for more than half a century, hitting a record high in 2021. Although the Covid-19 pandemic has delivered a large shock to the energy sector, it has not taken long for demand to bounce back to old levels (IEA 2021; Ritchie et al. 2022). At the same time, however, a number of factors — including increasingly frequent extreme climate events and international conflicts that made fossil fuel prices spike — have led several countries to struggle to balance energy supply with demand. For example, in 2021 China and Texas suffered from unprecedented power crunches that plunged millions into darkness, while in 2022 European governments urged households and firms to reduce their energy consumption over fears for gas supplies (European Commission 2022). Similar news frequently comes from poor economies and remote areas, where energy security remains a major issue (Cole et al. 2018; Longden et al. 2022).

Long-term interventions to prevent power shortages include investing to expand generation capacity and improve the distribution network. Shorter-term measures instead typically involve reducing strain on the power grid, for instance by capping heating and cooling or by cutting power through load shedding at times of grid stress. However, since cuts are mandatorily enforced and disrupt daily lives and businesses, they inevitably spark discontent, especially when they are not anticipated by the public. Moreover, there are policymakers who remain reluctant to introduce rationing measures and tell people how often and for how long they should use energy-consuming appliances, possibly due to concerns for their libertarian image (see e.g. Lawson 2022).

Building on these points, scholarly attention has turned to strategies that allow users to sort themselves into different energy consumption schemes. These notably include interruptible electricity contracts (Baldick et al. 2006), which provide rebates to those who accept outages. As argued by Allcott et al. (2016), if distribution companies gave users the choice between interruptible and uninterruptible contracts, then outages would be allocated towards those users who are least affected by them. Another approach is to accommodate individual preferences by means of tradable quotas, which combine a cap on overall consumption with the use of market mechanisms to allocate demand (Maurer et al. 2005). One downside of such a tradable system is that it is scarcely suitable for small consumers. Moreover, it can involve substantial costs associated with creating a new market and ensuring its smooth functioning.

This paper explores yet a different mechanism for addressing energy shortages, namely voluntary quotas that trade-off consumption for security. Voluntary quotas can be introduced as contracts by which users willingly limit their maximum consumption of energy in exchange for the guarantee that they will not suffer, or suffer as little as possible, from outages. Unlike mandatory rationing, voluntary quotas are not imposed without consent; unlike tradable quotas, they do not require the setting up of a new market. In addition, the characteristics of voluntary quotas make them inherently different in nature from price-based interventions aimed at reducing energy use, such as dynamic pricing schemes that make the price of energy increase at times of high demand. This is because the burden from higher prices falls disproportionately on low-income users, resulting in a more unequal access to energy, whereas quotas give all users the same opportunity to self-regulate and secure a sustainable consumption level.

We study the effects of voluntary quotas in the simplified framework of an incentivized online experiment. Eight hundred UK residents were recruited to play a Nash demand game that captures key features of electricity consumption decisions. Subjects were given a production task and decided independently how much energy to demand to carry out the task. Higher energy use resulted in increased production and a higher payoff. Subjects were told that energy was supplied by a generator and that if the sum of individual demands did not exceed generator capacity, then each subject would have their demand satisfied. Conversely, if total demand exceeded capacity, then each would suffer a power outage, produce nothing, and receive a payoff of zero. Demand decisions were made in two consecutive rounds. In Round 1 capacity was fixed, while in Round 2 it could either remain unchanged or decrease according to a known probability distribution as a result of a supply shortage.

The experiment consisted of three treatments which featured different voluntary quotas, plus an untreated baseline. Subjects who accepted the quota (henceforth acceptors) saw their maximum per-period demand reduced but were sure to get what they demanded. Subjects who did not accept the quota (henceforth rejectors) could demand and possibly earn more than acceptors but run the risk of coming up empty-handed. The first quota varied with capacity and was designed so that in the event of unanimous acceptance and maximum energy demand by all subjects, capacity was exactly exhausted both in the presence and absence of an energy supply shortage. The second quota varied with capacity too, but it entitled subjects to demand less energy than the first and it always resulted in some residual capacity. The third quota entitled subjects to demand a fixed amount of energy and could exhaust capacity only in the event of an energy shortage. These treatments allow us to investigate the impact of different quota schemes on overall consumption, which is not obvious a priori. For example, a quota that restricts demand by a small extent may be welcomed by the public but be of little help in curbing aggregate consumption, whereas a quota that considerably restricts demand may fail to meet consumption reduction objectives because it is not appealing to users.

The results suggest that although voluntary quotas do not suffice to prevent outages with certainty, they may significantly contribute to reducing demand and relieving the stress upon supply systems. Depending on treatment, between 53 and 77 percent of subjects accepted the quota that was proposed to them. When no supply shortage occurred, quotas resulted in a reduction of up to 35 percent in aggregate demand and up to 82 percent in the frequency of outages compared to the baseline condition. In case of shortage, outage frequency decreased by a less pronounced but still clearly discernible extent (up to 35 percent less than the baseline). As detailed in the discussion below, the choice of what kind of quota to introduce would ultimately depend on the energy provider's objective.

The remainder of the paper is organized as follows. Section Section 2 discusses some related literature. 3 introduces the experiment and Section 4 provides some benchmark theoretical results. Section 5 discusses our experimental findings. Section 6 provides concluding remarks and outlines directions for future research.

2 Related literature

Our work relates to several strands of literature. Since voluntary quotas are intended to limit energy demand while promoting energy security, they can be a means towards achieving sufficiency in energy consumption (Princen 2005; O'Neill et al. 2018). Put simply, the concept of sufficiency involves reducing the consumption of energy, and consequently its environmental impact, to a sustainable level that is consistent with equality and well-being. Voluntary quotas are also consistent with the idea of libertarian paternalism (Thaler and Sunstein 2003), as they are aimed at affecting behaviour while respecting freedom of choice.

In terms of aim and methodology, our paper shares similarities with the theoretical and experimental literature on collective climate action. This literature is heterogeneous in experimental design and relies on a diverse range of games, such as the dictator game (Czap et al. 2018), the public goods game (Alpizar and Gsottbauer 2015; Calzolari et al. 2018), and the collective-risk social dilemma (Milinski et al. 2008; Tavoni et al. 2011; Farjam et al. 2019). The approach closest in spirit to ours involves the use of a common pool dilemma, which has been widely employed to study the issue of resource over-exploitation (Bernard et al. 2013; Berger and Wyss 2021), particularly in contexts characterized by environmental uncertainty (Aflaki 2013; Anderies et al. 2013; Bochet et al. 2019), and the use of quotas for common pool management (Ostrom 1999; Cardenas et al. 2000). However, it is worth noting that in common pool dilemmas, the unconsumed resources are multiplied by a certain factor and redistributed among the players. Given our aim of reducing energy consumption, we made the decision to employ a Nash demand game instead, where the experimenter retains the resources that are not consumed.

3 Experimental design

3.1 Setup

Consider a finite group of $n \geq 2$ players who consume a limited shared resource of size C. Choices are made independently and simultaneously. Suppose initially that each player i can demand any amount $d_i \in \left[0, \overline{d}\right]$ of the resource, where $\overline{d} < C < n\overline{d}$. If the demands sum to C or less, then players receive what they demanded; otherwise they get nothing. Player i's payoff is therefore:

$$\pi_i\left(d_i, d_{-i}\right) = \begin{cases} d_i & \text{if} \quad \sum_{j=1}^{n} d_j \leq C \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

where d_{-i} denotes the vector of all demands excluding i's.

Now suppose instead that before demands are made, each player is offered the choice to either accept or reject a consumption quota. This choice has consequences on both the player's choice set and their payoff. If the quota is rejected, then the set of possible demands and payoff are as described above; conversely, accepting the quota means reducing the maximum possible demand to an amount $\underline{d} < \overline{d}$ in exchange for the guarantee that demand will be met. Assume that $n\underline{d} \leq C$, so that if all players accept the quota and demand as much as possible, then group demand does not exceed the available resource. We call this the non-exceedance condition. Let a be a dichotomous variable equal to Y if the quota is accepted and N if it is rejected. Thus, player i's demand is:

$$d_i(a_i) \in \begin{cases} [0, \underline{d}] & \text{if} \quad a_i = Y \\ \left[0, \overline{d}\right] & \text{otherwise} \end{cases}$$

and i's payoff function becomes:

$$\pi_i\left(a_i, d_i(a_i), a_{-i}, d_{-i}(a_{-i})\right) = \begin{cases} d_i & \text{if} \quad a_i = Y \quad \text{or} \quad \sum_{j=1}^{n} d_j \leq C \\ 0 & \text{otherwise} \end{cases}. \tag{2}$$

As an example, consider the case where k players accept the quota and $n - k$ players reject it. Let the subsets of acceptors and rejectors be S_1 and S_2, respectively. Acceptors can demand at most $k\underline{d}$ and their demands will be met with certainty, whereas rejectors will receive their demands if and only if:

$$\sum_{j \in S_2} d_j \leq C - \sum_{j' \in S_1} d_{j'}$$

that is if and only if their aggregate demand does not exceed the resource left after meeting acceptors' demands. If this is not the case, then rejectors will receive nothing regardless of how much they demanded.

3.2 Framing and treatments

Subjects were randomly assigned to groups of 10 people. To give the experiment a meaningful context and reduce confusion among participants, instructions were framed in terms of energy use (Alekseev et al. 2017; a full list of instructions is available in the Supplementary Material). Each subject was told to suppose that "You are a tailor making shirts. You use an electric sewing machine, the energy consumption of which is measured in 'Energy Units'. To produce one shirt you must use 1 Energy Unit. The more Energy Units you use, the more shirts you produce." For simplicity, Energy Units could only be consumed in integer amounts. Note that contextual instructions could alternatively have been given in terms of residential power consumption, in which case payoffs would have corresponded to the utility from using electric appliances net of electricity cost. We chose to focus on production for ease of exposition.

The shared resource was framed as the capacity of a generator: "Your machine is powered by an electric generator, and so are the machines of nine other tailors. Each tailor's machine consumes 1 Energy Unit per shirt, just like yours." Finally, each participant was informed that "If in a decision round the overall number of Energy Units used by the 10 of you exceeds the generator's capacity, then there will be a power failure, in which case you will produce nothing in that round."

Subjects had to decide how many Energy Units they would like to use in each of two rounds. In Round 1, generator capacity was 100 Energy Units. In Round 2, there was a 50 percent probability that capacity would remain steady at 100 Energy Units, and a 50 percent probability that it would decrease to 50 Energy Units due to an energy supply shortage. This information was given to each subject at the beginning of Round 1; at the beginning of Round 2, participants were then informed about the actual (realized) capacity, i.e. 50 or 100 Energy Units. Throughout the paper we write Round 2_x to denote the second decision round when capacity was $x \in \{50, 100\}$.

