Resource and Energy Economics 75 (2023) 101404

Contents lists available at ScienceDirect

Resource and Energy Economics

journal homepage: www.elsevier.com/locate/ree

Looking beyond time preference: Testing potential causes of low willingness to pay for fuel economy improvements

Féidhlim P. McGowan [a,b,*], Eleanor Denny [b], Peter D. Lunn [b,c]

[a] J.E Cairnes School of Business and Economics, University of Galway, Galway, Ireland
[b] Behavioural Research Unit, Economic and Social Research Institute, Dublin, Ireland
[c] Department of Economics, Trinity College Dublin, Ireland

ARTICLE INFO

Dataset link: https://osf.io/73psm/

Keywords:
Energy-efficiency gap
Experiment
Willingness-to-pay
Underestimation bias
Concentration bias
Discounting

ABSTRACT

Time preferences are considered a leading cause of the energy efficiency gap. We test two cognition-based mechanisms (concentration bias and underestimation bias) which are distinct from time preferences but can produce identical behaviour when costs are paid upfront and benefits are spread over time. We use an experiment that measures willingness-to-pay for an improvement in fuel economy to test the explanatory power of these mechanisms. The sample is large, nationally representative and comprised only of car buyers (n = 2368). The experiment varies between-subjects (i) the payment schedule for the fuel economy improvement, and (ii) the temporal framing of its monetary benefit. We combine the payment schedules and the benefit frames so that the pattern of results predicted by time preferences differs from the pattern predicted by cognitive mechanisms. Results support the preregistered hypotheses: willingness-to-pay increases as the payment schedule becomes more dispersed across time and decreases when the benefit is presented as more disaggregated (i.e. a monthly saving instead of annual or multi-year saving). The findings are consistent with the predictions of the two cognitive mechanisms, which may explain part of the energy-efficiency gap currently attributed to pure time preference.

1. Introduction

Adoption of greener technologies is crucial to avert a climate crisis (IEA, 2018). But incentivizing adoption through Pigouvian taxes, the method preferred by economists, is politically fraught. In circumstances where such taxes are deemed infeasible, there are a number of instruments policy makers can deploy. Minimum efficiency standards, such as the Corporate Average Fuel Economy, can effect improvements to some degree (Wang and Miao, 2021) but there are also numerous unintended – but by now, predictable – consequences, such as the rebound effect, distortions to the resale price of old fuel-efficient vehicles, and the creation of loopholes which car manufacturers exploit (Gillingham et al., 2016; Jacobsen and Van Benthem, 2015; Anderson and Sallee, 2011). Other options include "cash-for-clunkers" programs to increase the average fuel economy of the vehicle fleet, and direct subsidies to the most energy-efficient options. An alternative approach is nudging, namely changing the choice architecture of a decision without changing the financial incentives. Labelling interventions which highlight the long-run monetary savings of purchasing an energy-efficient product are a prime example of an informational nudge (Allcott and Knittel, 2019). However, in order to be effective, informational nudges must be designed to counter the mechanism(s) that underlie *why* existing behaviours occur. There is currently

* Correspondence to: J.E Cairnes School of Business and Economics, University of Galway, Galway, Ireland.
E-mail address: feidhlim.mcgowan@universityofgalway.ie (F.P. McGowan).

https://doi.org/10.1016/j.reseneeco.2023.101404
Received 3 February 2023; Received in revised form 24 July 2023; Accepted 24 September 2023
Available online 4 October 2023
0928-7655/© 2023 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

no consensus on the primary determinants of the energy-efficiency gap. Leading explanations include imperfect information, inattention, bounded-rationality, and short-sightedness or 'myopia'- in other words a high subjective discount rate (Gerarden et al., 2017). This last factor, a high discount rate, was the original explanation for the energy paradox (Hausman, 1979). However, decisions that appear to be a manifestation of time preference can have their roots in alternative cognitive mechanisms. We describe and test two such mechanisms, which are distinct from time preferences and also independent of simple inattention. Crucially, both mechanisms can produce behaviour that appears identical to present-biased choices.

The first mechanism we test is *concentration bias*, which is a feature of the focusing model of economic choice (Kőszegi and Szeidl, 2013). Concentration bias describes how in a multi-attribute choice setting, greater utility weight is given to attributes that differ more between options, because larger differences are more salient. According to concentration bias, incremental savings on fuel costs are underweighted because the magnitude of fuel cost differences is much smaller than the magnitude of the upfront price difference. People focus on the larger difference and this distorts the expected utility calculation. The cornerstone of this model is limited attention, which is distinct from simple inattention. Notably, there is no misperception of the actual fuel costs of each option over time. Allcott (2016) suggested testing concentration bias as a potential cause of the energy paradox in a review of the energy efficiency literature.[1] This paper, to the best of our knowledge, is the first to implement this suggestion.

The second alternative to time preference that we test is a systematic tendency to underestimate the sum of numerical sequences when people approximate the total intuitively. In other words, when people approximate a total in their heads, either by multiplication or addition, that estimated total tends to be less than the true sum on average. Intuitive summation generating an *underestimation bias* is a recent finding in psychology (Scheibehenne, 2019). A compressive internal scaling of numerical quantities is the most likely cause of this bias (Dehaene et al., 2008), which is only beginning to receive attention from economists. Recent work indicates that underestimation bias becomes stronger as the sequence to be summed increases in length (McGowan et al., 2022). This suggests that underestimation bias might be more pronounced for consumer durables like cars, where savings from investing in greener technology accumulate over a long period of time.

To illustrate how genuine discounting and underestimation bias can cause the same economic behaviour, consider the following example: a car buyer is choosing between two cars that are identical in every way, except one is more fuel-efficient. The buyer's expectations for annual mileage and fuel prices are accurate. Her accurate belief is that the more fuel-efficient car would save her €42 per month in fuel costs. She expects to own the car for exactly three years. What is her additional WTP for the better car? Suppose her annual discount rate is five per cent (so $\delta = 0.95$). For simplicity she evaluates savings on the final day of each year. Her net present value is: $\delta(42*12) + \delta^2(42 * 12)+\delta^3(42*12) =$ €1365. Now assume instead that $\delta = 1$ i.e. she cares equally about the future and present. However, she exhibits underestimation bias, and hence her intuited savings undershoot the true total by 10%. Her WTP will be equal to the perceived sum of savings: (€42*36)*0.9 = €1360. Thus, the behavioural impact of the two mechanisms would be indistinguishable. Similarly, a straightforward alternative explanation can be derived from concentration bias, if a price difference of perhaps €1,000 is more salient and hence given greater weight in the utility function than the €42 difference in fuel costs.

We employed a willingness-to-pay (WTP) experiment to test the explanatory power of concentration bias and underestimation bias. The sample was large and nationally representative ($n = 2368$), and all participants had bought a car in the past five years. The hypothetical scenario was very similar to the example directly above. Participants saw two cars, which were identical except for fuel efficiency, and expressed how much extra they would pay for the more fuel efficient car (hereafter 'better' car for short). The scenario was leasing a plug-in hybrid vehicle for three years. The payment schedule for the better car was varied between subjects, and so too was the temporal framing of the fuel cost difference between the cars. This variation is at the core of the design to separate valuations that result from application of time-preferences from those that derive from the cognitive limitations of boundedly-rational individuals. In addition, in order to test predictions and posited causes of concentration bias, we varied how information was presented on the label along two dimensions. These label manipulations were orthogonal to the primary cost–benefit manipulation designed to tease apart time preferences from the cognitive mechanisms as drivers of low WTP. One label manipulation varied the *accessibility* of fuel cost information (Kahneman, 2003). The other varied whether the fuel saving from the more efficient car was bundled or dispersed across different types of journey. This manipulation tested the concentration bias prediction that the bundled or more 'concentrated' saving should be valued more highly.

