

ARTICLE

Laypeople have difficulty processing efficiency when assessing environmental policies

Antoine Marie^{1,2} , Hugo Trad^{1,3}  and Brent Strickland^{1,3}

¹Institut Jean Nicod, Département d'Études Cognitives, ENS, EHESS, CNRS, PSL Research University, Paris, France; ²Center for Political Research (CEVIPOF), Sciences Po, Paris, France and ³Africa Business School, The School of Collective Intelligence, UM6P, Rabat, Morocco

Corresponding author: Antoine Marie; Email: antoine.marie.sci@gmail.com

(Received 31 January 2025; revised 28 June 2025; accepted 18 September 2025)

Abstract

What should mostly matter is how successful environmental policies are at satisfying citizens' policy preferences (e.g., reducing carbon emissions), relative to the policies' cost. Yet, across 6 studies ($N = 2759$, 2 pre-registered), we found that French citizens tended to be rather insensitive to policy efficiency. In Experiments 1a–d ($N = 854$), citizens regarded an environmental policy driven by an altruistic intention that turned out to be inefficient as being more commendable than a policy motivated by selfishness that dramatically reduced carbon emissions. In Experiment 2 ($N = 1105$), altruistic but low efficiency policies were supported only slightly less than selfish but highly efficient policies. Independent manipulation of intent and efficiency indicated low sensitivity to large differences in efficiency expressed numerically, and substantial sensitivity to actors' intentions. Moreover, moral commitment predicted stronger support for any environmental policy addressing the issue, regardless of its efficiency. Finally, Experiment 3 ($N = 800$) found that introducing reference points and qualitative appraisals of a policy's impact and financial cost can nudge participants towards greater attention to its efficiency. Our paper highlights the importance of using contextual and qualitative (vs. numeric) descriptions of policies to make citizens more focused on their efficiency.

Keywords: consequences; efficiency; environment; morality; public policy

Would you prefer a public policy that is highly efficient in taming CO₂ emissions but is motivated by profit? Or a policy that fares quite poorly in curbing emissions and at a high cost for society, but altruistically aims to reduce carbon emissions? This paper suggests that many laypeople might support the latter over the former option.

In modern societies, high amounts of executive power can be concentrated in the hands of individuals such as ministers and CEOs (Chief Executive Officers). This

power enables them to make policy decisions that can contribute to address pressing issues such as climate change, sustaining economic growth, or ensuring national defense, but which can also result in considerable *wasted* money and opportunities when the decisions are misguided. A fundamental principle in democratic governance and corporate management is that decision-making authority stems from a delegation of power, whether from citizens to government officials or from shareholders to corporate executives (Jensen and Meckling, 1976; Strøm, 2000). This delegation is accompanied by oversight mechanisms designed to align the decisions of those in power with the interests of their principals (i.e., citizens, stakeholders and workers). In the public sphere, oversight mechanisms include electoral accountability and legislative oversight (McCubbins *et al.*, 1987), while in the corporate world, they encompass board supervision, shareholder voting rights and market discipline (Fama and Jensen, 1998). From this normative standpoint, it can be argued that what should mostly matter is how efficient policies are at reaching citizens' policy preferences (e.g., as expressed in votes), while keeping in mind that policies' deployment can confiscate resources that could be allocated to other issues. Specifically, whether the ministers or CEOs who implement a given policy had noble or altruistic *intentions* or tried to benefit personally from implementing the policy should have secondary importance.¹ Abilities to judge public policies based on their concrete *results* seem particularly important in market societies, in which corporate innovation motivated by the search for profits can bring about positive externalities (Gilbert, 2006; Hall and Rosenberg, 2010; Smith, 2013).

Contrary to this normative ideal, however, cost-benefit analyses often seem counterintuitive to lay policy reasoning. Other considerations, such as inferences about 'altruistic' vs 'selfish' intentions behind the policies, appear powerful in influencing judgments. In the domain of charity giving, donors are relatively insensitive to the extent to which charities concretely increase human welfare. Donors' attention is instead focused on whether the charities pursue their preferred causes (Berman *et al.*, 2018), and donors strongly underestimate potentially large differences in effectiveness between charities (Caviola *et al.*, 2020). In laboratory experiments, cute, isolated children often attract more empathy and donations than larger communities of poor people in developing nations (Bloom, 2017). Likewise, the explicit belief that the selfish pursuit of monetary interest on the market is morally 'bad' and undermines the common good is widespread across cultures. It leads many citizens to oppose Adam Smith's notion that private self-regard can contribute to making society prosperous and efficient, often in contradiction with citizens' own and higher goods quality (Smith, 2013; Rubin, 2014; Boyer and Petersen, 2018). Many people have the folk economic intuition

¹Whenever strangers, friends, neighbors, or colleagues are in a position to deliver us benefits or to inflict costs on us, especially over repeated interactions, monitoring their intentions and values is adaptive as shortcuts to future cooperative prospects (Baumard *et al.*, 2013; Uhlmann *et al.*, 2015). What we are arguing here is that in modern mass societies, the concentration of power and the existence of large chains of command put powerful leaders – with whom one will never interact – in a position to potentially solve pressing societal issues, but also to waste precious public resources, through one-off decisions. By standards of policy efficiency and democratic sovereignty, it could be argued that the nobility of the motivations driving policymakers' when they implement a given policy should be secondary to the policy's ability to satisfy citizens' needs.

that doing well is somehow antithetical to doing good (to others) (Lin-Healy and Small, 2013).

Attitudes toward technologies are sometimes similarly marked by relative inattention to impact and opportunity costs, and a focus on intentions. Genetically modified organisms (GMOs) are widely seen by scientists as safe and useful in increasing agricultural yields and fighting against world hunger (AAAS, 2013; Blancke *et al.*, 2015; Medani *et al.*, 2024). Yet, those facts are rarely featured in folk reasoning about GMOs. Food GMOs tend to face substantial popular opposition (S. E. Scott *et al.*, 2016; Irsn, 2017) anchored in essentialist intuitions that scientists ‘shouldn’t play God’ (Blancke *et al.*, 2015; S. E. Scott *et al.*, 2016), and in beliefs that the profit motive driving their development makes them look suspicious (Bonny, 2003).

In this paper, we set out to explore another possible manifestation of efficiency neglect by examining how French laypeople factor in information about the *efficiency* of pro-environment public policies relative to cues about the *intentions* driving their implementation. We suspected that laypeople spontaneously give little weight to large differences in efficiency between public policies, potentially supporting policy decisions that are ineffective and wasteful if they seem driven by good intentions. This tendency appears problematic to the extent that it is desirable for citizens to try to convert their own policy preferences into tangible policy programs that satisfy those preferences (McCubbins *et al.*, 1987; Fama and Jensen, 1998; Strøm, 2000; Herzog and Hertwig, 2025).

As regards the scope of our investigation, we focus on environmental policies of emissions reduction. Addressing anthropogenic climate change by reducing carbon dioxide emissions is one of the most pressing policy and technological problems faced by modern societies. For that reason, we frame our studies around evaluations of policies aimed at reducing the CO₂ emissions of the French industry. Our approach consists of presenting French respondents with scenarios in which policymakers (ministers or CEOs) are portrayed as trying to reduce CO₂ emissions through the deployment of carbon capture policies that achieve high vs low efficiency, and are driven by ‘altruistic’ vs ‘selfish’ intentions in the policymakers who implement them.

To foreshadow, we observe that lay French participants pay little attention to the efficiency of environmental policies when it is described in numerical terms, and especially so when they highly moralize environmental protection (Experiments 1a–d and 2). We also find that lay respondents tend to reward altruistic intentions in policymakers, and to punish selfish intentions, in cases of one-off policy decisions where those intentions should arguably have little importance. As a result, respondents can end up supporting highly inefficient and costly policies seen as altruistically motivated almost to the same extent as highly efficient policies seen as motivated by personal profit (Experiments 1a–d and 2). Encouragingly, however, our final experiment (Experiment 3) finds that introducing reference points and qualitative appraisals of the policies’ effects and cost can help citizens judge them in more efficiency-focused ways than when efficiency information is given only numerically.

Overall, our experiments show that processing quantitative information about public policies’ efficiency is difficult for laypeople. In line with work on ‘boosting’ in behavioral policy research (see Herzog and Hertwig, 2025, for a review), we show the

importance of finding intuitive framings when conveying policy characteristics to the public.

Experiments 1a–d

Experiment 1a exposed participants to the descriptions of two policies deployed by a CEO and aiming to reduce the CO₂ emissions of his company through carbon capture and storage technology. Development of this technology is currently under growing market incentives (V. Scott *et al.*, 2013). It may deserve to be judged based on cost-benefit considerations about its ability to reduce emissions, rather than beliefs about the moral intentions underlying its deployment (Hall and Rosenberg, 2010). One policy was described as highly efficient in decreasing CO₂ emissions but selfish, the other as altruistic but quite inefficient. We asked how *morally commendable* respondents would see each environmental policy. People are moralistic animals, and their moral evaluations powerfully shape their voting and policy attitudes (Tetlock, 2002; Baumard *et al.*, 2013). We expected respondents to see the altruistic but low-efficiency policy as being more commendable than the selfish but high-efficiency policy. Experiments 1b and c then assessed the generality of our findings on three unrelated issues.