To isolate the effect of a fall in capacity, no feedback was given to participants between the first and second rounds. This is similar to asking subjects how they would behave in two distinct situations, the difference being that the choice whether to accept or reject the quota was made only once at the beginning of Round 1. Thus, we do not assess learning effects. We are comfortable that the experiment can nevertheless yield insightful results, as it captures aspects of strategic behaviour relevant to situations where agents face sharp reductions in energy supply and do not have the time to learn how to coordinate.

In the baseline condition, participants were not offered any quota and could demand any whole amount of Energy Units between 0 and $\overline{d} = 20$ per round. Subject i's payoff in Round $t = 1, 2$ was therefore $\pi_{i,t} = d_{i,t} \in \{0, 1, \ldots, 20\}$ if $\sum_{j=1}^{10} d_{j,t} \leq C_t$ and zero otherwise, where $C_1 = 100$ while C_2 could be either 50 or 100 (each with probability 1/2).

The design had three treatments (denoted by the superscripts F, HP and LP, as shorthands for *Fixed*, *High Proportional* and *Low Proportional*), each featuring a different quota scheme. Before making their first demand decision, subjects were asked whether they wanted to participate in an energy programme based on voluntary quotas. In Treatment F, each participant was informed that "If you accept the quota, then in each round you can use up to 5 Energy Units. In this case your energy use is guaranteed, meaning that you can operate your machine regardless of other participants' decisions. Conversely, if you don't accept the quota, then your energy use is not guaranteed, but you can use up to 20 Energy Units." In this treatment, accepting the quota meant therefore reducing the maximum possible demand per round from $\overline{d} = 20$ to

$$\underline{d}^F = 5$$

and being certain to earn a payoff $\pi_{i,t} = d_{i,t}(Y) \leq 5$ in each round.

Treatments HP and LP differed from Treatment F in that the quota was not fixed but proportional to generator capacity. In Treatment HP, acceptors were entitled to use a number of Energy Units equal to 10 percent of capacity. The maximum possible demand in each round was therefore:

$$\underline{d}^{HP} = 10 \times \mathbb{1}_{C=100} + 5 \times \mathbb{1}_{C=50}$$

where $\mathbb{1}$ denotes the indicator function. Finally, in Treatment LP the quota amounted to 6 percent of capacity, implying that:

$$\underline{d}^{LP} = 6 \times \mathbb{1}_{C=100} + 3 \times \mathbb{1}_{C=50}.$$

Quotas were designed to satisfy the non-exceedance condition, i.e. so that in the event of unanimous acceptance, capacity was enough to accommodate group demand both in the presence and absence of an energy shortage. If all players took the quota and demanded as much as possible, then in Treatment HP capacity would be exhausted both when $C = 100$ and when $C = 50$; in Treatment F it would be exhausted when $C = 50$ but not when $C = 100$; in Treatment LP it would be exhausted in neither case. It is also easy to check that an acceptor's overall expected payoff would be highest in Treatment HP (17.5) and similar in Treatments F and LP (10 and 10.5, respectively). This 4×2 design allowed us to assess and compare on the one hand the effects of a high and a low quota (HP vs. LP), and on the other the effects of a fixed quota and a quota that varies with capacity (F vs. LP). Table 1 shows the maximum possible demand in the baseline and treatments, while Figure 1 summarizes the timing of the game.

Table 1: Maximum possible demand by treatment and capacity

	Baseline	Treatment F		Treatment HP		Treatment LP	
		Quota accepted	Quota rejected	Quota accepted	Quota rejected	Quota accepted	Quota rejected
$C = 100$	20	5	20	10	20	6	20
$C = 50$	20	5	20	5	20	3	20

Figure 1: Timing

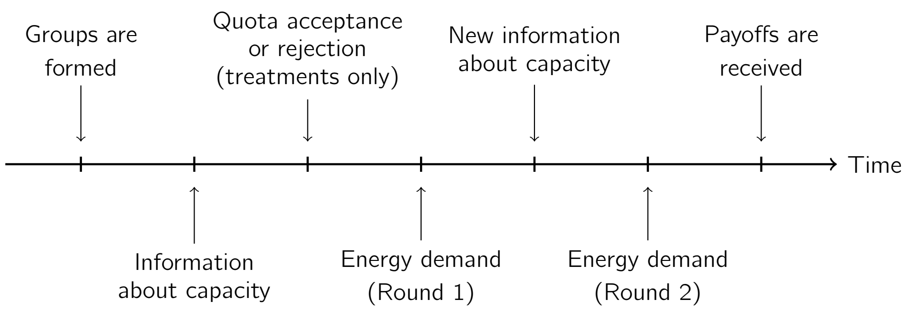

3.3 Implementation

The experiment was pre-registered at aspredicted.org (#78308) and programmed in oTree (Chen et al. 2016). Eight hundred UK residents between 18 and 45 years of age, equally divided between females and males, were recruited through Prolific (Palan and Schitter 2018) and randomly assigned to treatment conditions.

After the second decision round and before receiving any feedback about their earnings, participants were required to fill in a short questionnaire. Five questions of the kind "How often do you do the following?" were asked to assess subjects' energy saving habits. The questions concerned running the washing machine half empty, leaving water running while brushing teeth, showering for more than 10 minutes, leaving the lights on when leaving a room, and turning up the heat instead of putting on warmer clothes when it gets cold. Answers were given on a 5-point scale and were then averaged together to obtain a single energy saving score, with higher values indicating a more judicious energy saving behaviour. The questionnaire also used experimentally validated survey questions to elicit participants' self-reported altruism and willingness to take risks. The latter was measured by responses to the question "How do you see yourself: are you generally a person who is fully prepared to take risks or do you try to avoid taking risks?", while the former was measured by responses to "Imagine the following situation: today you unexpectedly received 1,000 GBP. How much of this amount would you donate to a good cause?" (Dohmen et al. 2011; Falk et al. 2023). Responses to both questions were given on a 0-10 scale, with higher values representing higher risk-taking and altruism.[1]

Additional individual characteristics were available from Prolific, including age, sex, ethnicity, education, employment status, subjective socio-economic status, political stance, and climate change beliefs. In the paper we use the term climate skeptic to refer to subjects who answered "no" or "don't know" to Prolific's question "Do you believe in climate change?". An overview of subject characteristics is given in Appendix A.

Subjects were paid at a pre-announced rate of 5 pence per payoff unit (that is 1 GBP per 20 payoff units), in addition to a fixed show-up fee of 0.40 GBP. On average the experiment took 3.7 minutes to complete. Average earnings were 0.98 GBP per subject, which amounts to roughly 16 GBP per hour.

4 Theoretical benchmarks

4.1 Baseline

In the baseline game, a strategy for player i is a triple:

$$(d_{i,1}, d_{i,2_{100}}, d_{i,2_{50}})$$

where $d_{i,t}$ denotes individual demand in Round t. Recall that players do not receive any feedback about play between the first and second rounds. The only new information available before the second demand decision is the realized capacity in Round 2, which does not however affect optimal behaviour in Round 1. Thus, in the baseline each round can be analyzed as an independent game.

It is also useful to note that players have no weakly dominated strategy except $d_i = 0$. This is because if $d_{-i} \geq C$, then player i is indifferent between any of their strategies (since i's payoff will be zero regardless of how much they demands), whereas if $d_{-i} < C$, then i's best response is to demand $\min\left\{\bar{d}; C - d_{-i}\right\}$ (since demanding less than this would yield a positive but lower payoff, while demanding more would yield a payoff of zero).

A similar argument shows that the game has multiple Nash equilibria, the characteristics of which are summarized in the lemma below.

[1] Eliciting risk and social preferences at the end of the experiment is common practice in behavioural and experimental economics, see e.g. Cohn et al. (2022) and Non et al. (2022).

Lemma 1: *The baseline game has two classes of Nash Equilibria. The first class consists of all strategy profiles that satisfy:*

$$\sum_{i=1}^{10} d_{i,t} = C_t \tag{3}$$

in which case a positive number of players receive a positive payoff. The second class instead satisfies:

$$\sum_{i=1}^{10} d_{i,t} \geq C_t + \max\{d_{1,t}, \ldots, d_{10,t}\} \tag{4}$$

in which case all players receive a payoff of zero.

Proof. Consider the first class of equilibria. If condition (3) holds, then each player is receiving their demand. Any unilateral downward deviation would reduce the deviating player's payoff, and any upward deviation would trigger an outage. In neither case would the deviating player be better-off.

Now consider the second class of equilibria. If condition (4) holds as an equality, then the player who demands $\max\{d_{1,t}, \ldots, d_{10,t}\}$ can prevent the outage. However, this requires demanding nothing at time t, which is never a profitable deviation. If instead condition (3) holds as a strict inequality, then no unilateral deviation can prevent capacity from being exceeded.

Any strategy profile that does not satisfy either condition (3) or condition (4) cannot be a Nash equilibrium. If $\sum_{i=1}^{10} d_{i,t} < C_t$, then each player has an incentive to increase their demand and earn a higher payoff. Conversely, if $C_t < \sum_{i=1}^{10} d_{i,t} < C_t + \max\{d_{1,t}, \ldots, d_{10,t}\}$, then the player who demands $\max\{d_{1,t}, \ldots, d_{10,t}\}$ has an incentive to reduce their demand to an extent so that the resource is exactly exhausted, thereby receiving a positive payoff rather than zero. ■

Thus, in equilibrium either total demand equals total supply or a tragedy of the commons-like situation occurs. There are several strategy profiles that result in each of these two outcomes. A particularly interesting equilibrium satisfying condition (3) is the one in which each player demands $C_t/10$ and the resource is shared equally. In contrast, an intuitive equilibrium satisfying condition (4) is the one in which all players demand $\overline{d} = 20$ but receive nothing. Note also that although the equilibria in the first class can be characterized by high levels of payoff inequality, all of them are nevertheless Pareto-optimal.

4.2 Treatments

In the three treatments, a strategy for player i becomes a quadruple:

$$(a_i, d_{i,1}, d_{i,2_{100}}, d_{i,2_{50}}) =: s_i$$

where $a_i \in \{Y, N\}$. In what follows we assume that players who accept the quota always demand $\underline{d}$. This is a harmless assumption as conditional on accepting the quota, payoff maximization always requires demanding as much as one can. Importantly, since the decision whether to accept or reject the quota affects demand in both rounds, it is no longer possible to consider the two rounds as separate problems.