Results supported the primary hypotheses. We hypothesized that the more disaggregated payment schedules would increase WTP. Point estimates indicate WTP being 35% higher in the monthly payment condition compared to scenarios where participants make an additional once-off payment or additional annual payment for the better car. The significant difference is robust to model specification. There was no significant difference in WTP between annual and once-off payment frames. Note that models of time preference predict the opposite, namely that an increase in WTP should be largest between the once-off payment and the annual payment conditions, with a smaller gap between annual and monthly payment conditions. As for the manipulation of temporal framing of fuel cost differences, results supported the hypothesis that framing over a longer time period would increase WTP. The point estimate on the main model indicated a 13% increase in WTP in the annual and three-year fuel cost frames relative to the monthly frame. This suggests that when people only see the monthly fuel cost, they underestimate accumulated fuel savings. Note that under a pure time-preference explanation, this framing manipulation should have no effect on WTP, because it does not affect when savings are realized, only the time frame over which they are aggregated.

[1] "Could bias towards concentration be relevant for some energy efficiency decisions? Because the benefits of energy efficiency come in small flows every time an electricity bill is paid or the gas tank is filled, this is a natural application of Koszegi & Szeidl's (2013) model. Even a lab experiment framed in an energy use context would be interesting if it measured and distinguished bias towards concentration from other potential mistakes". (Allcott, 2016, p. 27).

These findings have implications for policy. On a practical level, policymakers seek ways to make adoption of energy-efficient technologies as painless as possible, both in monetary and psychological terms. Creating opportunities for incremental repayment might be one way to achieve that. However, this issue requires further research, as payment schedules are a dimension of the purchase decision that have received relatively less attention in energy-efficiency research. This may be because payment schedule differences do not fit neatly within the 'information gap' paradigm.

The rest of the paper proceeds as follows. Section 2 reviews the most relevant literature. Section 3 describes the experiment design, hypotheses and procedure. Section 4 presents the results. Section 5 discusses the findings, proposes directions for future research, and concludes.

2. Related literature and decision framework

This section first provides background context on the range of factors that plausibly contribute to the energy-efficiency gap, then summarizes the most relevant literature on (i) the mechanisms we test and (ii) the effect of labels when people value energy-efficiency of consumer durables.

2.1. Leading causes of the energy-efficiency gap

Numerous causes have been forwarded as contributors to the energy-efficiency gap. Apart from the aforementioned time preferences and cognitive biases, other factors are: limited information, financial constraints, attention deficits, preferences for other appliance attributes, principal–agent problems,[2] uncertainty over future energy prices (and hence likely savings), and regulations that distort the prices people face. These causes have been categorized into market failures and "behavioural failures" (Gillingham and Palmer, 2020; Allcott, 2016), where the latter involve departures from the rational agent model.

A leading framework for categorizing such departures is the DellaVigna (2009) taxonomy of non-standard preferences, non-standard beliefs and non-standard decision-making. Regarding non-standard preferences, time-inconsistent preferences that arise from a quasi-hyperbolic discount function (Laibson, 1997) are considered relevant to the energy-efficiency gap, because they capture the self-control problem between wanting to reduce energy usage in the future but being unwilling to incur some sacrifice in the present (Bradford et al., 2017). Another non-standard preference is loss aversion, which has been posited as a reason why positive net-present-value investments might be willingly foregone: If consumers are loss averse, and future energy prices are uncertain, they may weigh the potential negative payoffs heavily enough to be deterred from purchase, even though they know the investment would probably have positive net benefits (Gillingham and Palmer, 2020). However, although theoretically plausible, the empirical evidence for loss aversion as a key wedge in the energy efficiency gap is weak (Gerarden et al., 2017).

There is evidence that beliefs over energy usage may be systematically misaligned. Attari et al. (2010) showed that consumers display a 'central tendency' in their perceptions of energy use, with severe underestimation of the energy usage from intensive sources such as heating and cooling, and minor overestimation for low-intensity products like a light bulb. Larrick and Soll (2008) documented the concept of the miles-per-gallon (MPG) illusion: people's intuition is that fuel costs scale linearly in MPG, which is not the case, as the relationship is strongly non-linear. The EU standard of "litres per 100 km" provides consumers with a linear relationship, though people may still underestimate fuel cost differences between vehicles if they underestimate their mileage.

The category of non-standard decision making includes all quirks, mistakes, shortcuts, misconceptions and other idiosyncrasies that an optimizing rational agent would not display, but may impact outcomes in markets populated by boundedly-rational consumers. Two features of this category, namely inattention and salience effects, are of primary importance for the energy-efficiency gap according to Gerarden et al. (2017). For instance, Turrentine and Kurani (2007) interviewed recent car buyers and found that many were inattentive to fuel costs when buying a car; others considered it in a simple way that did not reflect calculation of net-present value. Labelling interventions, reviewed below, attempt to increase the salience of fuel costs within the multiattribute choice context of vehicle purchase.

2.2. Experimental evidence on labelling and energy valuation

Undervaluation of energy efficiency is inferred from heightened responsiveness by consumers to a change in upfront costs relative to a change in future costs. This is often referred to as myopia, which in its usage covers several departures from the standard model — imperfect information, inattention, deliberation costs, high discount rates, and so on (Silvi and Rosa (2021) provide a comprehensive guide). There is a lack of consensus on the degree to which consumers are myopic (Greene, 2010). Two recent labelling experiments illustrate how this issue remains contested. First, Gillingham et al. (2021) leveraged a change in fuel labels, after the fuel economy of two popular small cars had been overestimated and subsequently corrected. The observed change in demand in this natural experiment implied that consumers were indifferent between $1 in discounted fuel costs and 15–38 cents in the vehicle purchase price (assuming discounting at 4%), which is stronger undervaluation than typically found by exploiting changes in fuel prices (Allcott and Wozny, 2014). In contrast, no undervaluation was recorded in a large-scale labelling intervention experiment (in both the field and online) that drew buyers' attention to fuel economy (Allcott and Knittel, 2019).

[2] And other instances of misaligned incentives which are not classic principal–agent problems, for example when a landlord has no reason to improve the insulation of a property because the heating bill is paid by the tenant.

The paper concluded that imperfect information and inattention do not have a significant systematic effect on vehicle markets after finding precise null effects for treatments. The authors drew the implication that fuel economy receives a low decision weight rather than being undervalued. However, Allcott and Knittel (2019) also entertained the possibility that the labels were not effective in *actually* making consumers factor in fuel costs to their decision. The labels contextualized the lifetime monetary savings by giving the equivalent quantity of clothes or airline tickets that could be purchased with the saving.[3] This aspect of the label may have distracted some participants from comparing the saving against the upfront price difference.

The null finding in a real-world decision stands in contrast to the wider literature on energy efficiency labelling interventions for consumer durables. Generally, WTP is higher for energy-efficient appliances when savings are framed in monetary terms, rather than in terms of energy output (Newell and Siikamäki, 2014; Min et al., 2014; Andor et al., 2020). WTP is also often higher when lifetime energy costs are shown on the label (Heinzle, 2012) but the field evidence suggests the success of this intervention may depend on the complementarity of staff training quality (Kallbekken et al., 2013; Carroll et al., 2016).

The unobserved nature of preferences makes it difficult to determine the role cognitive limitations play in valuations. One way to cut through this complexity is to design tasks with no role for preferences. This approach was taken in an online experiment that presented participants with similar appliances and asked them to choose the one that minimized lifetime cost (Blasch et al., 2019). The experiment varied whether annual energy consumption was displayed in monetary terms or physical units. The correct answer was chosen more often in the monetary label condition. However, this relative difference masked a wider trend — over half of participants did not report trying to solve any optimization problem. This speaks to the central role for aspects of bounded rationality.