Carbon capture technology and the environment

Experiments 1a, 2 and 3 examine judgments of policy decisions aiming to reduce CO₂ emissions through the deployment of carbon capture and storage technologies. Everyone is by now familiar with the Intergovernmental Panel on Climate Change report (IPCC, Pörtner *et al.*, 2022) insisting on the urgency of reducing atmospheric greenhouse gas concentrations to keep global warming under 1.5°C. One pathway to doing this is to reduce emissions through decreased production. But the IPCC report also clearly asserts that technologies aiming at CO₂ capture, storage and re-use (i.e., carbon capture) can help in meeting global climate change objectives² (V. Scott *et al.*, 2013; Pörtner *et al.*, 2022). Currently, many of these technologies are at an early stage of development (V. Scott *et al.*, 2013). As such, their long-term future depends on both government support and committed entrepreneurs, who in turn rely on voter and consumer backing. Thus, understanding public attitudes³ toward carbon capture

²Carbon capture technologies consist of extracting the carbon dioxide responsible for climate change from the air using turbines, chemical reactions, or algae, to transform it into solid material to be stored durably in cement or deep geological formations (V. Scott *et al.*, 2013; Wikipedia, 2024).

³Opinion polling on the topic suggests that public opinion is divided in Western industrialized societies, even among populations who are committed to combating climate change (Haszeldine, 2009; Otto, 2023). The question of the relative weight of good intentions vs efficiency in judgments of public policies in this domain is important given the underlying moral intuitions that create divides around the use of new technologies to fight climate change (including but not limited to carbon capture technologies). To illustrate the division among environmental attitudes in Western societies, Pinker (2021) has proposed that those committed to combating climate change can be divided into two broad camps (forming the endpoints of a continuum). A first camp sees humanity as the contaminator of an otherwise pure and pristine planet. On this view, economic growth might never be compatible with combating climate change objectives, technological progress cannot play a major beneficial role in reducing emissions, while 'altruistic' intentions to incur sacrifices to curb humanity's emissions should be the priority. Because this view focuses on the importance

technologies might be an important ingredient in charting a path forward compatible with IPCC objectives of carbon emissions reduction. This is why we framed our vignettes on environmental policies around the issue of carbon capture in Experiment 1a, but also Experiments 2 and 3.

Method

Data availability

All studies were designed in Qualtrics and recruited French respondents via Foule Factory, a French participant recruitment platform equivalent to Prolific, ensuring high data quality. All data and R files can be downloaded at https://osf.io/kzpwd/?view_only=80c7d92cebfa451996a017edcbd2f166.

Hypotheses

On all four issues of Experiments 1a–d, we predicted that respondents would judge an altruistic but low-efficiency policy as significantly more morally commendable than a highly efficient but selfish policy (H1). Experiments 1a–d were not preregistered, but their vignettes and designs were preset to be comparable across issues. Sample sizes and analyses in all four studies were defined in advance based on piloting work.

Participants

A total of 854 French participants were recruited for Experiments 1a–d through four distinct experiments. Data were collected between June and August 2019. All respondents were paid €0.50. No respondent was excluded because all data were of good quality and submissions were complete. Respondents in the final datasets had the following characteristics: Experiment 1a on reducing CO₂ emissions with carbon capture: $N = 247$; $Mean_{age} = 44.2$, Standard Deviation (SD)_{age} = 13.5, 61.5% female; Experiment 1b on promoting sex equality in the workplace: $N = 211$, $M_{age} = 43.8$, $SD_{age} = 12.7$, 59% female; Experiment 1c on reinforcing France's diplomatic authority: $N = 202$, $M_{age} = 42.5$, $SD_{age} = 13.6$, 61% female; Experiment 1d on regulating immigration: $N = 194$, $M_{age} = 43.8$, $SD_{age} = 13$, 64% female.

Materials and procedure

For Experiments 1a–d, we selected two issues considered, in France, as being important for left-wing people (CO₂ reduction, sex equality), and two issues viewed as more important for right-wing people (maintaining France's diplomatic authority, immigration regulation). Since our focus is on environmental attitudes in this paper, our description of the procedures is based on the issue of carbon capture (Experiment 1a), on which a CEO was the policymaker. Minor wording differences otherwise differentiated Experiments 1a–d (see Supplementary, A).

of well-intended efforts to reduce emissions and purify the planet, it tends to give consequentialist considerations of impact and efficiency a secondary role. On the other hand, a second, techno-optimist camp sees human ingenuity and engineering as capable of innovating technological solutions that could play a major role in reducing primary emissions but also those already released in the atmosphere, like carbon capture and storage. On this view, technological innovation is seen as a crucial and realistic solution for steering society toward developing impactful and cost-effective techniques to reduce carbon emissions.

Participants had to give their informed consent to participate. They were randomly presented with two consecutive vignettes (in a within-subjects design) that narrated the discussion between a policymaker (in Experiment 1a, a CEO) and his advisor on a new policy they were considering implementing. The style and structure of the vignettes were inspired by vignette studies by Knobe (2003).

In the Altruistic/Low-efficiency policy condition, the CEO reported being driven by the altruistic intention to help the issue when considering deploying the policy: 'I honestly care deeply about [fighting climate change].' However, his advisor informed him that the policy would cost 100 million euros to the company and that it would only do little to help the issue: '+10%' positive impact at the national level and '+0.001%' impact at the global level (where positive impact meant carbon emissions *reduction* in Experiment 1a).

By contrast, in the Selfish/High-efficiency condition, the policymaker pursued a selfish reward unrelated to tackling the policy issue: 'Honestly, I don't care about [fighting climate change]. What I care about is [the profits we can make and our company's success.]' However, his advisor notified him that the policy would allow for 100 million euros to be saved by the company, and that it would be highly effective at helping the issue: '+80%' positive impact at the national level, and '+1%' impact at the global level (i.e., impact meant emissions reduction in Experiment 1a). Although this summary information was not explicitly stated in the vignettes, the chosen figures implied that the Selfish/High-efficiency policy was 8 times more impactful at a national level, and 1000 times more impactful at the global scale, than the Altruistic/Low-efficiency policy, in addition to being vastly profitable financially as opposed to costing high amounts of money. Below is the translation to English of the vignettes from Experiment 1a. Vignettes were presented on two consecutive pages:

Participants were asked, 'To what extent would you say that the [CEO's] decision was commendable?' displayed under each vignette (1, 'Not at all commendable'; 7, 'Totally commendable'). They read the vignette and morally assessed the policy decision twice, once per condition.

Participants then reported their level of moral commitment to the issue addressed by the policy decision – environmental protection in Experiment 1a – by answering the question 'To what extent do you think that [protecting the environment] should be the government's priority?'. Responses were collected on a slider scale ranging from 0, 'I don't care' to 100, 'Absolute priority'. The questionnaires of the four studies all ended with demographic questions: political orientation on a one-item left-right axis, sex, age and level of education.

See Supplementary, A for vignettes and materials.

Results

All analyses were run in R (version 4.4.1) using R Studio (version 2024.12.0+467).

Effects of decision type on commendability

Figure 1 presents judgments of commendability of each policy decision were first regressed with standard Ordinary least squares (OLS) on data aggregating the four Experiments 1a–d with policy decision, issue and respondent's moral commitment

Altruistic/Low-efficiency policy	Selfish/High-efficiency policy
The CEO of a global cement and concrete company is considering implementing a new policy: directly capturing CO ₂ as it is leaving its factories' smokestacks, and then recycling it as an ingredient that makes the factories' concrete stronger. [Both conditions]	
His chief advisor told him, 'If you implement this policy, you will cut our company's CO ₂ emissions by 10%. But even with our company's size, this would make only a minor difference in the fight against climate change, as it would decrease the world's total CO ₂ emissions by only 0.001%. This policy would only help the world a little bit.'	His chief advisor told him, 'If you implement this policy, you will cut our company's CO ₂ emissions by almost 80%. Given our company's size, the impact on fighting climate change would be massive, as it would decrease the world's total CO ₂ emissions by 1%. This is the kind of policy that could really help the world.'
The CEO responded to the advisor, 'I honestly care deeply about fighting climate change. How much would the policy cost?'	The CEO responded to the advisor, 'Honestly, I don't care about fighting climate change. What I care about are the profits we can make and our company's success. How much would the policy cost?'
The advisor responded, 'By my calculations, our company would lose about 100 million euros with that policy, because the new technology would significantly increase production costs.'	The advisor responded, 'By my calculations, our company would make about 100 million euros with that policy, because the new technology would significantly decrease production costs.'
The CEO decided to implement the policy. He indeed reduced his company's CO ₂ emissions by 10%, thus reducing the world's total net CO ₂ emissions by 0.001%. However, because the CEO cared deeply about his action, he thought it was definitely worth the 100 million euros in lost profits.	The CEO decided to implement the policy. He indeed reduced his company's CO ₂ emissions by 10%, thus reducing the world's total net CO ₂ emissions by 0.001%. He indeed reduced his company's CO ₂ emissions by 80%, thus reducing the world's total net CO ₂ emissions by 1% and significantly contributing to the fight against climate change. However, more importantly to the CEO, he also made a nice profit of 100 million euros.

as predictors. The same models were then run separately on each issue, corresponding to each separate Experiment 1a–d (see Supplementary, C). We report standardized regression coefficients, 95% confidence intervals (CIs) between squared brackets, and *ps* from these models.