As before let S_2 be the subset of players who do not accept the quota. The following proposition characterizes the sets of Nash equilibria in all treatments.

Proposition 1: *A strategy profile is a Nash equilibrium of the game in Treatment $\tau \in \{F, HP, LP\}$ if and only if*

$$\pi_i\left(s_i, s_{-i}\right) \geq \pi_i\left(\left(Y, \underline{d}^{\tau}, \underline{d}^{\tau}, \underline{d}^{\tau}\right), s_{-i}\right) \quad \textit{for every } i \tag{5}$$

and

$$k\underline{d}^{\tau} + \sum_{j \in S_2} d_{j,t} = C_t \tag{6}$$

where k is the number of players who accept the quota.

Proof. If condition (5) is not satisfied, then at least one player who rejected the quota is receiving a payoff which is less than the payoff they would receive by accepting the quota and demanding $\underline{d}^{\tau}$. If condition (6) is not satisfied, then capacity is not exhausted and each acceptor has an incentive to deviate unilaterally and reject the quota, thereby earning a payoff equal to $\underline{d}^{\tau}$ plus all Energy Units left unused. ■

Condition (5) states that in a Nash equilibrium all players must have an expected payoff which is at least as high as the expected payoff received by accepting the quota and demanding as much as possible. Condition (6) requires total demand to equal total supply in each round. The main implication of Proposition 1 is that once a voluntary quota scheme is introduced, each strategy profile resulting in a tragedy of the commons can no longer be Nash, i.e. some if not all players receive a positive payoff in equilibrium. The intuitive reason for this result is that quotas establish a fallback position which allows players to earn a positive payoff with certainty. This observation naturally raises the question of how many players can accept the quota in equilibrium, which is addressed in Proposition 2 below.

Proposition 2: *The maximum number of acceptors in a Nash equilibrium is 6 in Treatment F, 7 in Treatment LP, and 10 in Treatment HP.*

Proof. Consider Treatment HP first. Since the quota amounts to 10 percent of capacity, it is easy to see that if all players accept the quota and demand $\underline{d}^{HP}$, then the two conditions of Proposition 1 are satisfied. Put differently, if 9 players accept the quota, then the tenth player is indifferent between accepting it (and demanding $\underline{d}^{HP}$) and rejecting it (and demanding $C_t - 9\underline{d}^{HP}$, which equals $\underline{d}^{HP}$ by construction).

Now consider Treatment LP. Recall that a player who accepts the quota plays strategy $(Y, 6, 6, 3)$ and has an expected payoff of 10.5.

- If 7 players accept the quota, then the number of Energy Units left unused after acceptors' demands are met is 58 when $C = 100$ and 29 when $C = 50$. Consider the following strategies for the remaining three players: two of them play $(N, 19, 19, 9)$, while the third plays $(N, 20, 20, 11)$. The resulting strategy profile is a Nash equilibrium because (i) players who reject the quota have an expected payoff which is higher than the maximum payoff they could receive by accepting the quota and demanding $\underline{d}^{LP}$, and (ii) the resource is exactly exhausted regardless of whether $C = 100$ or $C = 50$.
- If 8 players accept the quota, then the number of Energy Units left unused after acceptors' demands are met is 52 when $C = 100$ and 26 when $C = 50$. However, the two remaining players can demand at most 40 Energy Units, implying that the resource is exhausted neither in Round 1 nor Round 2_{100}. This makes it impossible to satisfy condition (6).
- A similar argument proves that if 9 or 10 players accept the quota, then condition (6) cannot be met.

Finally, consider Treatment F. In this case, a player who accepts the quota plays strategy $(Y, 5, 5, 5)$ and has an expected payoff of 10.

- If 6 players accept the quota, then the amount of Energy Units remaining after demands are met is 70 when $C = 100$ and 20 when $C = 50$. Consider the following strategies for the remaining four players: two of them play $(N, 17, 17, 5)$, while the other two play $(N, 18, 18, 5)$. The resulting strategy profile is a Nash equilibrium because both conditions (5) and (6) are met.
- If 7 players accept the quota, then the number of Energy Units left unused after acceptors' demands are met is 65 if $C = 100$ and 15 if $C = 50$. However, the three remaining players can demand at most 60 Energy Units, meaning that the resource is exhausted neither in Round 1 nor Round 2_{100}. This makes it impossible to satisfy condition (6).
- Similar reasoning shows that if 8 or more players accept the quota, then condition (6) cannot be met.

■

These results yield some general predictions about experimental outcomes. First, in view of Lemma 1 we expect the frequency of outages in the baseline game to be high. The straightforward explanation for this is that since the game has multiple equilibria, since in many of these equilibria all players get nothing, and since there is no time for trial-and-error learning, it is easy for subjects to miscoordinate and demand more than can be supplied.

From Proposition 1 we know that the games in Treatments HP, LP and F suffer from a similar issue, which again is likely to result in frequent miscoordination. Nevertheless, these games differ from the baseline game in two important respects. First, in the treatments there is no Nash equilibrium in which everyone receives a payoff of zero. Second, a significant portion of the set of equilibria consists of strategy profiles where a positive number of players accept the quota. Thus, we expect quotas to be taken by a non-negligible proportion of subjects, resulting in less frequent outages than in the baseline condition. It also seems safe to predict that subjects will be keener to accept the HP quota, which entitles them to consume a higher number of Energy Units. Moreover, Treatment HP is the only treatment in which a Nash equilibrium exists where all players accept the quota (see Proposition 2). Along a related line, we conjecture that subjects may tend to prefer the LP quota to the F quota, since the former yields a slightly higher conditional expected payoff and it can be accepted by a greater number of players in equilibrium.

Predicting which quota will be the most effective in reducing energy use and the frequency of outages is not obvious. On the one hand, we may expect Treatment HP to yield the best results, since it gives players the strongest incentive to accept the quota and it gives salience to the equilibrium outcome in which players share the resource equally. On the other hand, however, there is reason to expect that the HP quota will perform worst. To see this note that if in Treatment HP a large number of players accept the quota and demand as much as possible, then the resource is seriously depleted, implying that a high demand by a few lone rejectors can suffice to cause an outage.[2] Conversely, since the F and LP quotas place a tighter constraint on acceptors' demand, these quotas may prove more effective in curbing energy consumption despite being accepted less frequently by subjects.

5 Results

We begin by examining quota acceptances and their determinants. We then proceed to study demand decisions and their implications for individual and group outcomes.

[2]For example, suppose that $C = 100$ and that the quota is rejected by just one player out of 10. If this rejector demands 11 Energy Units or more, then capacity is exceeded.

5.1 Acceptors and rejectors

The decision to accept the quota was taken by the majority of subjects. Consistent with our discussion in Section 4, the acceptance rate was highest in Treatment HP (77.1 percent), followed by Treatments LP and F (64.1 and 53.0 percent, respectively).

Table 2 reports logit estimates of the probability of acceptance. The dependent variable takes value 1 if the quota was accepted and 0 otherwise. The coefficients are average marginal effects with robust standard errors in parentheses. Our main explanatory variables of interest are subjective Risk tolerance and Altruism, the Energy saving score, and dummies for Climate change skepticism and Green Party support.

The estimates shown in column 1 unsurprisingly indicate that risk-averse individuals were significantly more likely than risk-tolerant individuals to take the quota. On average, a one-point increase in Risk tolerance lowers the probability of acceptance by about 2 percent. Acceptance also tends to be positively associated with Altruism. This latter result can be explained by noting that the decision to take the quota leaves more energy for others to use, and may therefore be intended as an altruistic act. The coefficients on Energy saving score and Green Party support have the expected sign but are neither individually nor jointly significant (p-value = 0.57). Conversely, the coefficient on Climate change skepticism is significant at the 1 percent level and suggests that skeptics were considerably less likely to accept the quota. Interestingly, this result dovetails with the empirical finding of Volland (2017) that trust in others is negatively associated with household energy consumption. The underlying argument is that low trust is a known predictor of climate skepticism (Tranter and Booth 2015), which in turn results in a lower likelihood to accept the quota.

In column 2 we include controls for sex, age, age squared, ethnicity (1 if a subject belongs to a minority group and 0 otherwise), education (highest degree earned), employment status (1 if unemployed or non-working student and 0 otherwise), subjective socioeconomic status (scored on a 10-point scale with higher scores corresponding to a higher subjective status), and the experiment completion time in minutes. The results are not sensitive to the inclusion of these variables. Sex is the only control variable with significant predictive power: on average, women were 7.6 percent more likely than men to accept the quota. The non-significance of completion time suggests that subjects did not accept the quota simply in order to make some quick and easy money and without thinking about the task. Had this been the case, the estimated coefficient would have been negative and significant.

Finally, the regression in column 3 includes treatment dummies. The omitted reference category consists of subjects in Treatment LP. Compared to them, subjects

Table 2: Determinants of quota acceptance (logit estimates)

Dependent variable: 1 if the quota is accepted Average marginal effects, robust SEs in parentheses			
	(1)	(2)	(3)
Risk tolerance	−0.020**	−0.018**	−0.016*
	(0.008)	(0.009)	(0.009)
Altruism	0.024*	0.024*	0.029**
	(0.014)	(0.014)	(0.015)
Climate change skeptic	−0.251***	−0.256***	−0.229***
	(0.083)	(0.089)	(0.085)
Energy saving score	0.034	0.027	0.018
	(0.032)	(0.032)	(0.032)
Green Party supporter	0.016	0.017	0.019
	(0.059)	(0.058)	(0.057)
HP quota			0.133***
			(0.045)
F quota			−0.103**
			(0.048)
Completion time		0.008	0.007
		(0.007)	(0.007)
Control variables	no	yes	yes
Observations	601	601	601

One, two, and three asterisks denote significance at the 10, 5, and 1 percent levels, respectively.

in Treatments F and HP were on average 10 percent less likely and 13 percent more likely to accept the quota, respectively. Figure 2 further investigates this by showing how average marginal treatment effects vary with Risk tolerance and Altruism. Increases in Risk tolerance result in essentially no change in the marginal effect of the F quota (represented in blue), whereas an increase in Altruism from its minimum to its maximum raises the marginal effect by about 4 percentage points. In contrast, the marginal effect of the HP quota (represented in orange) decreases by 8 percentage points for a 10-point increase in Altruism and increases by 4 percentage points for a 10-point increase in Risk tolerance. Taken together, the results indicate that highly risk-averse altruists are more likely to accept any quota, whereas risk-tolerant and self-interested individuals are more sensitive to what kind of quota they are offered. This remark is substantiated by the average acceptance probabilities reported in Appendix B. For example, all else being equal, the difference between the probabilities of accepting the HP and F quotas is 17 percentage points higher when Risk tolerance = 10 and Altruism = 0 than when Risk tolerance = 0 and

Altruism = 10 (see Figure B2).