2.3. Concentration bias and underestimation bias

At the core of concentration bias is the idea that larger differences attract more attention than smaller differences. Due to this increased attention, attributes that have larger magnitude differences receive more decision weight. Concentration bias could explain the disparity between the null effect for fuel economy labels in Allcott and Knittel (2019) and the positive effect sizes for other consumer durables as a consequence of the much higher upfront price of vehicles relative to appliances. For a given proportional difference, this higher upfront price translates to a larger absolute difference in price, thereby attracting more attention and greater decision weight.

It is important to emphasize that while concentration bias and present-focused time preference may sometimes generate similar predicted outcomes, they are distinct mechanisms. For example, the utility value of a large benefit in the future – such as a promotion or sizable bonus – would be diminished by discounting, but attract focus-weighted utility under concentration bias. Many dispersed costs on the road to that attention-grabbing prize would be underweighted under the focusing model. However, in the case of pay-now, benefit-over-time investments which are typical of the energy paradox, concentration bias pushes choices in the same direction as present-focused time preferences.

The empirical evidence for concentration bias is presently mixed. Concentration bias predicted labour supply decisions in a real-effort task (Dertwinkel-Kalt et al., 2022). The effect was mediated by the 'accessibility' of the difference in workload between the two options. However, no evidence for concentration bias was found in a field experiment in Kenya on purchasing an energy efficient heat-stove (Berkouwer and Dean, 2022). More recently, in a consumer choice experiment, a model of relative thinking was found to better fit the data than concentration bias (Somerville, 2022).

By contrast, the empirical evidence for the second cognitive bias we study here is more consistent. Underestimation when intuitively summing multiple amounts has been recorded in a laboratory experiment and a field experiment in supermarkets (Scheibehenne, 2019). Others experiments have found that it partially explains undervaluation of lotteries relative to expected payoffs, which is typically attributed only to risk aversion (Olschewski et al., 2021). Studies that investigate the related concept of intuitive averaging also record a tendency towards underestimation, which is consistent with a compressive number line (Brezis et al., 2015). In recent research, McGowan et al. (2022) found that underestimation bias is present when participants sum sequences that are economically meaningful and familiar, such as utility bills. They also found the bias was present across elicitation methods as it occurred in both a judgement task, where the best guess for the total is typed, and a forced-choice task, where no number is generated.

3. The experiment

The introduction and literature review highlighted that not purchasing energy-efficient products, which is typically attributed to discounting, could in theory be caused by cognitive mechanisms of concentration bias or underestimation bias. We designed an experiment to test the predictive power of these mechanisms. The payment schedule for the energy efficient investment, and the temporal framing of its monetary benefit, were varied across conditions so that the WTP pattern predicted by leading models of time preference diverges from the prediction of the two cognitive mechanisms. As the two cognitive mechanisms are not mutually exclusive, a test to measure their relative predictive power was also preregistered, but this proved inconclusive. The variation in payment schedule and temporal framing was designed so that each mechanism made a different prediction regarding the relative size of WTP gaps between pairs of conditions. The details of the separating test are described in Section 6.2 of the Appendix.

[3] For example: "A Ford Fiesta FWD will save you $8,689 over its lifetime compared to a Ford Crown Victoria Ffv. That is the same as it would cost for: 17.4 iPads, 8.7 tickets to Hawaii, 174 pairs of Levi's jeans". (Allcott and Knittel, 2019, Figure 1, p. 8)

The experiment scenario involved leasing a plug-in hybrid car on a three-year contract. The lease contract took the form of an upfront deposit and 36 monthly payments.[4] The participant was shown two cars, labelled A and B, which were identical except Car B had better fuel efficiency. After viewing the 'personalized projected fuel cost' label for the two cars, participants inputted their additional willingness-to-pay for Car B. Participants did this twice, with two pairs of cars that had a similar difference in fuel costs: the annualized saving was €252 in one choice set, and €324 in the other. The order of presentation was counterbalanced. This is the basic overview - design and procedural details are described below.

3.1. Participants

The sample was recruited by a market research company to be representative of the population of car buyers in Ireland (n = 2368). All participants had either bought a new car in the previous five years (n = 2000) or were thinking about buying a new car (n = 368). The socio-demographic breakdown of the sample was as follows: Exactly half was female, 33% were in the 18–34 age bracket, 42% were aged 35–54, and 25% were aged 55+. In terms of educational attainment, 53% had at least a higher degree. This is in line with the national average. The household income distribution was also nationally representative. Each participant was randomly allocated to an experimental condition. There was no substantial imbalance across conditions on the above demographic variables.[5]

Power analysis (included in the preregistration) found that a minimum sample size of 65 was needed to detect a medium sized effect ($d = 0.5$) for a difference in mean willingness-to-pay between two groups. We initially planned for 2000 participants (83.33 per arm), which would have provided sufficient power. The final sample size of 2368 surpassed these minimum power requirements. Participants were paid to take part in this experiment and a preceding survey.

3.2. Design and hypotheses

The design, hypotheses and analysis plan were preregistered on the Open Science Framework website after ethical approval had been granted by the ethics committee at the authors' institution, and before data collection began.

3.2.1. Payment schedule — fuel cost savings combinations

The following factors were varied between-subjects: (i) The payment schedule for making the additional payment for Car B - higher monthly payment, additional payment annually, or once-off three-year payment (i.e. additional deposit payment) and (ii) the temporal framing of fuel cost information — monthly, annual, or three-year (i.e. the duration of the lease contract). Balanced combinations (e.g. Monthly WTP - Monthly Fuel Costs) were excluded because in these circumstances the WTP elicitation could be reduced to a simple matching task (i.e. after calculating the saving per month, offering the same as monthly WTP). In contrast, imposing imbalance allows for different focus-weights (concentration bias) and bias in intuitive summation (underestimation bias) to alter WTP. The six combinations are shown in Table 3.1 below.

Hypotheses The primary hypotheses were as follows:
H1: WTP will be decreasing in the temporal aggregation of the payment schedule.
H2: WTP will be increasing in the temporal aggregation of the fuel cost frame.

H1 is predicted by both time discounting and the two cognitive mechanisms, underestimation bias and concentration bias. H2 is predicted only by the cognitive mechanisms because framing does not affect when savings are realized; time discounting as a parsimonious model assumes descriptive invariance.

WTP gap predictions

Concentration bias and underestimation bias make different predictions to exponential discounting regarding how WTP will change across the payment frames, controlling for how the savings are described.

Concentration bias and underestimation bias predict, that (controlling for how savings are framed):

$WTP_{monthly}$ - WTP_{annual} > WTP_{annual} - $WTP_{deposit}$

because the difference in disaggregation is greater in the former (36 payments vs. 3 payments).

Exponential discounting makes the opposite prediction:

$WTP_{monthly}$ - WTP_{annual} < WTP_{annual} - $WTP_{deposit}$

This prediction stems from the fact that the stream of payments is discounted in a linear fashion. When the entire payment has to be paid immediately, the nominal value must equal the present-discounted value of the stream of savings. With a 5% annual discount rate, the nominal stream of savings of €972 over 3 years is worth €897 today, or €299 annualized. When the payments are spread over time, they are subject to the same discounting applied to the stream of savings. If the annual payment is made on the first day of the each year, so that initial payment is not discounted (like the deposit), then the formula for the annual payment (A) to match the present discounted value of the savings is: $A + \delta A + \delta^2 A = €897$. If the payment is made on the last day of the

[4] The lease scenario was chosen over a purchase scenario so that different beliefs about how long respondents would own the car would not be a factor in the WTP responses.

[5] No significant difference between groups on gender or education. On age, the only significant difference was that Group B had a lower proportion of 18–34 year olds (28%) than Group A (34%) ($p = 0.027$).