In analyzing aggregated data from all four studies, respondents' commendability ratings were highest regardless of policy decision on the issue of environmental protection, followed by sex equality (main effect of issue: $\beta = -0.17$, $[-0.29, -0.05]$, $p < 0.001$), promoting France's authority ($\beta = -0.51$, $[-0.64, -0.38]$, $p < 0.001$), and regulating immigration ($\beta = -0.63$, $[-0.76, -0.50]$, $p < 0.001$).

As we feared, respondents rated the Selfish/High-efficiency decision as substantially less morally commendable than the Altruistic/Low-efficiency policy decision (main effect of decision type: $\beta = -0.54$, $[-0.62, -0.45]$, $p < 0.001$ with the Altruistic/Low-efficiency policy decision defined as baseline).

We recognized that our within-subjects design could have incentivized participants to exaggerate the difference between their two judgments compared to a between-subjects design, simply because they saw the two policy scenarios sequentially. We thus compared again ratings of the two conditions while restricting the dataset to the first condition (i.e., policy) that participants had been exposed to. In this analysis,




Figure 1. Judgments of commendability of policy decisions as a function of decision type (within-subjects) and policy issue (separate experiments) in Experiments 1a–d. Black lines are medians, red dots are means and black whiskers are 95% CIs of the mean.

the Altruistic/Low-efficiency decision was still seen as substantially more commendable than the Selfish/High-efficiency decision (main effect of decision type: $\beta = -0.56$, $[-0.68, -0.44]$, $p < 0.001$).

We now turn to results focusing on each experiment separately. When the policymaker was a CEO who tried to reduce CO₂ emissions from his company through carbon capture, the Selfish/High-efficiency decision was judged as substantially less commendable than the Altruistic/Low-efficiency policy ($\beta = -0.99$, $[-1.15, -0.84]$, $p < 0.001$). A similar effect was found when the policymaker was a CEO who tried to reduce sex inequality in his company, with the Selfish/High-efficiency decision being rated as substantially less commendable than the Altruistic/Low-efficiency policy ($\beta = -0.82$, $[-0.99, -0.64]$, $p < 0.001$).

By contrast, when a minister implemented a policy aiming to promote France's authority in the world and to regulate immigration, the difference in perceived commendability between the Altruistic/Low-efficiency policy and the Selfish/High-efficiency decision was smaller. Still, the Altruistic/Low-efficiency policies were rated as being slightly more commendable than their Selfish/High-efficiency counterparts, not less, although the difference was not always significant (France's authority: $\beta = -0.13$, $[-0.33, 0.06]$, $p = 0.167$; immigration regulation: $\beta = -0.23$, $[-0.41, -0.04]$, $p = 0.016$).

Associations between moral commitment to the issue and commendability judgments

Figure 2 shows that participants were on average most morally committed to the issue of environmental protection ($M = 80.2$, median = 82, $SD = 18.8$), followed by promoting sex equality ($M = 68.9$, median = 71, $SD = 23.4$), France's authority ($M = 52.3$, median = 54, $SD = 21.7$) and regulating immigration ($M = 56.3$, median = 60, $SD = 28$).




Figure 2. Commendability judgments of policy decisions as a function of respondents' moral commitment to the issue, decision type, and breaking down by issue in Experiments 1a–d. Colored lines are simple linear regression slopes, and greyed areas are the 95% CIs.

On all four issues, greater moral commitment to the issue was associated with more praise of the Altruistic/Low-efficiency decision. However, on issues of reducing carbon emissions and fighting sex inequality (on which the policymaker was a CEO), greater moral commitment predicted less commendability of the Selfish/High-efficiency decision (reducing carbon emissions: main effect of commitment: $\beta = 0.00$; commitment $\times$ policy: $\beta = -0.36$, $[-0.51, -0.21]$, $p < 0.001$; promoting sex equality: main effect of commitment: $\beta = 0.07$, $[-0.01, 0.16]$, $p = 0.095$; commitment $\times$ policy: $\beta = -0.32$, $[-0.49, -0.14]$, $p < 0.001$). By contrast, on the issues of promoting France's authority and regulating immigration (where the policymaker was a minister), greater moral commitment was associated with higher perceived commendability of the Selfish/High-efficiency decision (promoting France's authority: main effect of commitment: $\beta = 0.21$, $[0.11, 0.30]$, $p < 0.001$; commitment $\times$ policy: $\beta = -0.06$, $[-0.25, 0.13]$, $p = 0.52$; regulating immigration: main effect of commitment: $\beta = 0.36$, $[0.27, 0.46]$, $p < 0.001$; commitment $\times$ policy: $\beta = 0.01$, $[-0.17, 0.20]$, $p = 0.9$).

Simple slope analyses of the associations between moral commitment and judgments of commendability of each policy decision considered separately can be found in Supplementary, C.

Discussion

A policy decision that only helped reduce CO₂ emissions through carbon capture a little bit at a high cost, but which was altruistically motivated, was judged as more commendable than a policy at least eight times more impactful at reducing emissions and financially profitable, but selfishly motivated. This result, based on efficiency information given in numeric format, was generalized to three other issues and was observed both when the policymaker was the CEO of a private company or a minister. Moreover,

the more participants were morally committed to an issue, the more commendable they saw an altruistic policy decision aiming to tackle it, despite its poor impact and high financial cost. Finally, we note that the Altruistic/Low-efficiency policy decisions were rated higher when a CEO implemented them. This might be due to French respondents demanding greater contributions to the common good from public servants than from CEOs (Zitellmann, 2023), resulting in the thought that CEOs deserve more praise than ministers for similar prosocial decisions.⁴

Experiment 2

The designs employed in Experiments 1a–d showed that well-intended but inefficient policies tended to be seen as more commendable than highly efficient policies motivated by self-interest, at least when policy information is provided in numeric format. However, the degree to which those findings are due to efficiency insensitivity or to intentions being given high importance is unclear. To answer this question, Experiment 2 independently manipulated the policy's efficiency, the decision maker's intent and their identity. Finally, to provide a more conservative measurement of people's (in)sensitivity to efficiency, the outcome variable was reframed in terms of *support* for the policy.

Method

Preregistration

The design, materials, sample size, hypotheses and analyses of Experiment 2 were pre-registered on OSF: https://osf.io/ry3qn/?view_only=1a1b3741cc114adb8ef3851072ad29c4.

Hypotheses

Based on Experiments 1a–d and a pilot study of Experiment 2 ($N = 661$), we made the following preregistered hypotheses:

H1: Participants will support more high efficiency policies than low efficiency policies *ceteris paribus*.

H2: Participants will support more altruistic policies than policies in which no intention is specified *ceteris paribus*.

H3: Participants will support less selfish policies than policies in which no intention is specified *ceteris paribus*.

H4: The altruistic but low efficiency policies will be supported more than the selfish but highly efficient policies (across levels of policymaker identity), like in Experiments 1a–d.

⁴ A second and complementary explanation is that participants were, on average, more morally committed to the former two issues than to the latter two, and this might have driven up ratings of commendability of the Altruistic/Low-efficiency policy.

Participants

A preregistered *a priori* power analysis based on a pilot study of $N = 661$ was used to estimate the target sample size of Experiment 2. We took a medium effect of policymakers' intention (altruistic vs no intention) on judgments of support of the policy, $d = 0.38$, as the effect we wanted to be able to detect in the final study. Gpower estimated that 87 participants per group were required to detect $d = 0.38$ at 80% power. Given our 12 experimental conditions design, this means that 1044 participants were required. A total of 1203 French participants were recruited on Foule Factory for Experiment 2, which lasted about three minutes, in exchange for €0.75. Ninety-eight respondents who failed the attention check were excluded from the data. Our final dataset retained 1105 respondents ($M_{\text{age}} = 42$; $SD_{\text{age}} = 13.2$; 50% female). Data were collected in October 2023.

Materials and procedure

Experiment 2 focused on the issue of protecting the environment using carbon capture, the type of environmental policy that was the focus of our interest in this project. The vignettes and procedure were the same as in Experiment 1a except for the following changes. Instead of comparing two decision types only, we independently varied the policymaker's intention (altruistic intent, no intention specified and selfish intent), the policy's efficiency (low vs high efficiency) and the policymaker's identity (CEO vs Minister). The design was now a $3 \times 2 \times 2$ design with 12 conditions, between-subjects to avoid repeated reading of the vignettes artificially inflating differences in policy support.

Altruistic motivations reported by the policymaker were the same as in Experiment 1a ('I honestly care deeply about fighting climate change'). The no intention conditions were created by suppressing the sentences mentioning the policymaker's reported motivations to implement the policy. We meant the 'no intention' baseline to gauge how responsive to policy efficiency participants would be in the absence of any cue to the decision maker's motivations. As regards the selfish motivations of the policymakers, we modified them slightly to make them more believable and less cynical. When a CEO implemented the policy, we used the formulation: 'Honestly, I don't care about fighting climate change. What I care about is to improve the image of our company' (instead of 'What I care about are the profits we can make and our company's success' in Experiment 1). When the policy maker was a minister, we used the formulation: 'Honestly, I don't care about fighting climate change. What I care about is to improve our political party's image to increase chances of winning the presidential elections' (instead of 'What I care about is making budget savings that could increase chances of our party winning the next election' in Experiment 1). Selfish motivations were thus reputational, rather than pecuniary, to avoid a tension with scenarios in which the policy costs a high amount of money (see below).