Figure 2: Average marginal treatment effects at representative values of Risk tolerance and Altruism, with 95 percent confidence intervals (reference category: LP quota)

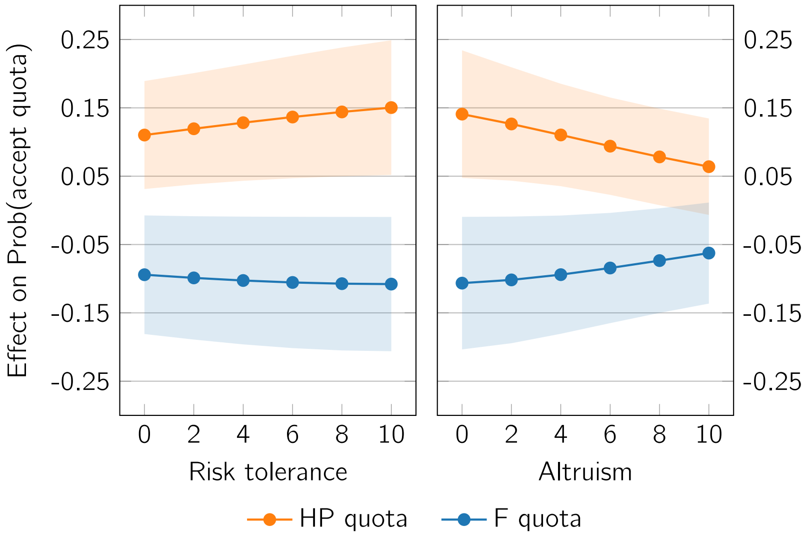

5.2 Energy demand decisions

Figure 3 shows mean individual demands by treatment, capacity level, and acceptance or rejection choice. The results of Rounds 1 and 2 are given in the left- and right-hand panels, respectively. Error bars represent 95 percent bootstrapped confidence intervals. In Round 1, mean demand in the untreated baseline was 13.69 Energy Units (red column). In Round 2, subjects in the baseline condition who did not experience a supply shortage demanded on average 13.22 Energy Units (straight-line red column), whereas subjects who did experience a shortage demanded 10.03 Energy Units (dashed-line red column).

In the treatments a further distinction can be made between the demands of acceptors and rejectors. We describe each in turn. Rejectors' mean demand in Round 1 was 15.0 in Treatment F (blue), 16.5 in Treatment HP (orange), and 15.7 in Treatment LP (green). In Round 2_{100} (i.e. under no-shortage conditions), rejectors' mean demand was 14.5 in Treatment F, 16.5 in Treatment HP, and 15.0 in Treatment LP. In Round 2_{50} mean demands were 11.9, 13.6 and 10.9, respectively.

Acceptors' demands tended to remain close to $\underline{d}^{\tau}$. Mean values ranged between 4.2 and 4.6 in Treatment F, between 4.6 and 9.2 in Treatment HP, and between 2.6 and 5.4 in Treatment LP. Note that acceptors were made aware that the decision to demand less than $\underline{d}^{\tau}$ would result in a forgone payoff but would reduce the likelihood that rejectors would experience an outage. We therefore have confidence that those who accepted the quota but demanded less than $\underline{d}^\tau$ did so deliberately. Overall, 62.5 percent of acceptors demanded as much as possible in both decision rounds, while 77.6 percent did so in at least one round.

In Appendix C we show that no significant difference is found between demand in Round 1 and demand in Round 2_{100}. This is evidence that participants did not adopt hedging strategies. If subjects hedged their demands, then they would behave differently in different rounds, e.g. by alternating low and high demands, to insure themselves against the risk of taking an action that results in a low payoff. Our findings indicate that this was generally not the case.

Figure 3: Mean individual demands by treatment, capacity level, and acceptance or rejection choice

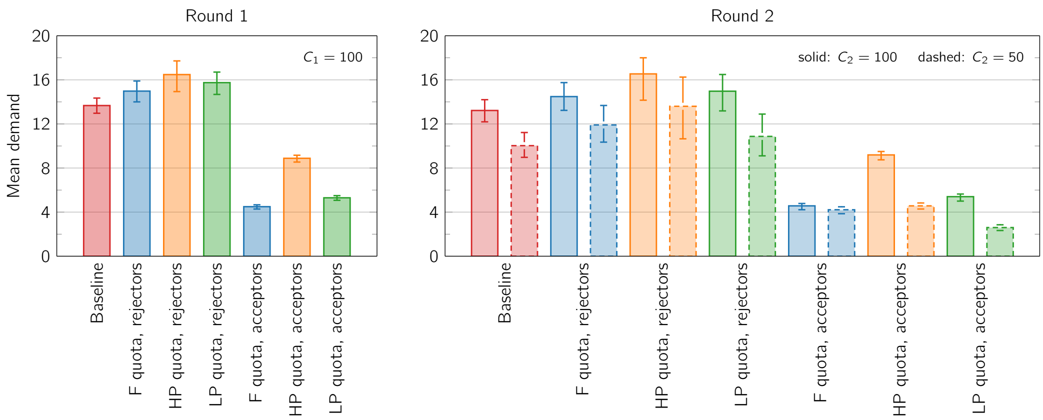

How did subjects react to a supply shortage? Figure 4 makes this clear by plotting the elasticity of individual demand with respect to capacity (calculated using the midpoint method), together with bootstrapped 95 percent confidence intervals. A value of 0 means that demand did not vary with capacity, whereas a value of 1 means that demand and capacity varied in the exact same proportion and direction, i.e. a 1-percent decrease in capacity led to a 1-percent decrease in demand.

Acceptors' elasticity takes values near 1 in Treatments LP and HP, while it does not significantly differ from zero in Treatment F. This is because the two proportional quotas make $\underline{d}^\tau$ change in a one-to-one ratio with capacity, whereas the fixed quota makes no such adjustment. More interestingly, the demand of rejectors and subjects in the baseline is always inelastic, that is the percentage decline in demand is less than the percentage decline in capacity. On average, a decrease in capacity from 100 to 50 Energy Units made demand fall by 42 percent in the baseline and by between 18 and 46 percent in the treatments. This relative insensitivity of demand makes miscoordination and outages considerably more of a problem under capacity shortages than under no-shortage conditions.

Table 3 shows Tobit estimates of the determinants of individual energy demand.

Figure 4: Elasticity of demand with respect to capacity

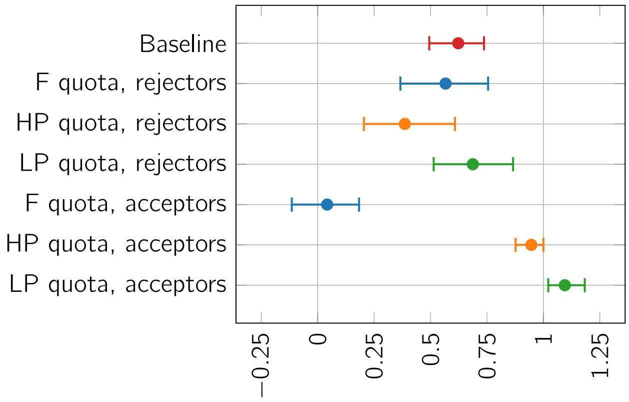

The lower- and upper-censoring limits are 0 and 20, respectively, as demand is bounded within this interval. In column 1 demand is regressed on the set of explanatory variables introduced in Section 5.1, including all control variables, plus treatment and round-capacity dummies. The coefficient on Energy saving score fails to reach statistical significance, while the coefficients on other variables are significant at the 5 percent level or better. All else constant, mean individual demand in Treatments F, HP and LP was −3.9, −3.2 and −4.8 Energy Units lower than in the baseline, respectively. Completion time has a significant but very small effect.

In column 2 treatment dummies are interacted with acceptance and rejection dummies. The reference category continues to be the untreated Baseline. Inclusion of the interaction terms makes the coefficients on Altruism and Climate change skepticism lose significance. Similar results are obtained when restricting the sample to rejectors and subjects in the baseline, i.e. to all those subjects who could demand up to 20 Energy Units. The estimates reported in column 3 show that a unit increase in Risk tolerance is associated with a 0.27 increase in demand, and that Green Party supporters demanded less than supporters of other parties.

Finally, in column 4 the sample of observations is restricted to acceptors. The only significant coefficient is that on Altruism, further confirming that acceptors' decision to demand less than $\underline{d}^{\tau}$ was driven by pure or warm-glow altruistic motives. On average, acceptors in Treatment HP demanded about 3 Energy Units more than acceptors in Treatment LP, whereas acceptors' demands in Treatments F and LP do not significantly differ from one another.

Figure 3 and Table 3 show that rejectors' mean demand was up to 5 Energy Units higher than demand in the baseline. This difference may be due to either or both of two factors. First, the baseline sample partly consists of subjects who would have accepted a quota if offered one. Since acceptors tend to be more risk averse and altruistic than rejectors, we expect these subjects to have demanded less than

Table 3: Determinants of individual demand (Tobit estimates)

Dependent variable: individual energy demand Pooled estimates, robust SEs in parentheses				
	(1)	(2)	(3)	(4)
Risk tolerance	0.229*** (0.075)	0.095* (0.050)	0.235** (0.116)	−0.019 (0.026)
Altruism	−0.310*** (0.104)	−0.102 (0.076)	−0.140 (0.164)	−0.089** (0.043)
Climate change skeptic	1.907** (0.792)	0.096 (0.614)	−0.044 (1.086)	0.120 (0.204)
Energy saving score	−0.374 (0.268)	−0.182 (0.190)	−0.765 (0.471)	−0.016 (0.090)
Green Party supporter	−1.049** (0.454)	−0.841** (0.314)	−2.072*** (0.772)	0.094 (0.152)
F quota	−3.978*** (0.477)			
HP quota	−3.252*** (0.422)			
LP quota	−4.810*** (0.451)			
F quota × rejector		1.644*** (0.535)	1.853*** (0.640)	
HP quota × rejector		4.078*** (0.803)	4.979*** (0.982)	
LP quota × rejector		2.081*** (0.598)	2.613*** (0.701)	
F quota × acceptor		−8.525*** (0.335)		−0.177 (0.125)
HP quota × acceptor		−5.097*** (0.321)		3.179*** (0.121)
LP quota × acceptor		−8.344*** (0.321)		
Completion time	−0.003*** (0.001)	−0.002*** (0.001)	−0.009*** (0.002)	−0.001 (0.001)
Control variables	yes	yes	yes	yes
Round-capacity fixed effects	yes	yes	yes	yes
Observations	1600[§,†]	1600[§,†]	822[§]	778[†]

[§]277 right-censored observations.