Table 3.1
Combinations of payment schedule and saving frame and orthogonal label manipulations.

Group	Payment schedule	Saving frame	Label type	N
A	Monthly	3-Year	LD (Low-Accessibility, Dispersed)	93
	Monthly	3-Year	HD (High-Accessibility, Dispersed)	91
	Monthly	3-Year	LC (Low-Accessibility, Concentrated)	101
	Monthly	3-Year	HC (High-Accessibility, Concentrated)	102
B	Monthly	Annual	LD	97
	Monthly	Annual	HD	98
	Monthly	Annual	LC	97
	Monthly	Annual	HC	99
C	Annual	3-Year	LD	108
	Annual	3-Year	HD	95
	Annual	3-Year	LC	101
	Annual	3-Year	HC	102
D	Annual	Monthly	LD	103
	Annual	Monthly	HD	95
	Annual	Monthly	LC	100
	Annual	Monthly	HC	93
E	Deposit	Annual	LD	105
	Deposit	Annual	HD	99
	Deposit	Annual	LC	98
	Deposit	Annual	HC	96
F	Deposit	Monthly	LD	104
	Deposit	Monthly	HD	97
	Deposit	Monthly	LC	98
	Deposit	Monthly	HC	96

year, the formula is $\delta A + \delta^2 A + \delta^3 A = €897$. This gives a range for A of €314 to €331. For the monthly payment, the equivalent annualized WTP is €322.

Note that applying quasi-hyperbolic preferences instead of exponential discounting will reduce all WTPs, but have a greatest *relative* effect on the deposit payment[6] and hence increase the expected size of the gap in the time preference prediction above.

3.3 Separating test between cognitive mechanisms

The separating test between concentration bias and underestimation bias is described below. First a common prediction is described. The three WTP pairs, where each combination of payment schedule and savings frame are inverses of each other, are AF, BD and CE. Within each pair, the first condition (A,B,C respectively) is expected to have the higher WTP than its inverse because its payment schedule is more disaggregated and savings frame more aggregated. The size of the gap can be measured by taking the WTP difference. Both concentration bias and underestimation bias predict that the AF gap should be largest, next the BD gap, and finally the CE gap. Concentration bias makes this prediction because the focus-weight ratio is greatest for the AF pair, namely 36:1, smaller for the BD pair (12:1), and smallest for the CE pair (3:1). Underestimation bias predicts this order because the strength of the bias is expected to be increasing in the length of the sequence to be summed. The disaggregated component is most granular (longest) in the AF pair and least granular (shortest) in the CE pair relative to the other component.

The test of relative predictive power relies on the fact that the strength of concentration bias depends on a ratio, whereas the strength of underestimation bias depends on an arithmetic difference. Concentration bias predicts that the intermediate BD gap should be closer in absolute size to the AF gap than the CE gap, because $36/12 < 12/3$. In contrast, underestimation bias predicts the BD gap should be closer in size to the CE gap, because $36\text{-}12 > 12\text{-}3$. This test is shown graphically in Fig. 1 below.

3.3.1 Label manipulations

Testing between concentration bias and underestimation bias is difficult, because both mechanisms make essentially the same predictions at an ordinal level. However, there are concrete predictions of concentration bias, and posited explanations for it, that can be tested orthogonal to the primary experimental manipulation above. First, we varied the *accessibility* of the fuel cost difference between the two cars. Accessibility is a term for the ease (or effort) with which particular mental contents come to mind: "Highly accessible features influence decisions, whereas features of low accessibility are largely ignored". (Kahneman 2003, pg 11). Dertwinkel-Kalt et al. (2022) found that accessibility mediated concentration bias. To vary accessibility, on half of the labels the fuel cost difference between the cars was implicit. On the other half, the precise difference was given; in other words, the arithmetic was performed for participants.

[6] For instance, a β of 0.7 would lead to a WTP of approximately €630 instead of €897.

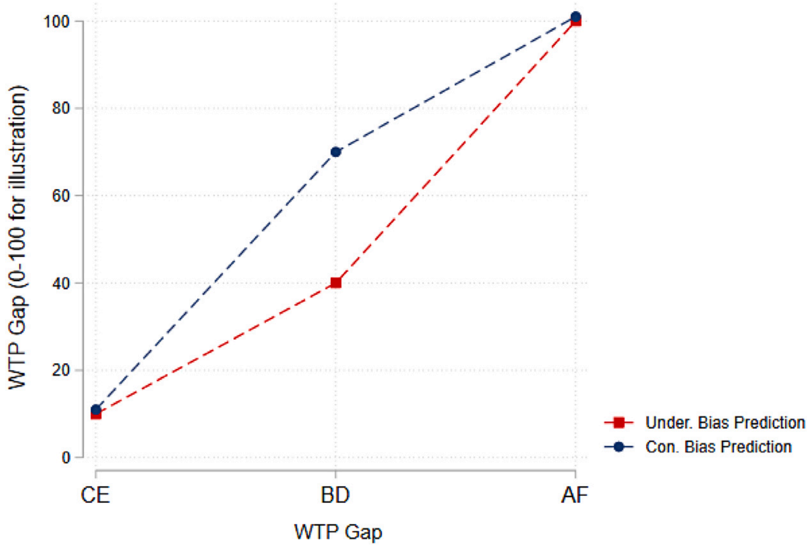

Fig. 1. Graphical Illustration of opposing predictions of Concentration Bias and Underestimation Bias regarding size of intermediate WTP gap BD relative to largest gap, AF, and smallest gap, CE.

Annual Running Cost
by Journey Type for A and B

Short Drives (less than 30mins)		Medium Drives (30mins – 1hr)		Longer Drives (1hr+)	
A	B	A	B	A	B
€252	€192	€432	€348	€576	€468

Annual Running Cost
by Journey Type for A and B

Short Drives (less than 30mins)		Medium Drives (30mins – 1hr)		Longer Drives (1hr+)	
A	B	A	B	A	B
€192	€192	€324	€324	€744	€492

Fig. 2. Dispersed vs. concentrated fuel cost savings in low accessibility condition. Both savings are €252 per year.

Annual Running Cost
by Journey Type for A and B

	Short Drives (less than 30mins)		Medium Drives (30mins – 1hr)		Longer Drives (1hr+)	
	A	B	A	B	A	B
	€252	€192	€432	€348	€576	€468
Annual Savings from B:	€60		€84		€108	

Annual Running Cost
by Journey Type for A and B

	Short Drives (less than 30mins)		Medium Drives (30mins – 1hr)		Longer Drives (1hr+)	
	A	B	A	B	A	B
	€192	€192	€324	€324	€744	€492
Annual Savings from B:					€252	

Fig. 3. Dispersed vs. concentrated savings in high accessibility condition where cost difference is explicit. Both savings are €252 per year.

The second manipulation was designed to test a prediction in Kőszegi and Szeidl (2013) that concentrated benefits (fuel savings in this setting) should be preferred to dispersed ones (Kőszegi and Szeidl, 2013, p.65). To test this, we varied the dispersion of fuel saving in the personalized fuel cost forecast. On half the labels the saving was dispersed across all journey types. In the other half it was concentrated in a single journey category. The size of monetary saving was equal across all conditions. These manipulations were tested in a 2x2 between-subjects design nested within the six combinations of payment schedule and temporal framing of fuel cost.