Second, we increased the contrast between the low and high levels of the efficiency factor compared to Experiments 1a–d to make it more salient. Moreover, policy efficiency was described at the national level only (contrary to Experiments 1a–d) to make it easier to represent for lay participants. Third, the money sum described as being saved vs wasted by the implementation of the policy was raised to a more ecological

figure: 13 billion euros. Although this was not said in the vignettes, this is approximately the amount that the French state spends each year in ‘sustainable development policies’ (<https://www.economie.gouv.fr/facileco/comptes-publics/budget-etat>). Thus, in the low-efficiency conditions, the policy was described as reducing the French industry’s carbon emissions by only 2% at the national level, and as costing 13 billion euros in the long run (to the state when it was a minister vs to the company when it was a CEO). By contrast, in the high-efficiency conditions, the policy was described as contributing to an 80% reduction of the French industry’s CO₂ emissions at the national level, and as allowing 13 billion euros to be saved in the long run. Although this was not explicitly said in the vignettes, the figures entailed that the high-efficiency policy was 40 times more impactful at a national level than the low-efficiency policy, in addition to being vastly financially profitable as opposed to costing a lot of money.

Fourth, we realized that the dependent variable framed in terms of policy ‘decision’ used in Experiments 1a–d could implicitly invite participants to focus on the *moral character* of the actor’s decision to the detriment of the policy’s efficiency and impact. This could artificially increase chances that their responses neglect costs and benefits. We thus asked: ‘To what extent do you support this policy?’ (1, ‘Not at all’; 7, ‘Totally’), immediately below the vignette.

On the following page, respondents were also asked to what extent they saw the CEO or minister as an ‘ally’ of environmental protection: ‘For you, what does the CEO [Minister] represent in the struggle to protect the environment in the long run?’ (0, ‘clearly an enemy’ to 100, ‘clearly an ally’). This additional question was introduced to measure the extent to which policy support tracks inferences about the trustworthiness of the policymaker.

Next, participants reported their level of moral commitment to environmental protection on a three-item scale: ‘Protecting nature is an absolute moral imperative,’ ‘The conviction that one must fight to protect the environment is central to my identity’ (inspired from Skitka *et al.*, 2005; Skitka, 2010); ‘Protecting the environment should be the government’s priority’ (0, ‘Totally disagree’; 50, ‘I don’t know’ as default position; 100, ‘Totally agree’).

The experiment ended with demographic questions: political orientation on a 1-item left-right axis, sex, age and education, an attention check (see Supplementary, B), and an invitation to send written feedback to the experimenters.

Results

Effects of policy efficiency, intent and policymakers’ identity on support for policy

The results shown in Figures 3 and 4 were analyzed using multiple regressions of policy support on policy efficiency (low vs high), policymaker intent (unspecified vs altruistic vs selfish) and policymaker identity (CEO vs minister) (see Supplementary, D). Levels listed first in parentheses were set as baselines for each factor in the main regression analyses. Main effects are reported from a model containing only the main effects, and the two-way interactions from a model containing the main effects and the interaction.

Policies implemented by a minister garnered less support across the board than those made by a CEO ($\beta = -0.35$, $[-0.45, -0.26]$, $p < 0.001$).




Figure 3. $N = 1105$. Support for policy in Experiment 2 as a function of intent and efficiency, breaking down by policymaker identity. Black lines are medians, red dots are means and black whiskers are 95% CIs of the mean.

The high-efficiency policies were supported more than the low-efficiency policies, regardless of intention and policymaker identity ($\beta = 1.04$, [0.95, 1.13], $p < 0.001$). This difference was greater when the policymaker was a minister than a CEO (efficiency $\times$ identity: $\beta = 0.31$, [0.13, 0.49], $p < 0.001$), an effect driven mainly by low-efficiency policies being supported less when a minister was implementing them than when it was a CEO.

Restricted analyses of conditions that did not ascribe any intention to the policymaker allow us to gauge the extent to which judgments of support were sensitive to a large contrast in policy efficiency. When no intent was ascribed to the policymaker, high-efficiency policies were more supported than low-efficiency policies, $\beta = 1.20$, [1.05, 1.35], $p < 0.001$. Although this latter effect was, statistically speaking, of a fairly large size, it is striking that participants' judgments showed so little responsiveness to such large differences in efficiency as described by the vignettes. While support for high-efficiency policies not mentioning any intent was around 'Clearly' on the support scale, support for low-efficiency policies not mentioning intent dropped by only 2 scale points to 'Undecided', despite the latter having a gigantic financial cost and tiny environmental protection benefit.

Mention of a selfish motivation in the policymaker decreased support for the decision compared to when no intention was specified (main effect of selfish intent: $\beta = -0.45$, [-0.56, -0.34], $p < 0.001$). By contrast, mentioning an altruistic motivation driving the policymaker slightly increased support for the policy (main effect of altruistic intent: $\beta = 0.11$, [0.00, 0.22], $p = 0.045$). The former effect was larger than the latter, perhaps due to a ceiling effect on policy support.




Figure 4. $N = 1105$. Support for policy in Experiment 2 as a function of intent and efficiency (aggregating across policymaker identity). Black lines are medians, red dots are means and black whiskers are 95% CIs of the mean.

Participants in Experiment 2 supported the Altruistic/Low-efficiency policies slightly less than the Selfish/High-efficiency policies. Though these results contrast with Experiments 1a–d, it is important to notice how small the preferences for the more efficient policies were. Those comparisons were run with regression analyses comparing only those two experimental conditions on a restricted dataset. Breaking down by policymaker identity, when the policy was deployed by a minister, the Altruistic/Low-efficiency policies were rated $M = 4.06$, $SD = 1.62$ and the Selfish/High-efficiency policies $M = 5.04$, $SD = 1.66$ ($\beta = 0.55$, $[0.29, 0.82]$, $p < 0.001$). When the policy was deployed by a CEO, ratings of support were even closer, with the Altruistic/Low-efficiency policies being rated $M = 5.08$, $SD = 1.51$ and the Selfish/High-efficiency policies $M = 5.65$, $SD = 1.34$ ($\beta = 0.29$, $[0.15, 0.43]$, $p = 0.007$). This means that large discrepancies in policy efficiency only had a small impact on policy support.

Associations between moral commitment to environmental protection and policy support

As shown in Figure 5, the more morally committed to environmental protection respondents were, the more they supported the CO₂-reduction policies. Pooling




Figure 5. $N = 1105$. Support for policy in Experiment 2 as a function of respondents' moral commitment to protecting the environment, policymakers' intent and policy efficiency, broken down by policymaker identity. Simple linear regressions with 95% CIs as greyed areas.

all policy conditions together, the overall effect of commitment on support was: $\beta = 0.23$, $[0.19, 0.28]$, $p < 0.001$ (we report here the most conservative estimate; see Supplementary for a comparison between models). The positive effect of commitment on policy support was slightly smaller on high (vs low) efficiency policies, presumably due to a ceiling effect on the policy support scale (commitment $\times$ high efficiency: $\beta = -0.15$, $[-0.24, -0.06]$, $p = 0.001$). Specifying that the policymaker was motivated by altruism rather than giving no intention-relevant cues did not moderate the positive effect of moral commitment on policy support (commitment $\times$ intent: $\beta = -0.01$). However, specifying that the policymaker was motivated by selfishness rather than providing no intention-related information decreased the positive effect of moral commitment on policy support (commitment $\times$ intent: $\beta = 0.24$, $[-0.35, -0.13]$, $p < 0.001$). This was particularly visible when the policymaker was a CEO. Simple slope analyses of moral commitment to environmental protection on policy support in each experimental condition are reported in Supplementary, D. Overall, moral commitment to environmental protection was associated with more support for *any* environmental policy, regardless of the policies' contrasts in efficiency – except for the two selfish policies deployed by a CEO.

Finally, we found that support for the environmental policies is strongly predicted ($\beta = 0.60$, $[0.55, 0.65]$, $p < 0.001$) by perceptions of the policymaker as being an *ally* (vs an enemy) of the cause of environmental protection; what may be called their trustworthiness. See end of Supplementary, D for visualizations and analyses, not reported here for reasons of space.

Discussion

As in Experiments 1a–d, large numeric discrepancies in environmental policy efficiency had relatively little effect on their popularity in Experiment 2. Although the

low-efficiency policies were about 40 times less impactful at reducing CO₂ emissions than their low-efficiency alternatives, the change in support between them was no greater than two points on a seven-point scale ($\beta = 1.04$, [0.95, 1.13], $p < 0.001$), and the low-efficiency policies were overall not disapproved of. Moreover, the altruistic but low-efficiency policies were only slightly less supported than the selfish but high-efficiency policies. Holding efficiency constant, altruistic policies were supported more than policies ascribing no motivation to the policymaker, and selfish policies were disapproved of the most.⁵

In contexts where policy efficiency is described numerically, these intuitions may lead to altruistic but inefficient policies being supported roughly to the same extent as efficient policies perceived as selfishly motivated.

Finally, greater moral commitment to environmental protection was associated with more support for *any* policy aiming to reduce emissions, regardless of its efficiency and of the motivations driving it. Strong moralization of a policy issue apparently makes people less likely to apply cost-benefit thinking to it (Tetlock, 2003; Skitka, 2010).