[†]3 left-censored observations.

One, two, and three asterisks denote significance at the 10, 5, and 1 percent levels, respectively.

potential rejectors, thereby lowering baseline mean demand. Second, rejectors may have tried to reap some advantage from acceptors' decision to reduce their maximum possible demand. To see this note that in line with the argument in the proof of Proposition 2, the possible spare capacity resulting from the introduction of a quota reduces the likelihood of causing an outage when making a high demand.

In Appendix D we use coarsened exact matching (CEM) to mitigate self-selection bias. The treated and control units here consist of rejectors and subjects in the baseline condition, respectively. The CEM-weighted estimates seem to confirm that the decision to reject a quota tends to increase demand relative to a counterfactual in which no quota is offered. Although this finding should be taken cautiously (see the discussion in the appendix), it provides an argument against the disclosure of information about quota acceptance rates. If an energy provider chose to make this information public, a potentially non-negligible fraction of users would consume more than they would otherwise, therefore making the quota scheme less effective.

5.3 Individual and group outcomes

We are now well-positioned to examine the effect of voluntary quotas on aggregate demand and consumption. To this end, since our observations are independent by design, we constructed ten-person simulated groups for each treatment condition (as in e.g. Niederle et al. 2013 and Buso et al. 2021). Each group is a unique random combination of subjects sampled without replacement. Given the high number of possible combinations (which is in the order of 10^{16}), for computational convenience we randomly selected 20,000 different groups per treatment, that is 10,000 different combinations of 10 subjects per subtreatment. By subtreatment we mean an element of $\{$Baseline, F, HP, LP$\} \times \{$Shortage, No shortage$\}$. For example, subjects in Subtreatment (HP, Shortage) were offered the high proportional quota and experienced a capacity shortage in Round 2. To achieve balance in covariates across all treatments, observations were post-stratified as described in Appendix A.1.

Table 4 lists summary results for group demands and payoffs, the fraction of groups that experienced an outage (henceforth outage rate), and payoff inequality. The table is complemented by Figures 5 and 6, which plot kernel density estimates of individual (top rows) and group (bottom rows) demands and payoffs. The first, second, and third columns of the figures present the results for Round 1, Round 2_{100}, and Round 2_{50}, respectively. Vertical dashed lines denote capacity.

We first consider the untreated baseline. Mean group demand ranged between 132 and 135 under no-shortage conditions, and decreased to just below 100 when capacity was halved. The density curves of individual demand have maxima at 5, 10, 15, and 20. The fraction of subjects who demanded exactly one tenth of capacity was 37.2 percent in Round 1 and 35.3 percent in Round 2, while the fractions of those

Table 4: Group outcomes

	Acceptance rate (%)	Round	Group demand	Outage rate (%)	Group payoff	Gini coefficient
Baseline		1	136.05±15.78	98.85	1.10±10.24	0.99
		2_{100}	132.17±15.33	98.22	1.69±12.59	0.99
		2_{50}	98.41±16.85	99.99	0.01±0.48	0.99
F quota	53.83	1	93.78±18.59	34.83	60.55±32.44	0.48
		2_{100}	95.60±18.25	37.82	58.90±33.93	0.49
		2_{50}	74.94±16.09	94.45	25.05±8.22	0.50
HP quota	76.91	1	105.31±12.90	62.16	75.34±16.97	0.28
		2_{100}	109.56±13.12	74.56	71.65±16.79	0.31
		2_{50}	62.68±13.46	77.10	38.07±6.57	0.24
LP quota	64.05	1	90.65±17.38	27.75	66.77±27.22	0.40
		2_{100}	86.98±16.36	20.18	70.20±24.07	0.36
		2_{50}	57.60±16.20	64.01	24.10±13.69	0.48

who demanded strictly more than one tenth were 52.3 and 55.6 percent, respectively. Consistent with the prediction in Section 4, the group demand curves in Figure 5 lie almost entirely to the right of their respective capacity line, that is, outages were the rule rather than the exception. Only about one percent of groups did not experience an outage when capacity was 100, and virtually no group succeeded in doing so when capacity dropped to 50. As a result, the individual and group payoff curves in Figure 6 skyrocket upwards as payoff approaches zero.

The above results serve as a first benchmark for assessing quotas. In Treatment F, the quota made mean group demand remain below the capacity threshold both in Round 1 and in Round 2_{100}, resulting in an outage rate of 34.8 and 37.8 percent, respectively. The distribution of group payoffs under no-shortage conditions is bimodal: 25.6 percent of groups earned an aggregate payoff between 15 and 25 (in which case capacity was typically exceeded and only acceptors had their demand met), while 55.6 percent earned a payoff of 70 or more (in which case individual demands summed to less than capacity). However, due to the high number of rejectors with an inelastic demand, in the event of a capacity shortage the outage rate rose to a lacklustre 94.4 percent, which makes the F quota the least effective in preventing outages. The share of individuals who received a payoff of zero ranged from 20.7 percent in Round 1 to 41.1 percent in Round 2_{50}.

The HP quota was the most effective in sustaining consumption and payoffs, which at the individual level averaged more than 7 when capacity was 100 and 3.8 when capacity was 50. The proportion of subjects who consumed no energy was 17.8 percent in Round 1, 23.0 percent in Round 2_{100}, and 17.3 percent in Round 2_{50}. These results were obtained at the expense of demand reduction and grid stress alleviation, as mean group demand exceeded capacity both under no-shortage

Figure 5: Individual and group demands (kernel density estimates)

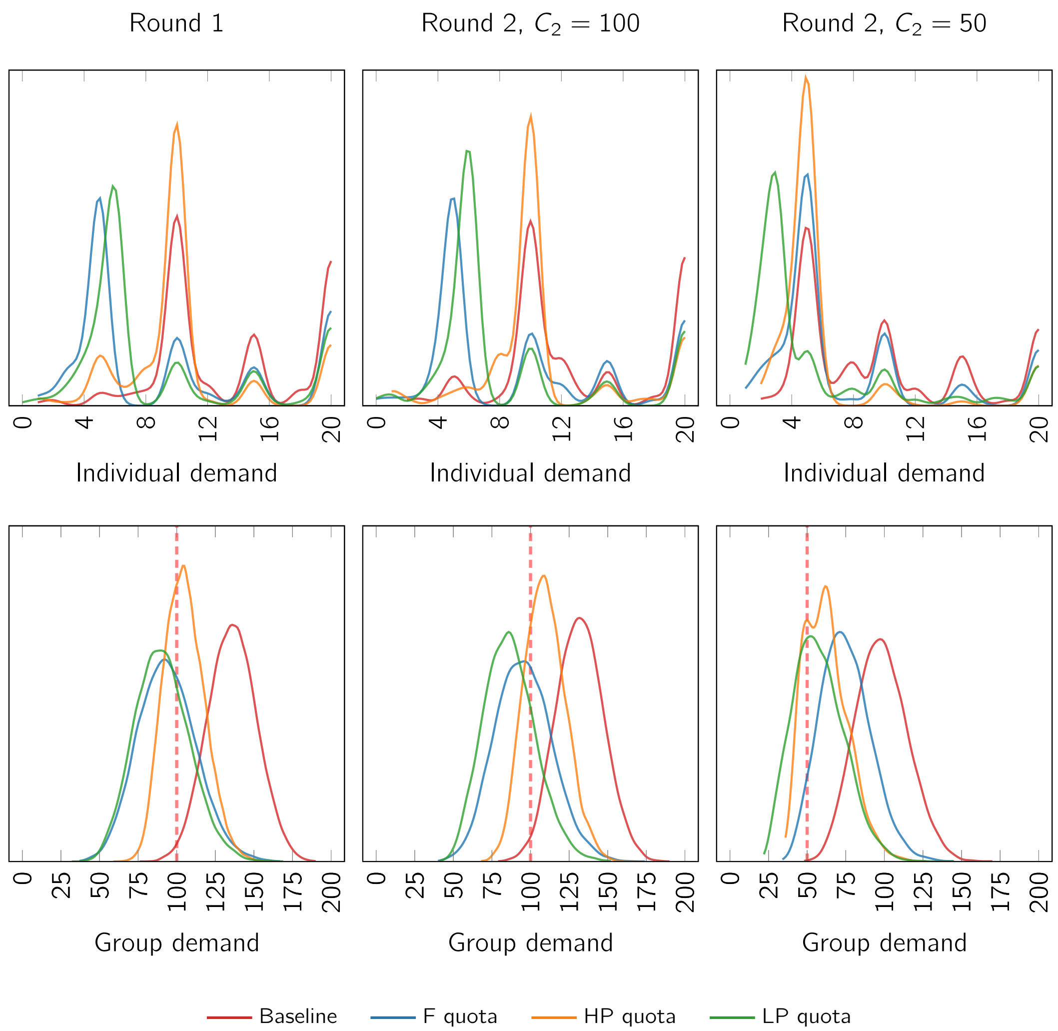

and under shortage conditions. The outage rate thus remained considerably high, ranging between 62.2 and 77.1 percent.