Fig. 2 shows the dispersed saving (left) and concentrated saving (right) in the low accessibility condition, where arithmetic must be performed to figure out the precise saving. Concentration bias predicts that WTP will be higher in the concentrated saving condition. Underestimation bias makes no clear prediction, but note that less arithmetic is required in the concentrated saving frame, which conceivably leaves more cognitive bandwidth for performing intuitive estimation. However, if the total difference

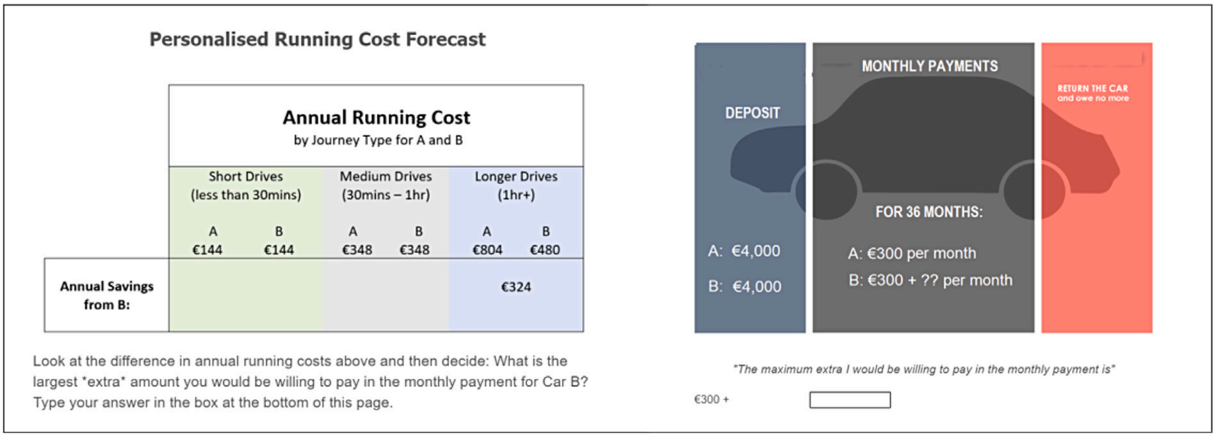

Fig. 4. WTP figure inputted after looking at fuel forecast.

in the dispersed condition is estimated by looking at the difference in leading digits, then WTP may be higher in the dispersed condition. Fig. 3 above shows the same labels except in the high accessibility treatment. If concentration bias is mediated by accessibility (Dertwinkel-Kalt et al., 2022) concentration bias should be reduced in this treatment arm.

3.3.2 Procedure

The experiment was conducted online using the Qualtrics platform. On the introduction page, the hypothetical scenario was explained; participants were told to imagine they were about to lease a plug-in hybrid car on a three-year contract. The contract took the form of an upfront deposit and 36 monthly payments. They were then shown the two options, A and B, which were identical except Car B had better fuel economy. Participants were told that they would be asked to input an additional willingness-to-pay for Car B but first they would be shown fuel cost estimates for both cars.

On the second page, the contract for Car A and B was shown, with a question mark to indicate the additional payment. The deposit was always €4000 and the base monthly payment was €300. Participants were told that after they viewed the fuel cost forecast, they would be asked to input a figure for the question mark to indicate their additional WTP for Car B. They viewed the fuel cost forecast and then inputted their WTP at the bottom of the same screen. These steps (apart from the introduction screen) were then repeated at the second hypothetical car dealership. The order of car pairs was counterbalanced across participants. The main response screen is shown in Fig. 4. Screenshots of other stages are shown in Figures 11–15 in section 6.4 of the Appendix.

4 Results

This section describes how WTP differed across the different conditions. First, descriptive results are presented in Section 4.1. Then in Section 4.2 we present regression results, beginning with the preregistered method of analysis, namely mixed-effect linear regression. The results of the label manipulations are also described in Section 4.2.

4.1 Descriptive results

Recall that across the different payment schedule conditions, participants inputted their additional WTP in (i) the once-off deposit, (ii) in the annual payment, or (iii) in the monthly payment. In order to allow direct comparison between these conditions, an annualized willingness-to-pay was calculated from every response: monthly payment responses were multiplied by twelve, deposit-payment responses divided by three, and annual payment responses left unchanged.

The preregistration stated that implausibly high WTP figures would be excluded. Inputting a WTP that exceeds the monetary saving can reflect additional valuation of the carbon emissions savings.[7] However, this argument has a limit: at some point, a WTP may be considered ‘unreasonably’ high. We applied a maximum cutoff at an annualized WTP of €650. This figure is double the maximum annual monetary saving, and thus allows the monetary value of the utility gain from being more environmentally friendly to equal the utility of the monetary savings (i.e. v = s in the decision framework terminology which is presented in Section 6.1 of the Appendix).[8]

[7] The annual monetary saving from the more fuel efficient car was €252 on one trial or €324. However four of the 48 labels had a typo which meant the actual difference was €312 instead of €324. See Table 6.1 in the Appendix for all combinations and respective mean WTPs.

[8] This is a simplifying assumption; the relative concavity of utility for money vs. carbon emissions savings is an interesting topic which is beyond the scope of this paper.

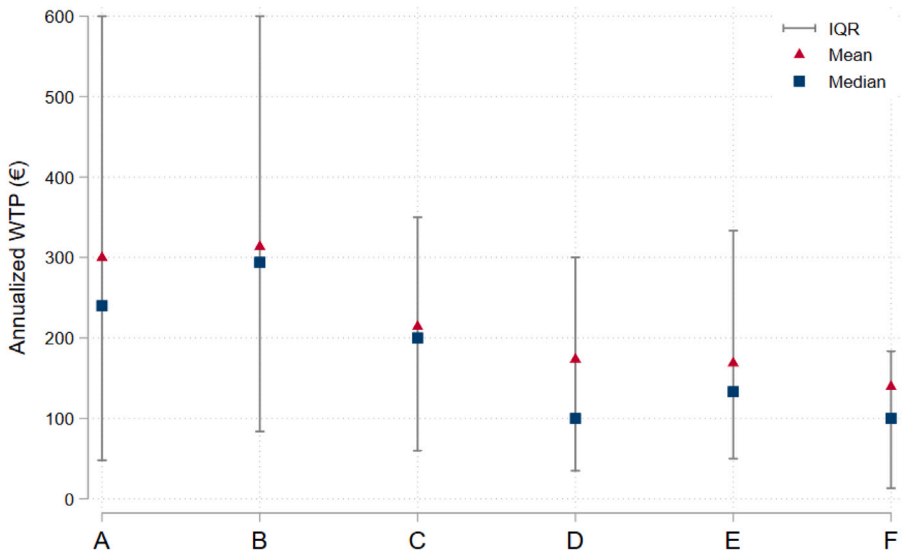

Fig. 5. Mean, Median and Interquartile Range of WTP for each of the six payment schedule-fuel cost frame combinations. Admissible WTP from €0 to €650 inclusive. Conditions A, C and E had more aggregated savings than their payment schedule counterpart.

Fig. 5 above shows the mean, median and interquartile range for annualized WTP for each of the six combinations of payment schedule and savings frame (see Table 3.1 for definition of each group). To test H1 (that WTP would be decreasing in the temporal aggregation of the payment schedule), we compare conditions that had different payment schedules but the same savings frame (A vs. C, B vs. E, D vs. F). For A vs. C, the former had significantly higher average WTP with a difference between conditions of €87 (t = 6.53, df = 981, p < 0.001). Group B, which like A had a monthly payment schedule, recorded significantly higher WTP (by €144) than Group E (t = 11.52, df = 913, p < 0.001). Lastly, the average WTP in Group D was €34 higher than in F (t = 3.77, df = 1206, p = 0.0001). Taken together, these results constitute strong support for H1.

To gauge empirical support for H2 (that WTP will be increasing in the temporal aggregation of the savings frame) we examine the difference in WTP within each payment schedule. In the monthly payment condition, the difference between A and B is not significant (two sample t-test, p = 0.45) and the point estimate is higher for B, in which savings were less aggregated (annually vs. over three years in A). In the annual payment pair, savings were framed over a three–year period (group C) or monthly (group D). Group C recorded significantly higher WTP (two-sample t-test, t = 4.42, p < 0.001). Lastly, in the deposit payment condition, group E, where fuel savings were shown annually, recorded higher WTP than group F, where savings were framed monthly (t = 3.46, p < 0.001). In summary, H2 receives weaker empirical support compared to H1.