Experiment 3

Experiments 1a–d and 2 suggest that respondents give relatively little importance to large discrepancies in efficiency in reducing carbon emissions when efficiency information is provided numerically. At the same time, they give substantial importance to the personal motivations driving the policymakers – which arguably should matter little given they are one-off decisions. With Experiment 3, we asked whether being able to compare an altruistic but inefficient policy to a highly efficient but selfish policy by seeing them *simultaneously* (vs in isolation) could reduce efficiency neglect. We expected a simultaneous presentation of the two scenarios to make the policies' characteristics more salient, thus making respondents more attentive to their efficiency, and less to the personal motivations driving them. Crucially, Experiment 3 also introduced reference points and qualitative appraisals of the policies' cost and efficiency to help respondents process efficiency contrasts more intuitively.

Method

Preregistration

The design, materials, sample size, hypotheses and analyses of Experiment 3 were preregistered on OSF: https://osf.io/qps6b/?view_only=67dc8635ba394678b65385dc5cd423ee.

Hypotheses

We made the following preregistered hypotheses:

H1: The altruistic but low efficiency environmental policies will be supported more, or to a similar extent, as the selfish but highly efficient policies.

⁵On high-efficiency policies, the no intentions conditions tended to pattern the same as the altruistic intent policies, possibly because respondents tended to implicitly infer prosocial intentions even when they were not specified.

H2: Seeing descriptions of the policies simultaneously will increase support for the selfish but highly efficient policy compared to when the policy descriptions are seen sequentially.

H3: Seeing descriptions of the policies simultaneously will decrease support for the altruistic but low efficiency policy compared to when the policy descriptions are seen sequentially.

Participants

We powered Experiment 3 over a small effect of $d = 0.15$ of the interaction between scenario (Altruistic/low-efficiency vs Selfish/high-efficiency, within-subjects) and mode of presentation (simultaneous vs sequential, between-subjects). A power analysis in GPower estimated that 400 participants per group in an independent groups design were required to detect $d = 0.15$ at 80% power. 816 French participants were recruited on Foulé Factory for about 4 minutes, paid €1.00. Sixteen participants who failed the attention check were excluded. Our final dataset contains 800 ($M_{\text{age}} = 44$; $SD_{\text{age}} = 13$; 50% female). Data were collected in April 2024.

Materials and procedure

The design of Experiment 3, which focused on carbon capture technology, was based on that of Experiment 1a. Contrary to Experiment 1a, Experiment 3 only portrayed policy decisions as being made by a minister (not a CEO) because judgments of support in Experiment 2 displayed higher variance with a minister. Importantly, Experiment 3 also manipulated whether respondents would consider the Altruistic/Low efficiency and Selfish/High-efficiency scenarios *simultaneously* (on the same survey page) vs *sequentially* (on two consecutive pages, like in Experiments 1 and 2). All participants saw both policy scenarios (Altruistic/Low-efficiency vs Selfish/High-efficiency shown as within-subjects factor and sequential vs simultaneous presentation presented as between-subjects factor).

In the *sequential* presentation condition, the survey had the same structure as in prior studies. Respondents were randomly exposed to one of the two scenarios first, asked to what extent they supported it, and then on the next page, they saw the description of the second policy and reported their support for it. In the *simultaneous* presentation condition, information about both policies was presented in two columns of text side by side on the same page (counterbalanced). We reasoned that putting policy characteristics side by side might make the (normatively important) information about policies' impact on emissions and financial cost easier to remember and to compare than when considered in sequential order. In this simultaneous presentation condition, the two columns of text containing the policy descriptions were introduced to participants as two distinct dialogues between a minister and his advisor, which they were asked to read 'very carefully, from left to right, one after the other.' In all conditions, to facilitate discrimination of the two scenarios, both ministers were given a neutral French family name: Mr. Lefevre or Mr. Ligneul. Our dependent variable was: 'To what extent do you support the public policy proposed by Mr. Lefevre [Mr. Ligneul]?' (1, 'Not at all' to 7, 'Totally').

The selfish motivation driving the minister in one of the two scenarios was again to improve the party's image. As regards policy efficiency, the dialogues in the vignettes stipulated 80% reduction (high efficiency) vs 1% reduction (low efficiency) of the French industry's CO₂ emissions at the national level (vs 2% reduction in Experiment 2, the goal here being to make the low-efficiency policy appear even more inefficient). Like in Experiment 2, the scenarios were described as entailing 13 billion euros to be saved (high efficiency) vs wasted (low efficiency), a figure equivalent to the French budget for sustainable development.

Experiment 3 introduced another crucial difference in the policy description: numerical information provided by the advisor about the financial cost of the policies was accompanied by *reference points for the policy's cost*, and *qualitative appraisals* of its efficiency, to make the policies' efficiency more intuitive to process. In the high-efficiency scenarios, 'The policy alone would *allow the entire annual state budget for sustainable development to be saved*, which is *colossal*. The policy would therefore be *highly efficient* in terms of its cost-benefit ratio'. By contrast, in the low-efficiency policies, 'The policy alone would *cost the entire annual state budget for sustainable development*, which is *colossal*. The policy would therefore be *quite inefficient* in terms of its cost-benefit ratio'. We expected those cues (in bold text in vignettes) to reduce participants' propensity to neglect the policies' efficiency by providing *intuitive* ways of judging what may be considered as a high vs a low policy cost. Contrary to prior studies, key information about the ministers' motivations and the policies' impact and financial cost was emphasized in bold in the questionnaire. See vignettes in Supplementary, A.

Experiment 3 used the same three-item measure of moral commitment to environmental protection, demographic questions and the same attention check as in Experiment 2. These were positioned on a page following the vignettes.

Results

Effects of policy scenario and presentation mode on support for policy in Experiment 3

Tests of our main preregistered hypotheses ($\beta_{\text{preregistered}}$) were performed on a dataset containing judgments of support for the first policy viewed in the sequential condition, and judgments of support for both policies in the simultaneous presentation condition. We also report results of those tests on the whole dataset ($\beta_{\text{all data}}$). Interested readers can find mean levels of support for the policies in Supplementary, E, together with regression tables.

Figure 6 shows results from Experiment 3. Contrary to previous studies, participants supported the Altruistic/Low-efficiency scenarios substantially *less* than the Selfish/High-efficiency policies (main effect of scenario: $\beta_{\text{preregistered}} = 1.11$, [1.02, 1.20], $p < 0.001$; $\beta_{\text{all data}} = 1.13$, [1.05, 1.21], $p < 0.001$). For the first time in our series of studies, participants seemed to slightly reject the Altruistic/Low-efficiency scenario, with average support in the 'somewhat not' region. The fact that the Selfish/High-efficiency policy was now supported more than the Altruistic/Low-efficiency policy can most likely be attributed to the reference points for the policies' cost and qualitative appraisals of policy efficiency introduced in Experiment 3. As we show in more detail below, introducing those qualitative cues likely helped respondents grasp the contrast in efficiency between the scenarios more intuitively compared to Experiments 1 and 2.




Figure 6. $N = 800$. Support for Altruistic/Low-efficiency and Selfish/High-efficiency policies in Experiment 3 as a function of whether the two policies are presented sequentially or simultaneously to participants (mode of presentation). Importantly, following our preregistered analyses, only ratings of the first policy viewed in the sequential presentation mode are plotted here. Black lines are medians, red dots are means and black whiskers are 95% CIs of the mean.

Effect of the sequential vs simultaneous presentation in Experiment 3

See Figure 6, which shows that the manipulation of the mode of presentation of the policy scenarios had no detectable effect. Mode of presentation and scenario did not interact significantly ($\beta_{\text{preregistered}} = -0.15, [-0.35, 0.05], p = 0.154$; $\beta_{\text{all data}} = -0.14, [-0.30, 0.02], p = 0.088$). We then examined the effect of the manipulation on each scenario considered separately using *post hoc* analyses. Support for the Selfish/High-efficiency policy was not greater when both policies were presented simultaneously than sequentially ($\beta_{\text{preregistered}} = 0.01, [-0.16, 0.18], p = 0.914$; $\beta_{\text{all data}} = -0.06, [-0.20, 0.08], p = 0.387$). Moreover, rather than decreasing support for the Altruistic/Low-efficiency policy as we expected in H3, presenting both policies simultaneously marginally increased support for it compared to when the two policies were considered one at a time ($\beta_{\text{preregistered}} = 0.18, [0.01, 0.34], p = 0.038$; $\beta_{\text{all data}} = 0.11, [-0.03, 0.25], p = 0.129$).

Effect of introducing reference points and qualitative appraisals of policy efficiency in Experiment 3 vs 2

Let us now turn to what is arguably the most important result of Experiment 3.