Finally, The LP quota performed best in terms of outage prevention. Due to the tight constraint on demand and the relatively high number of acceptors, mean group demand remained below capacity both in Round 1 and in Round 2_{100}. In the latter case the outage rate reached its minimum of 20 percent. Under shortage conditions, group demand exceeded capacity by some 8 Energy Units and the outage rate rose to 64 percent, which is rather high but more than 35 percentage point lower than in the baseline (and roughly 13 percentage point lower than in Treatment HP). The proportion of individuals who consumed no energy was 14.8 percent in Round 1, 11.1 percent in Round 2_{100}, and 28.6 percent in Round 2_{50}. It is also worth mentioning that the LP quota outperformed the F quota in terms of both group payoff and outage rate under no-shortage conditions, and in terms of only the outage rate

Figure 6: Individual and group payoffs (kernel density estimates)

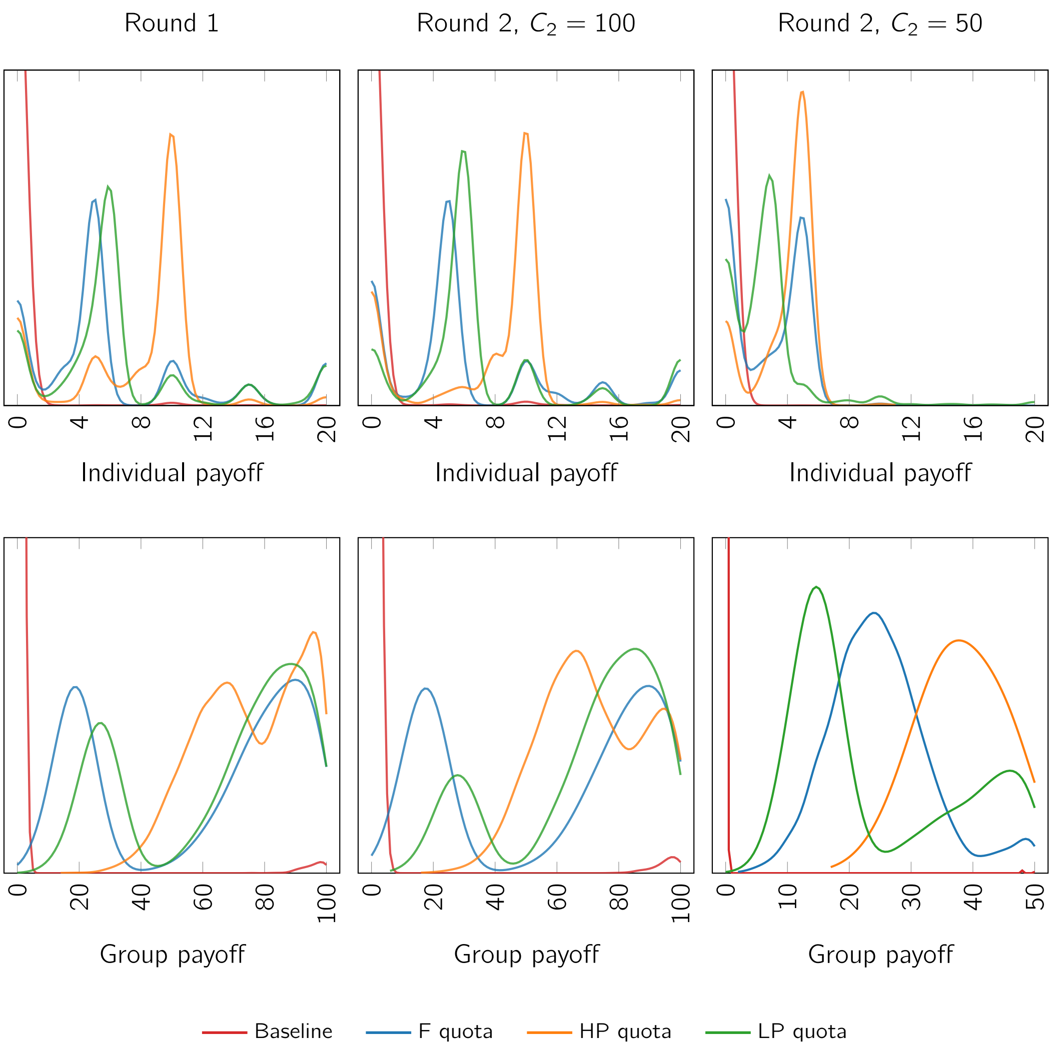

under shortage conditions (since the mean group payoff in Round 2_{50} is slightly higher in Treatment F than in Treatment LP). A Mann-Whitney test found that although small, this difference in group payoffs is statistically significant (p-value < 0.001). To see why this is the case, note that while on the one hand Treatment LP had a lower number of rejectors who received a payoff of zero, on the other hand acceptors' maximum possible payoff in Round 2_{100} is lower with the LP quota than with the F quota. On average, the first effect outweighed the second.

The last column of Table 4 reports the Gini coefficient of payoff inequality, calculated using all simulated group observations. The baseline values of the coefficient are all close to unity, since only a handful of subjects received a positive payoff while others got nothing. The Gini calculated for the treatments instead ranges between 0.28 and 0.50, with lower values corresponding to a higher proportion of subjects who accepted the quota and secured a stable consumption level for themselves. Inequality is highest in Treatment F, due to the considerable number of rejectors who experienced one or more outages, and lowest in Treatment HP, which maximizes the quota acceptance rate.

6 Concluding remarks

This paper explores the potential of voluntary quotas as a means of curbing resource consumption while enhancing consumption security. To this end, we conducted an experiment with a Nash demand game in which subjects are given the option to limit their maximum consumption in exchange for being assigned a priority status and being certain to receive their demands. Individuals are under no obligation to accept the quota, and the decision to do so allows them to self-regulate their consumption to a sustainable level.

Our main findings provide support for the following. First, individuals are slow to adjust their consumption to sudden reductions in supply. For instance, in our baseline condition, when energy capacity was reduced by half, we observed an average decrease in demand of around one-third. This inelasticity highlights the potential benefits of implementing measures aimed at reducing demand at times of supply shortage. Second, although voluntary quotas seem not to be sufficient to prevent outages entirely, they might nevertheless play a valuable role in reducing aggregate consumption and grid stress. Our experimental results show that depending on the specific quota scheme implemented, group consumption and the outage frequency decreased by up to 30 and 80 percent compared to the baseline, respectively. Third, the choice of which type of quota to implement would largely depend on the energy provider's objective. For example, if the provider were mainly interested in sustaining consumption, then a quota of the HP type would likely be the most effective option. If, instead, the provider were keen about reducing the frequency of outages, then implementing a quota of the LP type may be more appropriate.

Several research questions remain. Here we outline some of them. An avenue for further research is to examine the influence of experience on quota acceptance and demand decisions. This could be achieved through an experiment comprising numerous rounds where participants receive information on past play and are periodically given the chance to reconsider their decision to accept or reject the quota. Another important question related to the discussion in Section 5.2 is whether and how the provision of public information about quota acceptance rates affects individual behaviour. This too may be investigated using an experiment specifically designed for this purpose. One may furthermore devise an incentive compatible menu of quotas that impose different restrictions on demand and bestow users with different priority statuses. Finally, a field experiment could be used to seek further support for the external validity of our results. Work along each of these lines would certainly pave the way to a better understanding of the effects of voluntary quotas.

Replication files

The preregistration document and the data and code for replicating the results of this paper are available at https://osf.io/6rvxz/?view_only=9e5b558380224e299f604603b9ca26c1. All files are licensed under a Creative Commons Attribution 4.0 International (CC BY 4.0) license.

References

Aflaki, Sam (2013), "The effect of environmental uncertainty on the tragedy of the commons", *Games and Economic Behavior*, vol. 82, pp. 240-253.

Alekseev, Aleksandr, Gary Charness, and Uri Gneezy (2017), "Experimental methods: When and why contextual instructions are important", *Journal of Economic Behavior & Organization*, vol. 134, pp. 48-59.

Allcott, Hunt, Allan Collard-Wexler, and Stephen D. O'Connell (2016), "How do electricity shortages affect industry? Evidence from India", *American Economic Review*, vol. 106, no. 3, pp. 587-624.

Alpizar, Francisco and Elisabeth Gsottbauer (2015), "Reputation and household recycling practices: field experiments in Costa Rica", *Ecological Economics*, vol. 120, pp. 366-375.

Alvarez, Fernando and David Argente (2022), "On the Effects of the Availability of Means of Payments: The Case of Uber", *Quarterly Journal of Economics*, vol. 137, no. 3, pp. 1737-1789.

Anderies, John M., Marco A. Janssen, Allen Lee, and Hannah Wasserman (2013), "Environmental variability and collective action: Experimental insights from an irrigation game", *Ecological Economics*, vol. 93, pp. 166-176.

Baldick, Ross, Sergey Kolos, and Stathis Tompaidis (2006), "Interruptible electricity contracts from an electricity retailer's point of view: Valuation and optimal interruption", *Operations Research*, vol. 54, no. 4, pp. 627-642.

Berger, Sebastian and Annika M Wyss (2021), "Measuring pro-environmental behavior using the carbon emission task", *Journal of Environmental Psychology*, vol. 75, p. 101613.

Bernard, Mark, Anna Dreber, Pontus Strimling, and Kimmo Eriksson (2013), "The subgroup problem: When can binding voting on extractions from a common pool resource overcome the tragedy of the commons?", *Journal of economic behavior & organization*, vol. 91, pp. 122-130.

Bochet, Olivier, Jeremy Laurent-Lucchetti, Justin Leroux, and Bernard Sinclair-Desgagné (2019), "Collective risk-taking in the commons", *Journal of Economic Behavior & Organization*, vol. 163, pp. 277-296.

Buso, Irene M., Daniela Di Cagno, Werner Güth, and Lorenzo Spadoni (2021), *Voluntary Partnerships for Equally Sharing Contribution Costs: Theoretical Aspects and Experimental Evidence*, CESARE working paper 1/2021, LUISS Guido Carli.

Calzolari, Giacomo, Marco Casari, and Riccardo Ghidoni (2018), "Carbon is forever: A climate change experiment on cooperation", *Journal of Environmental Economics and Management*, vol. 92, pp. 169-184.

Cardenas, Juan C., John Stranlund, and Cleve Willis (2000), "Local environmental control and institutional crowding-out", *World Development*, vol. 28, no. 10, pp. 1719-1733.

Chen, Daniel L., Martin Schonger, and Chris Wickens (2016), "oTree—An open-source platform for laboratory, online, and field experiments", *Journal of Behavioral and Experimental Finance*, vol. 9, pp. 88-97.

Cohn, Alain, Tobias Gesche, and Michel André Maréchal (2022), "Honesty in the digital age", *Management Science*, vol. 68, no. 2, pp. 827-845.

Cole, Matthew A., Robert J.R. Elliott, Giovanni Occhiali, and Eric Strobl (2018), "Power outages and firm performance in Sub-Saharan Africa", *Journal of Development Economics*, vol. 134, pp. 150-159.

Czap, Natalia V., Hans J. Czap, Marianna Khachaturyan, and Mark E Burbach (2018), "Comparing female and male response to financial incentives and empathy nudging in an environmental context", *Review of Behavioral Economics*, vol. 5, no. 1, pp. 61-84.

Dohmen, Thomas, Armin Falk, David Huffman, Uwe Sunde, Jürgen Schupp, and Gert G. Wagner (2011), "Individual risk attitudes: Measurement, determinants, and behavioral consequences", *Journal of the European Economic Association*, vol. 9, no. 3, pp. 522-550.

European Commission (2022), *Save Gas for a Safe Winter: Commission proposes gas demand reduction plan to prepare EU for supply cuts*, press release, https://ec.europa.eu/commission/presscorner/detail/en/ip_22_4608 (visited on 09/30/2022).