In terms of the respective predictions of concentration bias, underestimation bias and discounting for the size of the WTP gaps, the descriptive evidence does not support the discounting prediction that:

WTP_{annual} - $WTP_{deposit} < WTP_{monthly}$ - WTP_{annual}. Instead, the point estimate on the annual-deposit difference in WTP is €38, compared to €114 for the monthly–annual difference. Wald tests of coefficient equality are conducted in the regression analysis below.

4.2 Regression results

Regression analysis was conducted to test whether the significant differences recorded in Section 4.1 are robust to controlling for label types and background characteristics. We also account for the repeated-measures nature of the data (each participant made two responses). The primary analysis omits WTP above €650. We perform multiple sensitivity checks on this decision, the main one being complementary logistic regression analysis, where the binary dependent variable takes a value of one when the inputted response was at least as large as the annual fuel cost saving. As shorthand, these responses will be referred to as *sufficient-WTP*. They could also be conceived of as responses that do not fall foul of the energy paradox. Transforming a continuous variable in this way discards useful variation, but this method ensures results are not driven by a subset of the data near an arbitrary threshold.

The first specification in Table 4.1 is a random-intercept linear regression. This model specification was preregistered. The dependent variable is the annualized WTP. This first specification includes only the categorical variable for payment schedule-fuel savings frame combination. The reference category is A, the condition with the monthly payment schedule and three-year fuel savings. A categorical variable for label type is also included, but these coefficients are omitted as the results of this manipulation are analysed in Table 4.2 (as these manipulations were orthogonal it is highly unlikely their influence drives the main results). The results in Column 1 align with the results illustrated in Fig. 5. There is no significant difference between condition A and B, the two monthly payment conditions. All other combinations record significantly lower WTP than A. Tests of equality for linear

Table 4.1
Columns 1 and 2 display mean annualized WTP by condition. WTP responses above €650 have been removed. Columns 3 and 4 model the predictors of the WTP response equalling the annual fuel savings using a logistic-link function.

	Mixed OLS	Mixed OLS	Logit	Logit
Payment-Savings Combination (ref = A)				
B	15.68	14.60	−0.06	−0.04
	(26.44)	(26.54)	(0.18)	(0.18)
C	−96.34***	−102.04***	−1.53***	−1.56***
	(23.63)	(23.24)	(0.15)	(0.15)
D	−138.91***	−144.44***	−2.06***	−2.08***
	(23.76)	(23.63)	(0.15)	(0.14)
E	−142.11***	−146.74***	−1.25***	−1.25***
	(23.39)	(23.58)	(0.15)	(0.15)
F	−168.90***	−173.57***	−1.33***	−1.32***
	(22.84)	(22.78)	(0.13)	(0.13)
Age 35–54 (ref. = 18–34)		−40.28***		−0.38***
		(8.59)		(0.14)
Age 55+		−42.77***		−0.53***
		(11.43)		(0.15)
Degree Holder		−3.91		−0.05
		(8.69)		(0.08)
Female (ref = not female)		−14.39*		−0.20**
		(8.27)		(0.10)
Constant	328.43***	371.58***	1.40***	1.84***
	(23.15)	(24.69)	(0.12)	(0.13)
Observations	3106	3106	4736	4736
Groups	1645	1645		

Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

combinations of the coefficients can be used to test the opposing WTP gap predictions of the cognitive mechanisms vs. discounting. The test rejects the null of: $(\beta_A + \beta_B)/2$ - $(\beta_C + \beta_D)/2 = (\beta_C + \beta_D)/2 - (\beta_E + \beta_F)/2$ in favour of the alternative that the monthly–annual difference is greater (β_{diff} = €87, $z = 5.27$, 95% C.I = €55,€120). Recall that if discounting was the primary explanation for the variation in WTP, the test of equality would have been rejected in the *opposite* direction.

The second column adds socio-demographic characteristics. The randomization process means one would not expect the addition of these variables to change the pattern of results, but their inclusion allows more precise estimation of the experimental effects of interest. The size of the WTP difference between A and B and the other conditions increases slightly. Notably, the reference age category of 18–34 has a significantly higher WTP than the older age groups (p < 0.01). Female participants have marginally lower WTP (p<0.1).

A logistic link function was used to regress the binary *sufficient-WTP* variable on the experimental categorical variables and demographic factors. The results are shown in Columns 3 and 4 of Table 4.1. Coefficients are shown in log-odds (as there are no continuous variables in the specification, there is no benefit to reporting marginal effects). The pattern of results in Column 3 is qualitatively similar to the mixed linear models. The log-odds of a sufficient-WTP response is significantly lower for all conditions relative to the two monthly payment ones, A and B. However, the order of coefficients in terms of magnitude changes. The annual payment conditions (C and D) have the lowest probability of a sufficient-WTP response. Demographic factors are added in Column 4 but the pattern of results does not change. Again, older age groups are significantly less likely to record a sufficient-WTP response, as are female participants (p < 0.05).

Table 6.2 in the Appendix reports results from a number of robustness tests using different inclusion criteria, the results of which we summarize here: First, excluding zero WTP responses and keeping the limit at €650 annualized-WTP produces minor changes to the results. The magnitude of the WTP difference between the two monthly payment conditions (A and B) and the other combinations becomes larger, with D, E and F having point estimates slightly over €200 below the reference category of A. The age and gender differences remain as in Table 4.1. Columns 2–5 of Table 6.2 use WTP thresholds of €500, €750, €1000 and €1250 respectively. Only when the limit is €500 is there a qualitative change in results. In this specification, only combination F (deposit payment, monthly saving frame) has a significantly lower WTP than A (p < 0.01). The change in the pattern of coefficients is due the fact that with a threshold of €500, the frequent response of €50 per month is excluded (13% of respondents in the monthly payment condition gave this response). The pattern of coefficients in Columns 3–5 matches the main specification, although the magnitudes of the coefficients generally increase. In all five specifications, the 35–54 and 55+ age groups record significantly lower WTPs (p < 0.01).

Table 4.2
Mixed linear models. Dependent variable is Annualized WTP. The six conditions (A–F) have been pooled by savings frame in Column 1, and by payment schedule in Column 2.

	(1)	(2)
More Disaggregated Payment	88.47***	
	(9.41)	
More Aggregated Savings Frame		24.16***
		(8.56)
Annual Saving Frame (ref = monthly)	87.22***	
	(10.70)	
Three-Year Savings Frame	108.94***	
	(10.75)	
Annual Payment (ref = monthly)		−126.33***
		(13.42)
Deposit Payment		−163.96***
		(13.20)
Labels (ref = Low Accessibility Dispersed Saving)		
High-Accessibility Dispersed Saving	−5.45	−6.27
	(12.24)	(12.08)
Low-Accessibility Concentrated Saving	−8.10	−9.94
	(12.16)	(12.06)
High-Accessibility Concentrated Saving	−29.66**	−32.85***
	(12.22)	(12.19)
Constant	124.73***	325.23***
	(11.07)	(14.79)
Observations	3106	3106
Groups	1645	1645

Robust Standard errors in parentheses.*$p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 4.2 presents mixed-effects linear regressions to summarize the separate effect of the payment schedule and savings frame manipulations. Two dummy variables were created: first, whether a condition had a more disaggregated payment structure, given its savings frame (Combinations A,B,D = 1; C,E,F = 0); and second, whether its savings frame was more temporally aggregated, given its payment structure (Combinations A,C,E = 1; B,D,F = 0). The first column shows that on average, the increase in WTP from a more disaggregated payment schedule is 88.47 (p < 0.001). This effect size is equivalent to a 0.48 standard deviation change in WTP. This specification also includes a categorical variable for savings frame, with monthly savings as the reference category. The effect of more aggregated savings is monotonic but concave, with point estimates of €87 and €108 for the Annual and Three Year frames respectively. The difference between these frames is marginally significant ($p = 0.062$). This can be interpreted as robust support for H1. In the second column, the savings dummy shows a coefficient of €24.16 (p < 0.001), which is equivalent a 0.13 standard deviation change in WTP. In logistic regression models (not reported) using the same explanatory variables, the saving manipulation is approximately three-fifths the effect size of the payment schedule manipulation. Taken together, these results indicate that both H1 and H2 are supported, with the former having a larger effect size. Figure 7 in the Appendix illustrates the sensitivity of the empirical support for H1 and H2 to the choice of WTP threshold. The effect size for H1 is larger but more volatile. The savings frame effect size (H2) is stable across a wide range of thresholds. Figure 8 in the Appendix breaks down this effect by specific condition (coefficients on annual and three-year conditions relative to monthly reference category).