Table 1. Summary of the differences between the policy scenarios most comparable to each other in Experiments 2 and 3

Decision maker	Experiment 2 (policy scenarios were always presented sequentially, in between-subjects design)		Experiment 3 (focusing on policy scenarios viewed sequentially, between-subjects)	
	Minister		Minister	
Intention	Altruistic	Selfish	Altruistic	Selfish
Efficiency	Low efficiency:	High efficiency:	Low efficiency:	High efficiency:
	2% reduction	80% reduction	1% reduction	80% reduction
	–13 billion euros	+13 billion euros	–13 billion euros	+13 billion euros
			Policy would ‘allow the entire annual state budget for sustainable development to be saved, which is colossal’	Policy would ‘cost the entire annual state budget for sustainable development, which is colossal’

In Experiment 2 (which had a between-subjects design), the Altruistic/Low-efficiency and Selfish/High-efficiency scenarios involving a Minister were highly similar to their counterparts viewed in a sequential order in Experiment 3. The key differences were the introduction in Experiment 3 of the reference point for the policies’ cost in terms of the state’s budget for sustainable development, and the qualitative appraisals of the policy’s efficiency. (The only other difference was a 2% vs 1% reduction in carbon emissions between the two experiments in the low-efficiency conditions). Table 1 below summarises the main differences between the information contained in the vignettes of Experiment 3 and 2. Focusing analyses on just these pairs of conditions allows us to approximately assess the effect of introducing the reference point and qualitative appraisals in Experiment 3.

Here is the key result: Introducing the reference point and qualitative appraisals of the policies’ efficiency in Experiment 3 made respondents *much* more responsive to policy efficiency compared to Experiment 2. See Figure 7 for visualizations. Comparing policy support in Experiment 3 to that in Experiment 2, respondents became markedly *less* supportive of the Altruistic/Low-efficiency policies, and slightly *more* supportive of the Selfish/High-efficiency policies (policy scenario $\times$ experiment: $\beta = 0.73$, $[0.47, 0.99]$, $p < 0.001$; *post hoc* analyses restricted to Altruistic/Low-efficiency policies: $\beta = -0.61$, $[-0.83, -0.39]$, $p < 0.001$; *post hoc* analyses restricted to Selfish/High-efficiency policies: $\beta = 0.26$, $[0.03, 0.48]$, $p = 0.027$).

Associations between moral commitment to environmental protection and policy support

We now go back to the analysis of Experiment 3 alone. Since the manipulation of presentation mode had no detectable effect on judgments of support (against our




Figure 7. Support for Altruistic/Low-efficiency and Selfish/High-efficiency policies implemented by a Minister in Experiment 3 vs Experiment 2, on the issue of carbon capture and storage. Focusing on just these conditions allows us to approximately gauge the effect of introducing a reference point for the costs of the policies (i.e., the Budget of the French state for sustainable development) and of the qualitative appraisals of the policies' efficiency. Black lines are medians, red dots are means and black whiskers are 95% CIs of the mean.

preregistered hypotheses), we report associations between policy support and moral commitment to environmental protection based on the whole dataset, visualized in Figure 8. Higher moral commitment to environmental protection was associated with higher support for the public policies aiming to reduce CO₂ emissions when averaging across them (main effect: $\beta = 0.09$, [0.05, 0.13], $p < 0.001$). More specifically, positive effects of moral commitment to the environment were markedly stronger on the less popular Altruistic/Low-efficiency scenarios than on the more popular Selfish/High-efficiency scenarios (commitment $\times$ policy: $\beta = -0.15$, [-0.23, -0.07], $p < 0.001$). Whether the two policies were presented simultaneously or in isolation did not affect the positive influence of moral commitment to the environment on policy support (commitment $\times$ presentation: $\beta = 0.01$, [-0.07, 0.09], $p = 0.778$). See Supplementary, E, for *post hoc*, simple slopes analyses of the associations between moral commitment and policy support in each condition.

Discussion

In contrast with Experiments 1a–d and 2, the altruistic but inefficient policy aiming to reduce CO₂ emissions was supported much *less* than the highly efficient but selfish policies ($\beta_{\text{preregistered}} = 1.11$, $p < 0.001$). This reversal in respondents' judgments can




Figure 8. N = 800 respondents. Support for Altruistic/Low-efficiency and Selfish/High-efficiency policies in Experiment 3 as a function of respondents' moral commitment to environmental protection and whether the two policies are presented sequentially or simultaneously to participants (mode of presentation). All data from Experiment 3 is shown. Simple linear regressions with 95% CIs as greyed areas.

likely be attributed to our introduction of a reference point and qualitative appraisals of the policies' impact and cost. Those qualitative cues, present only in Experiment 3, likely allowed respondents to process more intuitively the difference in policy efficiency (e.g., 'very efficient' vs 'very inefficient'; 'colossal' amounts saved or lost) and in more context (reference to the state's annual budget), rather than based on abstract numbers as in previous studies. This pattern of results gives hope that lay citizens can be nudged toward more attention to the efficiency of public policies if context and intuitive verbal summaries of their impact are provided. Contrary to expectations, presenting the characteristics of the policies simultaneously to facilitate comparisons of their efficiency levels did not make respondents more efficiency-focused in their judgments.

General discussion

Ministers and CEOs sit at the top of organizational chains of command that potentially allow them to turn citizens' policy preferences into concrete effects that benefit the common good. From standards of democratic sovereignty and corporate accountability, what should matter most is how successful the policies they deploy are

at reaching citizens' policy choices, while keeping track of opportunity costs (Herzog and Hertwig, 2025). This is especially relevant in market societies in which pecuniary rewards can be an engine of socially beneficial innovation (Rubin, 2003; Hall and Rosenberg, 2010; Smith, 2013).

Contrary to this ideal, our six experiments ($N = 2759$) found that French people's judgments of policy decisions were only moderately responsive to large differences in efficiency, at least when expressed in numerical format. In Experiments 1a–d ($N = 854$), the fact that high-efficiency policy decisions were at least 8 times more impactful than low-efficiency policies, and that they would make society save rather than lose 100 million euros, never sufficed to make participants judge high-efficiency decisions as more commendable than low efficiency ones. Experiment 2 ($N = 1105$, preregistered) showed that in the absence of information on the policymaker's motivations, high-efficiency policies 40 times more impactful than low-efficiency policies – and hugely profitable financially rather than costly – received only 1.5 scale points more support than low-efficiency policies, on a seven-point scale ($d = 1.12$). As a stark illustration of this trend in Experiment 2, the altruistic but low-efficiency policies were supported only slightly less than the selfish but high-efficiency policies.

Encouragingly, however, Experiment 3 ($N = 800$, preregistered) suggests that when providing a standard to which the policy's cost can be compared – that is, the state's budget for environmental transition – and qualitative appraisals of its efficiency, it is possible to nudge people toward placing a higher emphasis on policy efficiency.

How should one interpret our respondents' relatively low responsiveness to efficiency contrasts? One possibility is that, while citizens do care about policy outcomes, they are not always cognitively equipped to process large differences in impact or financial cost when this information is presented in purely numerical terms (e.g., percentages). Indeed, research on 'number numbness' has shown that numerical information is not intuitive and can fail to trigger emotional engagement (Peters *et al.*, 2006; Slovic, 2007).

At the same time, comparing results from Experiment 3 and 2 suggests that the lack of standards for assessing policy efficiency played a role too. Participants became more discerning of efficiency contrasts when reference values (i.e., state's budget for environmental transition) and qualitative appraisals were introduced in Experiment 3. What appears as efficiency neglect may thus also spring, in part, from respondents lacking intuitive reference points for assessing what 'counts' as an effective (vs ineffective) and costly vs profitable environmental policy.⁶

The upshot is that designing more persuasive political communication campaigns might require conveying efficiency information using more contextual and qualitative appraisals, and more intuitive graphical ways of representing quantitative differences. Both approaches have received empirical support prior to our studies and may nudge people to be more efficiency-aware (Cleveland and McGill, 1984; Chong and Druckman, 2007; Keren, 2010; Aarøe and Petersen, 2018).

⁶It remains difficult to isolate the independent contribution of the reference point and the qualitative appraisals on judgments of policy support introduced in Experiment 3. The budgetary reference point likely served as a concrete numerical anchor, inviting further contextual interpretation, while the qualitative appraisals provided an explicit cue that required no additional interpretation. Future research would benefit from disentangling the specific effects of these cues.

Against a pragmatic focus on policy efficiency, we also found that cues to the presence of altruistic vs selfish intentions behind the policies' implementation have substantial influence on how morally commendable they are perceived as being (Experiments 1a–d), and how much popular support they get (Experiment 2). In Experiment 2, when the policy's efficiency was low (i.e., when support for it was not already near the scale's maximum), adding cues to the policymaker's altruism had a moderate increasing effect on support compared to when no intention-relevant information was provided. This trend was particularly prominent among respondents who had a strong moral conviction about the necessity of environmental protection. Conversely, portraying the policy as motivated by selfish concerns – financial or electoral – decreased support for it compared to when information on intentions was absent, regardless of how efficient the policy was.

The weight of intentions throughout our studies confirms that moral evaluations and public opinion are powerfully shaped by intuitive cognitive processes (Haidt, 2001; J. D. Greene *et al.*, 2009; J. Greene, 2014; Uhlmann *et al.*, 2015). Policymakers, if they are to mobilize the public in favor of their policies, should signal prosocial intentions both in themselves and in the programs they promote. This idea has been documented already in the case of attitudes toward welfare, where beliefs about recipients' deservingness condition popular support for redistribution (Aarøe and Petersen, 2014; Petersen and Arceneaux, 2020).

While our approach emphasizes the importance of evaluating policy programs based on their actual impact and efficiency – that is, adopting a more consequentialist lens to achieve key societal goals – we acknowledge that paying attention to decision-makers' intentions can be rational. Individuals driven by narrow self-interest may behave unpredictably, supporting policies or decisions only when they align with their personal goals of the moment. In contexts where decision-makers are responsible for multiple decisions over time – as is often the case for politicians or CEOs – this unpredictability can be risky. By contrast, a policy-maker motivated by prosocial values may offer greater consistency and alignment with the public interest over time, whereas someone driven by financial gain only may eventually make decisions that undermine collective welfare.