Falk, Armin, Anke Becker, Thomas Dohmen, David Huffman, and Uwe Sunde (2023), "The preference survey module: A validated instrument for measuring risk, time, and social preferences", *Management Science*, vol. 69, no. 4, pp. 1935-1950.

Farjam, Mike, Olexandr Nikolaychuk, and Giangiacomo Bravo (2019), "Experimental evidence of an environmental attitude-behavior gap in high-cost situations", *Ecological Economics*, vol. 166, p. 106434.

Iacus, Stefano M., Gary King, and Giuseppe Porro (2011), "Multivariate Matching Methods That Are Monotonic Imbalance Bounding", *Journal of the American Statistical Association*, vol. 106, no. 493, pp. 345-361.

— (2012), "Causal Inference without Balance Checking: Coarsened Exact Matching", *Political Analysis*, vol. 20, no. 1, pp. 1-24.

IEA (2021), *World Energy Outlook 2021*, International Energy Agency, https://www.iea.org/reports/world-energy-outlook-2021 (visited on 01/31/2023).

Lawson, Alex (2022), "Don't mention rationing: Why energy crisis may need another Truss U-turn", *The Guardian*, https://www.theguardian.com/business/2022/sep/09/rationing-energy-crisis-liz-truss-blackouts-usage (visited on 09/30/2022).

Longden, Thomas, Simon Quilty, Brad Riley, Lee V. White, Michael Klerck, Vanessa Napaltjari Davis, and Norman F. Jupurrurla (2022), "Energy insecurity during temperature extremes in remote Australia", *Nature Energy*, vol. 7, no. 1, pp. 43-54.

Lyons, Elizabeth and Laurina Zhang (2017), "The Impact of Entrepreneurship Programs on Minorities", *American Economic Review*, vol. 107, no. 5, pp. 303-307.

Maurer, Luiz, Mario Pereira, and José Rosenblatt (2005), *Implementing Power Rationing in a Sensible Way: Lessons Learned and International Best Practices*, Energy Sector Management Assistance Program (ESMAP) Technical Paper 305/05, World Bank.

Milinski, Manfred, Ralf D Sommerfeld, Hans-Jürgen Krambeck, Floyd A Reed, and Jochem Marotzke (2008), "The collective-risk social dilemma and the prevention of simulated dangerous climate change", *Proceedings of the National Academy of Sciences*, vol. 105, no. 7, pp. 2291-2294.

Niederle, Muriel, Carmit Segal, and Lise Vesterlund (2013), "How Costly Is Diversity? Affirmative Action in Light of Gender Differences in Competitiveness", *Management Science*, vol. 59, no. 1, pp. 1-16.

Non, Arjan, Ingrid Rohde, Andries de Grip, and Thomas Dohmen (2022), "Mission of the company, prosocial attitudes and job preferences: A discrete choice experiment", *Labour Economics*, vol. 74, no. 102087.

O'Neill, Daniel W., Andrew L. Fanning, William F. Lamb, and Julia K. Steinberger (2018), "A good life for all within planetary boundaries", *Nature Sustainability*, vol. 1, no. 2, pp. 88-95.

Ostrom, Elinor (1999), “Coping with tragedies of the commons”, *Annual Review of Political Science*, vol. 2, pp. 493-535.

Palan, Stefan and Christian Schitter (2018), “Prolific.ac—A subject pool for online experiments”, *Journal of Behavioral and Experimental Finance*, vol. 17, pp. 22-27.

Princen, Thomas (2005), *The Logic of Sufficiency*, MIT Press, Cambridge (MA).

Ritchie, Hannah, Max Roser, and Pablo Rosado (2022), *Energy*, Our World in Data, https://ourworldindata.org/energy-production-consumption (visited on 01/31/2023).

Tavoni, Alessandro, Astrid Dannenberg, Giorgos Kallis, and Andreas Löschel (2011), “Inequality, communication, and the avoidance of disastrous climate change in a public goods game”, *Proceedings of the National Academy of Sciences*, vol. 108, no. 29, pp. 11825-11829.

Thaler, Richard H. and Cass R. Sunstein (2003), “Libertarian paternalism”, *American Economic Review*, vol. 93, no. 2, pp. 175-179.

Tranter, Bruce and Kate Booth (2015), “Scepticism in a changing climate: A cross-national study”, *Global Environmental Change*, vol. 33, pp. 154-164.

Volland, Benjamin (2017), “The role of risk and trust attitudes in explaining residential energy demand: Evidence from the United Kingdom”, *Ecological Economics*, vol. 132, pp. 14-30.

Appendices

A Sample characteristics

Table A1 summarizes the characteristics of our sample. Small differences in the number of observations per treatment are due to the algorithm used by Prolific to replace dropouts. Energy saving behaviour, socio-economic status, risk tolerance, and altruism were assessed as described in the main text. The last column reports the p-values of tests for differences across the four treatment groups. Group comparisons for continuous and binary variables were performed using the Kruskal-Wallis test and Fisher's exact test, respectively. The null hypotheses are never rejected, meaning that characteristics tend to be balanced across treatments.

Table A1: Sample characteristics

	Baseline	Treatment F	Treatment HP	Treatment LP	Kruskal-Wallis or Fisher's exact test (p-value)
Age	32.36±7.26	32.25±6.88	31.62±6.84	32.19±7.04	0.654
Females (%)	45.73	50.99	49.25	54.04	0.413
Ethnic minority members (%)	12.56	17.33	16.42	18.18	0.425
University graduates (%)	67.84	68.32	71.14	71.72	0.782
Unemployed	6.53	6.93	5.47	7.07	0.912
Believe in climate change (%)	91.46	91.09	95.52	93.94	0.243
Green Party supporters (%)	11.56	12.38	11.44	12.12	0.991
Subjective socio-economic status	5.26±1.47	5.10±1.53	5.23±1.47	5.29±1.62	0.676
Energy sav ing score	1.97±0.57	1.94±0.54	2.04±0.61	2.05±0.61	0.310
Risk tolerance	5.43±2.16	5.46±2.19	5.26±2.36	5.41±2.23	0.923
Altruism	1.09±1.60	1.14±1.57	0.92±1.43	1.09±1.35	0.227
Completion time (min.)	3.10±1.70	3.90±2.79	3.98±2.67	3.85±1.95	
Observations	199	202	201	198	

The most notable between-group differences concern Sex and Climate skepticism, as both variables are significant predictors of quota acceptance (see Section 5.1). The share of females is lowest in the untreated baseline and highest in Treatment 3, while the share of climate skeptics is lowest in Treatment HP and highest in Treatment F. In order to maximize the power to detect a difference, we compared these group pairs using a one-sided Fisher's exact test without multiple-testing corrections. The difference in sex composition between the baseline and Treament LP is significant at the 10 percent level (p-value = 0.060), and so is the difference in climate beliefs between Treatments HP and F (p-value = 0.056). In Section 5.3 we use post-stratification to minimize the confounding effect of imbalances in these two variables. The stratification procedure is outlined below.

A.1 Post-stratified random group formation

Post-stratification involves assigning sampling weights to observations according to the joint distribution of some variables. The simulated groups used for analysis in Section 5.3 were formed so that the joint distribution of Sex and Climate skepticism in each subtreatment matches the distribution in the whole sample. Post-stratification weights were thus constructed based on four variable combinations: non-skeptic females, skeptic males, and so on. As an illustrative example, let w be the fraction of climate skeptic males in the overall sample, let w_s be the fraction of skeptic males in Subtreatment s, and let n_s be the number of observations in Subtreatment s. The sampling weight assigned to each skeptic male in this subtreatment is the ratio w/w_s, which implies that each has a probability of $(10w)/(n_s w_s)$ to be in each of the 10,000 combinations drawn for Subtreatment s.

Table A2: Sample characteristics after post-stratified random group formation

	Baseline	Treatment F	Treatment HP	Treatment LP
Age	32.21±7.22	32.24±6.80	31.73±6.90	32.42±7.05
Females (%)	50.09	50.07	50.09	48.69
Ethnic minority members (%)	12.03	16.95	15.72	18.62
University graduates (%)	67.77	68.31	70.94	72.43
Unemployed	6.56	6.80	5.37	6.59
Believe in climate change (%)	92.95	92.98	92.98	93.02
Green Party supporters (%)	11.70	12.63	11.08	11.87
Subjective socio-economic status	5.24±1.48	5.12±1.53	5.23±1.51	5.30±1.63
Energy saving score	1.97±0.57	1.94±0.54	2.03±0.61	2.04±0.61
Risk tolerance	5.37±2.16	5.45±2.17	5.26±2.36	5.48±2.22
Altruism	1.10±1.59	1.12±1.52	0.91±1.42	1.10±1.37
Completion time (min.)	3.11±1.69	3.90±2.81	3.97±2.63	3.86±1.92
No. of groups[§]	20,000	20,000	20,000	20,000

[§]Each simulated group is a different random combination of 10 subjects, sampled without replacement and with weights calculated as described in the text. The values in this table are therefore based on 200,000 data units per treatment.

The characteristics of post-stratified random groups are reported in Table A2. The fraction of females now ranges between 48.69 and 50.09 percent, while the fraction of subjects who believe in climate ranges between 92.95 and 93.02 percent.

Importantly, the stratification procedure did not cause any substantial change in the distribution of other variables. For instance, the fraction of Green Party supporters ranges between 11.08 and 12.63 percent, while mean Risk tolerance and Altruism range between 5.26 and 5.48 and between 0.91 and 1.12, respectively.