Table 4.2 also shows how WTP changed across the orthogonal label manipulations, which were described in Section 3.2.2.[9] In the reference case, the saving was dispersed across three categories of journey and the fuel cost difference was not calculated on the label (low accessibility). The only label type that recorded a significantly different WTP was the high-accessibility concentrated saving label, where the fuel saving was in one category of journey and the precise difference in fuel costs was calculated. The difference is significant at the five-percent level in Column 1 ($\beta = -29.66$, p = 0.015) and the one-percent level in Column 2 ($\beta = -32.85$, p = 0.007). This effect is consistent across a range of WTP thresholds (see Figure 8 in the Appendix). In unreported models, when labels are pooled by saving dispersion and accessibility, the concentrated savings manipulation is associated with significantly lower WTP, but the effect size is small. This result does not support the supposition of Kőszegi and Szeidl (2013) about how to make an option more attractive. That high accessibility is associated with a marginally significant decrease in WTP does not align with the recent findings of Dertwinkel-Kalt et al. (2022), but could be explained by left-digit bias making the low-accessibility savings seem more attractive (Busse et al., 2013).

5 Discussion and conclusion

This experiment tested the explanatory power of two novel cognitive mechanisms – concentration bias and underestimation bias – that plausibly contribute to the energy-efficiency gap. At an applied level, this is the first study to simultaneously explore the effect of payment schedules and the temporal framing of fuel savings on willingness-to-pay.

9 How WTP varied by label type for each of the six combinations of payment schedule and savings frame is shown in Figure 5 in the Appendix.

The strongest finding to emerge is that a monthly payment schedule increases willingness-to-pay relative to making annual payments or a once-off deposit payment. The magnitude of the difference in WTP between the three payment schedules supports the common prediction of the cognitive mechanisms, underestimation bias and concentration bias, that the WTP gap should be greatest between the monthly and annual payment schedules. In contrast, discounting predicts the gap should be greatest between the annual payment and deposit payment, because a larger proportion of the payment is shifted further to the future.[10] This finding suggests that people (i) underestimate how much they are paying in total when they pay per month (the underestimation bias explanation) and/or (ii) find it less psychologically painful to pay in multiple small instalments (the focus-weight utility explanation of concentration bias).

Framing fuel costs over a longer period of time also increases WTP for a fuel efficiency upgrade. When people only see the monthly fuel cost, they may underestimate the accumulated savings from choosing the more efficient car. However, the effect size of this manipulation was smaller than variation in payment schedule: less than one-third the size in the mixed linear models, and three-fifths the size in the logistic regressions using the binary sufficient-WTP variable.

There were some limitations to this experiment. First, we did not directly measure time preferences. While random assignment to conditions makes it highly unlikely that differences in discount rates were a factor in the results, controlling for this factor may have resulted in more precise estimates of treatment effects. Moreover, we cannot speak to correlation between time preference and biases arising from bounded rationality, though it seems reasonable that when undervaluation appears severe, the explanation may be a combination of the two channels. The absence of a time preference measure means policy implications, which we mention briefly below, ought to be considered tentative. Another limitation is the hypothetical nature of the willingness-to-pay elicitation. Hypothetical bias can amplify effect sizes (Murphy et al., 2005). This is a relatively minor concern in the current context because while people may appear more virtuous when it is free to do so, there is no reason to think hypothetical bias should be vary in strength *between* experimental conditions. Another hypothetical element which may have introduced noise to WTP responses was the "personalized fuel cost forecast". It was necessary to set savings exogenously, but participants who disagreed with the implied driving behaviour (i.e. the ratio of short, medium and long drives) may have adjusted their WTP based on how they actually thought they would drive. These caveats notwithstanding, hypothetical studies provide guidance on which interventions are most likely to succeed in changing behaviour.

In light of the aforementioned limitations and need for future research, we must proceed with caution when drawing implications for policy; however, some indicative directions are worth noting. Altering payment schedules would appear to have larger effects than changing the temporal framing of savings. The energy paradox literature has mostly focused on role of fuel cost framing, from the perspective that given the right information, people will make the welfare-maximizing choice. In contrast, the marketing literature on how payment structures influence demand (e.g. Gourville (1998)) has received much less attention. A tentative implication of these findings is that offering disaggregated payment schedules may offer more traction for increasing the uptake of energy-efficient technologies. Future research should test whether this finding generalizes to non-hypothetical field settings. Second, it may be fruitful to directly investigate the significant and consistent difference in WTP across age brackets, with small payments spread over time apparently of greater appeal to the youngest cohort (aged 18–39). Plausible reasons include younger consumers being more accustomed to recurring payments across a range of consumption domains, and second, that members of this cohort are more likely to be liquidity constrained.

More research is also needed to determine the relative strength of underestimation bias and concentration bias in this setting. As discussed throughout, these cognitive mechanisms make very similar predictions, and the test of relative predictive power in this experiment was inconclusive. This may have been due to the realism constraint, namely the need to use a leasing scenario that was familiar to people. However, future experiments with less of an applied focus may succeed in comparing relative predictive power. Such a design may utilize sequences of costs and benefits that are specifically designed to exploit the differences in functional forms of underestimation bias and concentration bias. This research question has applied relevance because the mechanisms have different implications for policy. Underestimation bias is a clear mistake which drives a wedge between true preferences and choices. Policy interventions that shoulder some of the arithmetic burden for consumers would be valid to correct this 'internality'. For instance, providing lifetime fuel costs instead of monthly fuel costs could be welfare-improving in correcting the bias that causes choices to diverge from true preferences (Beshears et al., 2008). However, there is an ethical tension in leveraging this bias to enhance overall WTP by providing a monthly payment option. This tension must be tempered against the knowledge that the private sector choice architects may operate by inducing the same mistake (e.g. through six-year finance plans). Concentration bias, on the other hand, is not a mistake at the time the decision is made. Kőszegi and Szeidl (2013) explain that choices are made based on focus-weighted consumption utility, but these focus-weights do not reflect welfare from experienced utility. A policy maker would need to navigate this nuance, for instance by assuring the vehicle seller and buyer that imposing a life time fuel cost label – which may well make *any* purchase seem less appealing – is justified by the buyer's future interests, as judged by experienced utility. This is similar to the argument to restrict the temptations available to present-focused individuals. But preoccupation with policy difficulties is premature before future research is conducted on the role of these mechanisms in valuation of future energy savings.