At a more fundamental level, human social cognition evolved in small groups in which individuals would engage in long-term, mutually beneficial partnerships. Selecting partners with consistently prosocial commitments was adaptively key to benefit from cooperation (Baumard *et al.*, 2013; Uhlmann *et al.*, 2015). It is therefore not surprising that our respondents' concern for the moral character of other people manifests strongly with respect to political and economic leaders today.

Tracking moral character can serve as a reasonable heuristic for anticipating future behavior. However, our point in this paper was to suggest that this attentiveness to intentions may spill over into evaluations of one-off policy scenarios where such heuristics are arguably less relevant. In such cases, it may lead citizens to reward policies that signal moral commitment, even when those policies are objectively less effective at achieving their goals.

Experts' and politicians' discussions of the costs and benefits of policies typically rely on figures and percentages, which are then communicated by mass media and discussed by the public in informal conversations (Lippmann, 1929; Katz and Lazarsfeld,

1955). Those conversations are bound to be biased in favor of bits of information people can represent and remember better and find relevant to pass along (Wilson and Sperber, 2006; Aarøe and Petersen, 2018; Marie *et al.*, 2020; Mercier, 2020). To the extent, then, that policymakers' decisions respond to citizens' and consumers' preferences and reactions – as they should in a democracy – the combination of low responsiveness to efficiency and attention to leaders' intentions can pave the way to expensive policy programs that concretely achieve little so long as they appear driven by values citizens hold dear. Conversely, citizens may turn their back on efficacious and profitable policies whose motivation is perceived as lacking nobility.

Finally, the positive effects of moral commitment on policy support documented across our studies confirm prior observations that people's ability or motivation to engage in cost-benefit thinking appears lowered when they see the issue in *moralized* terms (Baron, 1994; Baron and Spranca, 1997; Tetlock *et al.*, 2000; Tetlock, 2003; Skitka, 2010). Experiments 1a–d showed that strong moral commitment to environmental protection predicted greater support for well-intended but highly inefficient and expensive policies aimed at reducing CO₂ emissions, and similar trends were seen on three other issues. Similarly, Experiment 2 showed that high moralization of environmental protection increases support for almost *any* policy that somehow contributes to reducing carbon emissions, even if it is highly inefficient. Granted, protecting the environment and reducing carbon emissions are among the biggest challenges of our time and warrant prioritization over some other goals. But other policy issues are essential to human wellbeing as well (e.g. for France, funding the health care system, education, investing in national defense, etc.). Pursuing those competing goals requires the ability to engage in pragmatic give and take.

We hope our work prompts further research on the communication formats most conducive to the cost-benefit thinking our modern societies require (Pinker, 2018; Rosling, 2023), in particular by developing compelling graphic displays of differences in efficiency.

Supplementary material. To view supplementary material for this article, please visit <https://doi.org/10.1017/bpp.2025.10025>.

Acknowledgments. Our studies received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013)/ERC Grants 324115–FRONTSEM (PI: Schlenker) and 313610–SEMEXP (PI: Chemla), and ERC H2020 Grant Agreement No. 788077–Orisem (PI: Schlenker). Most of the research was conducted at Institut d'Etudes Cognitives (ENS), which is supported by grants ANR-10-IDEX-0001-02 PSL*, ANR-10-LABX-0087 IEC and ANR-17-EURE-0017 FrontCog. Experiment 2 was funded by an endowment from Aarhus University to Antoine Marie. We thank Lene Aarøe for feedback on the project at an earlier stage.

Author contributions. All three authors contributed to the design of the studies. A. Marie and H. Trad jointly developed and implemented the surveys in Qualtrics. A. Marie wrote the preregistrations with oversight from H. Trad, collected the data, performed the statistical analyses, reported the results and drafted the initial version of the manuscript. H. Trad contributed to the conceptual framing and, together with B. Strickland, to the writing and revision of multiple sections. All authors reviewed and approved the final version of the manuscript for submission.

Competing interests. The authors declare no conflict of interest.

Data availability statement. Data and code can be downloaded at: https://osf.io/kzpwd/?view_only=80c7d92cebfa451996a017edcbd2f166.

References

- AAAS (2013), *Statement by the AAAS Board of Directors on Labeling of Genetically Modified Foods*, Washington D.C.: American Association for the Advancement of Science (AAAS) [15 January 2025]. <https://www.aaas.org/news/statement-aaas-board-directors-labeling-genetically-modified-foods>.
- Aarøe, L. and M. B. Petersen (2014), 'Crowding out culture: scandinavians and americans agree on social welfare in the face of deservingness cues', *Journal of Politics*, **76**(3): 684–697. doi:[10.1017/S002238161400019X](https://doi.org/10.1017/S002238161400019X).
- Aarøe, L. and M. B. Petersen (2018), 'Cognitive biases and communication strength in social networks: the case of episodic frames', *British Journal of Political Science*, 1–21. doi:[10.1017/S0007123418000273](https://doi.org/10.1017/S0007123418000273).
- Baron, J. (1994), 'Nonconsequentialist decisions', *Behavioral and Brain Sciences*, **17**(1): 1–10. doi:[10.1017/s0140525x0003301x](https://doi.org/10.1017/s0140525x0003301x).
- Baron, J. and M. Spranca (1997), 'Protected values', *Organizational Behavior and Human Decision Processes*, **70**(1): 1–16. doi:[10.1006/obhd.1997.2690](https://doi.org/10.1006/obhd.1997.2690).
- Baumard, N., J. B. André and D. Sperber (2013), 'A mutualistic approach to morality: the evolution of fairness by partner choice', *Behavioral and Brain Sciences*, **36**(1): 59–122. doi:[10.1017/S0140525X11002202](https://doi.org/10.1017/S0140525X11002202).
- Berman, J. Z., A. Barasch, E. E. Levine and D. A. Small (2018), 'Impediments to effective altruism: the role of subjective preferences in charitable giving', *Psychological Science*, **29**(5): 834–844. doi:[10.1177/0956797617747648](https://doi.org/10.1177/0956797617747648).
- Blancke, S., F. Van Breusegem, G. De Jaeger, J. Braeckman and M. Van Montagu (2015), 'Fatal attraction: the intuitive appeal of GMO opposition', *Trends in Plant Science*, **20**(7): 414–418. doi:[10.1016/j.tplants.2015.03.011](https://doi.org/10.1016/j.tplants.2015.03.011).
- Bloom, P. (2017), 'Empathy and its discontents', *Trends in Cognitive Sciences*, **21**(1): 24–31. doi:[10.1016/j.tics.2016.11.004](https://doi.org/10.1016/j.tics.2016.11.004).
- Bonny, S. (2003), 'Why are most Europeans opposed to GMOs?: factors explaining rejection in France and Europe', *Electronic Journal of Biotechnology*, **6**: 7–8. doi:[10.4067/S0717-34582003000100008](https://doi.org/10.4067/S0717-34582003000100008).
- Boyer, P. and M. B. Petersen (2018), 'Folk-economic beliefs: an evolutionary cognitive model', *Behavioral and Brain Sciences*, **41**(e158): 1–51. doi:[10.1017/S0140525X17001960](https://doi.org/10.1017/S0140525X17001960).
- Caviola, L., S. Schubert, E. Teperman, D. Moss, S. Greenberg and N. S. Faber (2020), 'Donors vastly underestimate differences in charities' effectiveness', *Judgment and Decision Making*, **15**(4): 509–516. doi:[10.1017/S1930297500007452](https://doi.org/10.1017/S1930297500007452).
- Chong, D. and J. N. Druckman (2007), 'Framing public opinion in competitive democracies', *American Political Science Review*, **101**(4): 637–655. doi:[10.1017/S0003055407070554](https://doi.org/10.1017/S0003055407070554).
- Cleveland, W. S. and R. McGill (1984), 'Graphical perception: theory, experimentation, and application to the development of graphical methods', *Journal of the American Statistical Association*, **79**(387): 531–554. doi:[10.2307/2288400](https://doi.org/10.2307/2288400).
- Fama, E. F. and M. C. Jensen (1998). Separation of Ownership and Control (SSRN Scholarly Paper No. 94034). Social Science Research Network. doi:[10.2139/ssrn.94034](https://doi.org/10.2139/ssrn.94034).
- Gilbert, R. (2006), 'Looking for Mr. Schumpeter: where are we in the competition–innovation debate?', *Innovation Policy and the Economy*, **6**: 159–215. doi:[10.7551/mitpress/3787.003.0007](https://doi.org/10.7551/mitpress/3787.003.0007).
- Greene, J. (2014), *Moral tribes: emotion, reason and the gap between us and them*, London: Atlantic Books Ltd.
- Greene, J. D., F. A. Cushman, L. E. Stewart, K. Lowenberg, L. E. Nystrom and J. D. Cohen (2009), 'Pushing moral buttons: the interaction between personal force and intention in moral judgment', *Cognition*, **111**(3): 364–371. doi:[10.1016/j.cognition.2009.02.001](https://doi.org/10.1016/j.cognition.2009.02.001).
- Haidt, J. (2001), 'The emotional dog and its rational tail: a social intuitionist approach to moral judgment', *Psychological Review*, **108**(1): 83–95. doi:[10.1037/0033-295X](https://doi.org/10.1037/0033-295X).
- Hall, B. H. and N. Rosenberg (2010), *Handbook of the Economics of Innovation*, Amsterdam, The Netherlands: Elsevier, North-Holland. https://books.google.com/books?hl=en&lr=&id=4nZTCD_zjN4C&oi=fnd&pg=PP1&dq=handbook+innovation&ots=5YEdwBC3bD&sig=KXsXGAlaKwLaeadhKvqnPDBa4Ls.