B Mean acceptance probabilities

Figure B1: Mean acceptance probabilities by treatment, Risk tolerance, and Altruism

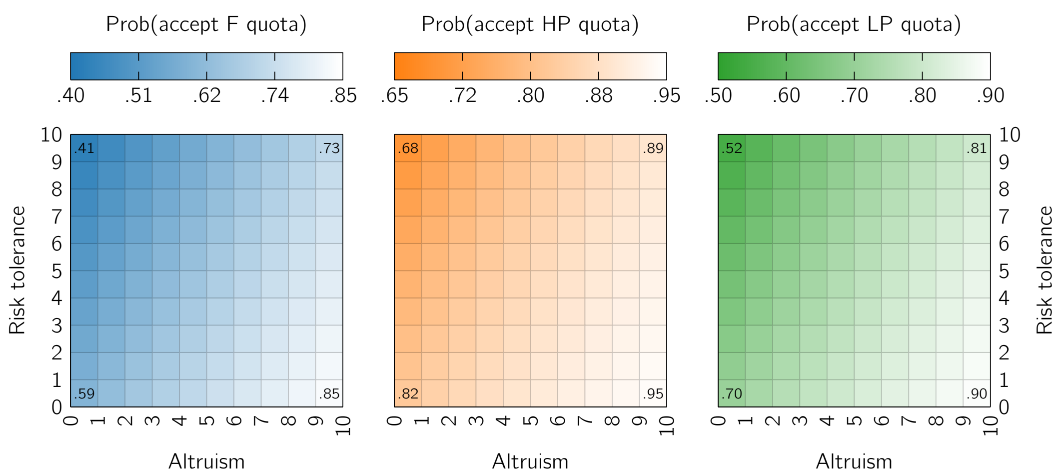

Figure B2: Differences in mean acceptance probabilities (HP quota vs. F quota)

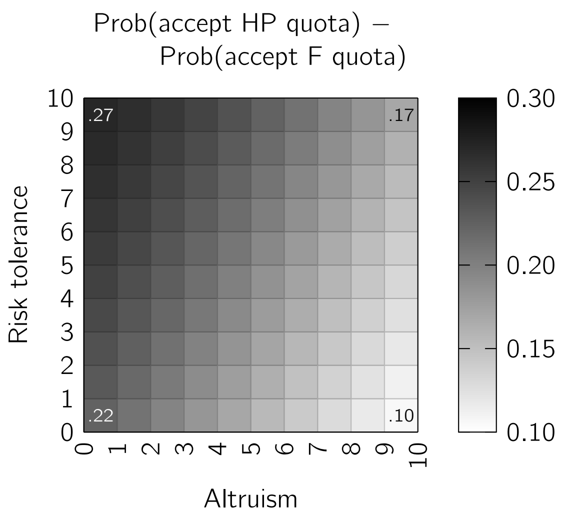

Figure B1 shows mean estimated acceptance probabilities by treatment at different values of Risk tolerance and Altruism (while holding all other variables constant). Estimated probabilities are based on the logistic regression reported in the third column of Table 2. Brighter and darker colours represent higher and lower probabilities of acceptance, respectively. Predictions at extreme values are given at the four corners of each heatmap; for instance, the estimated probability of accepting the F quota when Risk tolerance = 10 and Altruism = 0 is approximately 41 percent. In Figure B2 we show that consistent with our discussion in Section 5.1, the difference between the probabilities of accepting the HP and F quotas increases with Risk tolerance and decreases with Altruism.

C Changes in individual demand

Figure C1 plots the distributions of demand changes between Round 2 and Round 1 for subjects who did not experience a capacity shortage. The horizontal axis measures the difference between $d_{i,2_{100}}$ and $d_{i,1}$.

Figure C1: Changes in individual demand by treatment and acceptance or rejection choice

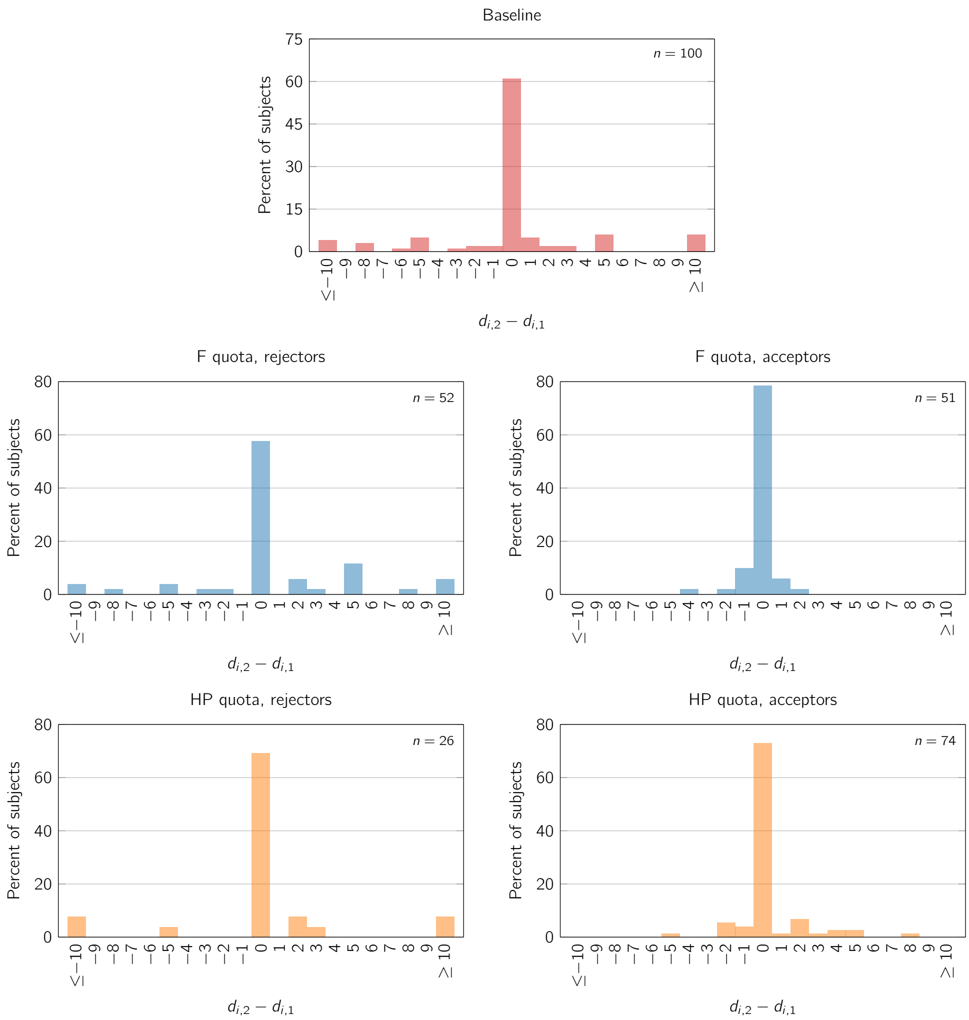

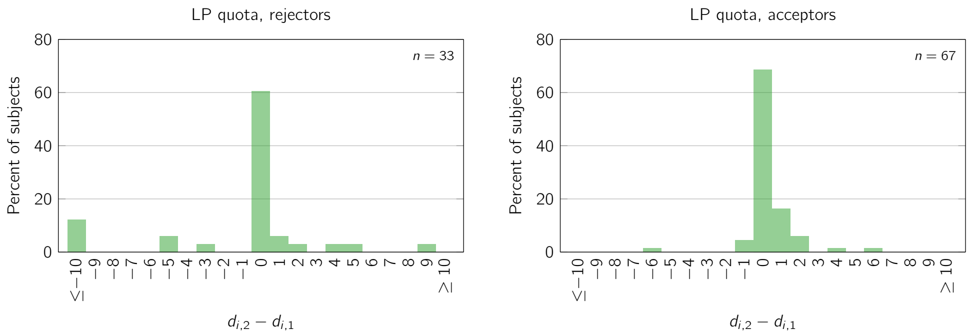

For each treatment and acceptance or rejection choice, we performed a non-parametric sign test of the null hypothesis that the median value of $d_{i,2_{100}} - d_{i,1}$ is zero. As shown in Table C1, the null is rejected only for acceptors in Treatment LP. However, the distribution of demand changes is in this case so tightly concentrated around zero that it would be difficult to explain it as due to hedging. Moreover, since acceptors can obtain a payoff of $\underline{d}$ with certainty, they have no reason to hedge their demands. Overall, we interpret these results as indicating that subjects tended not to exploit hedging opportunities.

Table C1: Sign test p-values

	p-value
Baseline	0.7493
F quota, rejectors	0.1338
F quota, acceptors	0.5488
HP quota, rejectors	0.7266
HP quota, acceptors	0.5034
LP quota, rejectors	0.9999
LP quota, acceptors	0.0072

Null hypothesis: the median value of $d_{i,2_{100}} - d_{i,1}$ is zero.

D Coarsened exact matching

To assess the effect on demand of self-selection in the rejectors group, we used coarsened exact matching to match subjects in the baseline condition to rejectors in Treatments F, HP and LP. CEM reduces bias by bounding the maximum imbalance in the empirical distributions of pretreatment variables between the treated (rejectors) and control (baseline) groups (Iacus et al. 2011, 2012; for applications see e.g. Lyons and Zhang 2017 and Alvarez and Argente 2022). The idea is that matched units in the control group would have likely rejected a quota if it had been offered to them, and can therefore be used as counterfactuals. The CEM procedure consists of the following steps. First, each pretreatment variable is temporarily coarsened into bins for matching purposes (for example, we coarsened education data into tertiary and non-tertiary education). Second, observations are sorted into mutually exclusive strata, each of which has the same values of all coarsened variables. Third, all observations within any stratum which have no matches on pretreatment variables in both the treated and the control groups are pruned from the data set. Finally, the original uncoarsened values of the matched data are used to estimate the sample average treatment effect on the treated, with weights that adjust for differences in the numbers of treated and control units within each stratum. Stratum weights are defined as 1 for treated units and $(m_C/m_T)(m_T^\psi/m_C^\psi)$ for control units, where m_T and m_C are the numbers of treated and control units in the sample, and m_T^ψ and m_C^ψ are the numbers of treated and control units in stratum ψ.

Our set of pretreatment variables consists of the covariates which were found to be significant in the first regression of Table 3, namely Altruism, Climate skepticism, Education, Green Party support, Risk tolerance, and Sex. In the case of Altruism and Risk tolerance, bin sizes were determined using the Sturges binning algorithm (Iacus et al. 2012). Approximately 74 percent of observations were successfully matched using this procedure.

Table D1: Effect of quota rejection on demand (CEM-weighted estimates)

	(1) OLS	(2) Tobit	(3) Tobit
Rejector	1.822*** (0.476)	2.538*** (0.712)	2.454*** (0.676)
Round-capacity fixed effects	no	no	yes
Observations	610	610[§]	610[§]

[§]219 right-censored observations.

One, two, and three asterisks denote significance at the 10, 5, and 1 percent levels, respectively.

Table D1 presents pooled estimates of the sample average effect of rejecting a quota. The dependent variable is individual demand, while the variable Rejector is a dummy that equals 1 if a subject was offered a quota and rejected it. All models were weighted based on CEM results. We stress that these findings lie on the untestable assumption that, conditional on matched observables, the reason why an observation belongs to the treated or control groups is not due to variables which are correlated with demand. As such they should not be taken as definitive evidence of a causal effect. In column 1 we report a simple difference in means, while in columns 2 and 3 we report Tobit estimates with and without round-capacity fixed effects. The coefficient on Rejector is always significant and suggests that on average rejecting the quota led to an increase in demand by about 2 Energy Units.