To conclude, the 'energy paradox' is one of the most pernicious puzzles in the energy economics literature, and its unresolved nature strongly hints at some pieces being overlooked. Behaviour that appears to be the expression of an individual time preference parameter may have a cognitive antecedent. Allcott (2016) suggested testing concentration bias, which is founded on the concept of

[10] Moreover, as shown in Section 3.2.1, depending on exactly when in the future the annual payment is made, pure discounting could actually predict the monthly — annual WTP gap to be negative.

limited attention (Kőszegi and Szeidl, 2013). We additionally tested the predictive power of underestimation bias, which is caused by a compressed internal number line (Scheibehenne, 2019). The experimental results show that in the current setting, both mechanisms can explain patterns of WTP not rationalizable under standard discounting alone. To make progress on solving the energy paradox, the set of behaviours broadly characterized as 'myopia' must be placed under the microscope. To that extent, these findings should help calibrate the lens.

Declaration of competing interest

None.

Data availability

Data and preregistration available at https://osf.io/73psm/.

Acknowledgments

The authors gratefully acknowledge PhD funding from the Irish Research Council, grant number GOIPG/2019/66. This research was also part funded by the Sustainable Energy Authority of Ireland (SEAI) Research, Development and Demonstration funding programme under grant agreement 19/RDD/525.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.reseneeco.2023.101404.

References

Allcott, H., 2016. Paternalism and energy efficiency: An overview. Annu. Rev. Econ. 8, 145–176.
Allcott, H., Knittel, C., 2019. Are consumers poorly informed about fuel economy? Evidence from two experiments. Am. Econ. J.: Econ. Policy 11 (1), 1–37.
Allcott, H., Wozny, N., 2014. Gasoline prices, fuel economy, and the energy paradox. Rev. Econ. Stat. 96 (5), 779–795.
Anderson, S.T., Sallee, J.M., 2011. Using loopholes to reveal the marginal cost of regulation: The case of fuel-economy standards. Amer. Econ. Rev. 101 (4), 1375–1409.
Andor, M.A., Gerster, A., Sommer, S., 2020. Consumer inattention, heuristic thinking and the role of energy labels. Energy J. 41 (1).
Attari, S.Z., DeKay, M.L., Davidson, C.I., De Bruin, W.B., 2010. Public perceptions of energy consumption and savings. Proc. Natl. Acad. Sci. 107 (37), 16054–16059.
Berkouwer, S.B., Dean, J.T., 2022. Credit, attention, and externalities in the adoption of energy efficient technologies by low-income households. Amer. Econ. Rev. 112 (10), 3291–3330.
Beshears, J., Choi, J.J., Laibson, D., Madrian, B.C., 2008. How are preferences revealed? J. Public Econ. 92 (8–9), 1787–1794.
Blasch, J., Filippini, M., Kumar, N., 2019. Boundedly rational consumers, energy and investment literacy, and the display of information on household appliances. Resour. Energy Econ. 56, 39–58.
Bradford, D., Courtemanche, C., Heutel, G., McAlvanah, P., Ruhm, C., 2017. Time preferences and consumer behavior. J. Risk Uncertain. 55 (2), 119–145.
Brezis, N., Bronfman, Z.Z., Usher, M., 2015. Adaptive spontaneous transitions between two mechanisms of numerical averaging. Sci. Rep. 5, 10415.
Busse, M.R., Lacetera, N., Pope, D.G., Silva-Risso, J., Sydnor, J.R., 2013. Estimating the effect of salience in wholesale and retail car markets. Amer. Econ. Rev. 103 (3), 575–579.
Carroll, J., Denny, E., Lyons, S., 2016. The effects of energy cost labelling on appliance purchasing decisions: Trial results from Ireland. J. Consumer Policy 39 (1), 23–40.
Dehaene, S., Izard, V., Spelke, E., Pica, P., 2008. Log or linear? Distinct intuitions of the number scale in Western and Amazonian indigene cultures. Science 320 (5880), 1217–1220.
DellaVigna, S., 2009. Psychology and economics: Evidence from the field. J. Econ. Literature 47 (2), 315–372.
Dertwinkel-Kalt, M., Gerhardt, H., Riener, G., Schwerter, F., Strang, L., 2022. Concentration bias in intertemporal choice. Rev. Econom. Stud. 89 (3), 1314–1334.
Gerarden, T.D., Newell, R.G., Stavins, R.N., 2017. Assessing the energy-efficiency gap. J. Econ. Lit. 55 (4), 1486–1525.
Gillingham, K.T., Houde, S., Van Benthem, A.A., 2021. Consumer myopia in vehicle purchases: evidence from a natural experiment. Am. Econ. J.: Econ. Policy 13 (3), 207–238.
Gillingham, K., Palmer, K., 2020. Bridging the energy efficiency gap: Policy insights from economic theory and empirical evidence. Rev. Environ. Econ. Policy.
Gillingham, K., Rapson, D., Wagner, G., 2016. The rebound effect and energy efficiency policy. Rev. Environ. Econ. Policy.
Gourville, J.T., 1998. Pennies-a-day: The effect of temporal reframing on transaction evaluation. J. Consum. Res. 24 (4), 395–408.
Greene, D.L., 2010. How consumers value fuel economy: A literature review.
Hausman, J.A., 1979. Individual discount rates and the purchase and utilization of energy-using durables. Bell J. Econ. 33–54.
Heinzle, S.L., 2012. Disclosure of energy operating cost information: A silver bullet for overcoming the energy-efficiency gap? J. Consumer Policy 35 (1), 43–64.
Jacobsen, M.R., Van Benthem, A.A., 2015. Vehicle scrappage and gasoline policy. Amer. Econ. Rev. 105 (3), 1312–1338.
Kahneman, D., 2003. Maps of bounded rationality: Psychology for behavioral economics. Am. Econ. Rev. 93 (5), 1449–1475.
Kallbekken, S., Sælen, H., Hermansen, E.A., 2013. Bridging the energy efficiency gap: A field experiment on lifetime energy costs and household appliances. J. Consumer Policy 36 (1), 1–16.
Kőszegi, B., Szeidl, A., 2013. A model of focusing in economic choice. Q. J. Econ. 128 (1), 53–104.
Laibson, D., 1997. Golden eggs and hyperbolic discounting. Q. J. Econ. 112 (2), 443–478.
Larrick, R.P., Soll, J.B., 2008. The MPG illusion. Science 320 (5883), 1593–1594.
McGowan, F., Denny, E., Lunn, P., 2022. Consistent underestimation in the intuitive summation of monetary amounts. Available at SSRN 4021708.
Min, J., Azevedo, I.L., Michalek, J., de Bruin, W.B., 2014. Labeling energy cost on light bulbs lowers implicit discount rates. Ecol. Econom. 97, 42–50.
Murphy, J.J., Allen, P.G., Stevens, T.H., Weatherhead, D., 2005. A meta-analysis of hypothetical bias in stated preference valuation. Environ. Resour. Econ. 30 (3), 313–325.

Newell, R.G., Siikamäki, J., 2014. Nudging energy efficiency behavior: The role of information labels. J. Assoc. Environ. Resour. Econ. 1 (4), 555–598.
Olschewski, S., Newell, B.R., Oberholzer, Y., Scheibehenne, B., 2021. Valuation and estimation from experience. J. Behav. Decis. Mak. 34 (5), 729–741.
Scheibehenne, B., 2019. The psychophysics of number integration: Evidence from the lab and from the field. Decision 6 (1), 61.
Silvi, M., Rosa, E.P., 2021. Reversing impatience: Framing mechanisms to increase the purchase of energy-saving appliances. Energy Econ. 103, 105563.
Somerville, J., 2022. Range-dependent attribute weighting in consumer choice: An experimental test. Econometrica 90 (2), 799–830.
Turrentine, T.S., Kurani, K.S., 2007. Car buyers and fuel economy? Energy Policy 35 (2), 1213–1223.
Wang, Y., Miao, Q., 2021. The impact of the corporate average fuel economy standards on technological changes in automobile fuel efficiency. Resour. Energy Econ. 63, 101211.