- Haszeldine, R. S. (2009), 'Carbon capture and storage: how green can black be?', *Science*, **325**(5948): 1647–1652. doi:[10.1126/science.1172246](https://doi.org/10.1126/science.1172246).
- Herzog, S. M. and R. Hertwig (2025), 'Boosting: empowering Citizens with Behavioral Science', *Annual Review of Psychology*, **76**(1): 851–881. doi:[10.1146/annurev-psych-020924-124753](https://doi.org/10.1146/annurev-psych-020924-124753).
- Irsn, B. (2017). Baromètre IRSN.
- Jensen, M. C. and W. H. Meckling (1976), 'Theory of the firm: managerial behavior, agency costs and ownership structure', *Journal of Financial Economics*, **3**(4): 305–360. doi:[10.1016/0304-405X\(76\)90026-X](https://doi.org/10.1016/0304-405X(76)90026-X).
- Katz, E. and P. F. Lazarsfeld (1955), *Personal Influence*, New York: The Free Press.
- Druckman J. N. Keren, G. (ed.). (2010), 'What's it all about?: framing in political science', *Perspectives on Framing*, New York: Psychology Press 1–344.
- Knobe, J. (2003), 'Intentional Action and Side Effects in Ordinary Language', *Analysis*, **63**(3): 190–194. <https://www.jstor.org/stable/3329308>.
- Lin-Healy, F. and D. A. Small (2013), 'Nice guys finish last and guys in last are nice: the clash between doing well and doing good', *Social Psychological and Personality Science*, **4**(6): 692–698. doi:[10.1177/1948550613476308](https://doi.org/10.1177/1948550613476308).
- Lippmann, W. (1929), *Public Opinion: By Walter Lippmann*, New York: Macmillan Company.
- Marie, A., S. Altay and B. Strickland (2020), 'The cognitive foundations of misinformation on science', *EMBO Reports*, **21**. doi:[10.15252/embr.202050205](https://doi.org/10.15252/embr.202050205).
- McCubbins, M. D., R. G. Noll and B. R. Weingast (1987), 'Administrative procedures as instruments of political control', *The Journal of Law, Economics, and Organization*, **3**(2): 243–277. <https://academic.oup.com/jleo/article-abstract/3/2/243/859002>.
- Medani, K. R., A. Neill, G. Garrod, M. Ojo and C. Hubbard (2024), 'Societal perceptions and attitudes towards genetically modified (GM) crops, feed, and food products in the Middle East, North Africa, and Turkey (MENAT) region: a systematic literature review', *Food Quality and Preference*, **117**: 105148. doi:[10.1016/j.foodqual.2024.105148](https://doi.org/10.1016/j.foodqual.2024.105148)
- Mercier, H. (2020), *Not Born Yesterday: The Science of Who We Trust and What We Believe*, Princeton: Princeton University Press.
- Otto, D. (2023). Public perception of CO2 capture and storage. Helmholtz Climate Initiative. <https://www.helmholtz-klima.de/en/aktuelles/public-perception-co2-capture-and-storage> [15 January 2025].
- Peters, E., D. Västfjäll, P. Slovic, C. K. Mertz, K. Mazzocco and S. Dickert (2006), 'Numeracy and decision making', *Psychological Science*, **17**(5): 407–413. doi:[10.1111/j.1467-9280.2006.01720.x](https://doi.org/10.1111/j.1467-9280.2006.01720.x).
- Petersen, M. B. and K. Arceneaux (2020), 'An intuitionist theory of argument strength in politics: how intuitive cognitive biases produce universally strong arguments', *Political Psychology*, **41**(6): 1113–1131. doi:[10.1111/pops.12668](https://doi.org/10.1111/pops.12668).
- Pinker, S. (2018), *Enlightenment Now: The Case for Reason, Science, Humanism, and Progress*, New York, NY: Penguin.
- Pinker, S. (2021), 'Enlightenment environmentalism: a humanistic response to climate change', *Journal of Applied Corporate Finance*, **33**(2): 8–23. doi:[10.1111/jacf.12455](https://doi.org/10.1111/jacf.12455).
- Pörtner, H.-O., D. C. Roberts, M. M. B. Tignor, E. S. Poloczanska, K. Mintenbeck, A. Alegría, M. Craig, S. Langsdorf, S. Löschke, V. Möller, A. Okem and B. Rama (Eds). (2022). Climate change 2022: impacts, adaptation and vulnerability. contribution of working group ii to the sixth assessment Report of the Intergovernmental Panel on Climate Change.
- Rosling, H. (2023), *Factfulness*, New York, NY: Flatiron Books.
- Rubin, P. H. (2003), 'Folk Economics', *Southern Economic Journal*, **70**(1): 157–171. doi:[10.1002/j.2325-8012.2003.tb00561.x](https://doi.org/10.1002/j.2325-8012.2003.tb00561.x).
- Rubin, P. H. (2014), 'Emporiophobia (Fear of Markets): cooperation or Competition?*', *Southern Economic Journal*, **80**(4): 875–889. doi:[10.4284/0038-4038-2013.287](https://doi.org/10.4284/0038-4038-2013.287).
- Scott, S. E., Y. Inbar and P. Rozin (2016), 'Evidence for absolute moral opposition to genetically modified food in the united states', *Perspectives on Psychological Science*, **11**(3): 315–324. doi:[10.1177/1745691615621275](https://doi.org/10.1177/1745691615621275).
- Scott, V., S. Gilfillan, N. Marksson, H. Chalmers and R. S. Haszeldine (2013), 'Last chance for carbon capture and storage', *Nature Climate Change*, **3**(2): 105–111. doi:[10.1038/nclimate1695](https://doi.org/10.1038/nclimate1695).
- Skitka, L. J. (2010), 'The Psychology of Moral Conviction', *Social and Personality Psychology Compass*, **4**(4): 267–281. doi:[10.1111/j.1751-9004.2010.00254.x](https://doi.org/10.1111/j.1751-9004.2010.00254.x).

- Skitka, L. J., C. W. Bauman and E. G. Sargis (2005), 'Moral conviction: another contributor to attitude strength or something more?', *Journal of Personality and Social Psychology*, **88**(6): 895–917. doi:[10.1037/0022-3514.88.6.895](https://doi.org/10.1037/0022-3514.88.6.895).
- Slovic, P. (2007), '“If I look at the mass I will never act”: psychic numbing and genocide', *Judgment and Decision Making*, **2**(2): 79–95. doi:[10.1017/S1930297500000061](https://doi.org/10.1017/S1930297500000061).
- Smith, A. (2013), *On the Wealth of Nations*, New York, New York: Simon and Schuster.
- Strøm, K. (2000), 'Delegation and accountability in parliamentary democracies', *European Journal of Political Research*, **37**(3): 261–289. doi:[10.1111/1475-6765.00513](https://doi.org/10.1111/1475-6765.00513).
- Tetlock, P. E. (2002), 'Social functionalist frameworks for judgment and choice: intuitive politicians, theologians, and prosecutors', *Psychological Review*, **109**(3): 451–471. doi:[10.1037/0033-295X.109.3.451](https://doi.org/10.1037/0033-295X.109.3.451).
- Tetlock, P. E. (2003), 'Thinking the unthinkable: sacred values and taboo cognitions', *Trends in Cognitive Sciences*, **7**(7): 320–324. doi:[10.1016/S1364-6613\(03\)00135-9](https://doi.org/10.1016/S1364-6613(03)00135-9).
- Tetlock, P. E., O. V. Kristel, S. B. Elson, M. C. Green and J. S. Lerner (2000), 'The psychology of the unthinkable: taboo trade-offs, forbidden base rates, and heretical counterfactuals', *Journal of Personality and Social Psychology*, **78**(5): 853–870. doi:[10.1037/0022-3514.78.5.853](https://doi.org/10.1037/0022-3514.78.5.853).
- Uhlmann, E. L., D. A. Pizarro and D. Diermeier (2015), 'A person-centered approach to moral judgment', *Perspectives on Psychological Science*, **10**(1): 72–81. doi:[10.1177/1745691614556679](https://doi.org/10.1177/1745691614556679).
- Wikipedia. (2024), Carbon capture and storage. https://en.wikipedia.org/w/index.php?title=Carbon_capture_and_storage&oldid=1254502191 [15 January 2025].
- Wilson, D. and D. Sperber (2006), 'Relevance Theory', *The Handbook of Pragmatics*. L. R. Horn & G. Ward (ed.) Blackwell, 606–632.
- Zitlmann, R. (2023), 'Attitudes towards capitalism in 34 countries on five continents', *Economic Affairs*, **43**(3): 353–371. doi:[10.1111/ecaf.12591](https://doi.org/10.1111/ecaf.12591).