Journal of Behavioral and Experimental Finance 43 (2024) 100941

Contents lists available at ScienceDirect

Journal of Behavioral and Experimental Finance

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

Income volatility and saving decisions: Experimental evidence

Nathan Wang-Ly [a,*], Ben R. Newell [a,b]

[a] School of Psychology, UNSW Sydney, NSW 2052, Australia
[b] UNSW Institute of Climate Risk & Response, UNSW Sydney, 2052, Australia

A R T I C L E I N F O

JEL classification:
D14
G51

Keywords:
Income volatility
Precautionary saving
Uncertainty
Financial decision making

A B S T R A C T

Around the world, it is becoming increasingly common for individuals to have volatile incomes. Previous research offers mixed evidence on whether uncertainty about one's income may increase or decrease saving behaviour. Across four incentivised online experiments (N = 712), we examine the relationship between income volatility and saving behaviour in a novel financial decision making task. In this task, participants receive hypothetical income that is either consistent or that varies to different degrees. We capture participants' perceptions of how volatile their income is and observe how this influences their decision to spend the income or save it towards a hypothetical impending emergency. Our results indicate that receiving a more volatile income, as measured by its coefficient of variation (CV), leads to higher savings within our task. However, there appears to be a threshold level of volatility that must be exceeded before participants save differently relative to receiving a stable income.

1. Introduction

No one can be certain about how much they will earn in the future. However, not all uncertainties are equal; some individuals have greater confidence in their future incomes than others. An individual working in a traditional salaried role can reasonably expect to earn the same amount each pay cycle. In contrast, an individual who is employed casually or who freelances may find that their earnings fluctuate unpredictably. Around the world, economic trends suggest that this latter scenario is becoming increasingly common. In Australia, permanent full-time employment rates have declined from previous decades, insecure part-time employment has become prevalent, and there has been a dramatic rise in platform work (often referred to as 'gig work') (Senate Select Committee on Job Security, 2022). Similar trends have been observed across other advanced economies in Asia and Europe (Hipp et al., 2015; Kalleberg & Hewison, 2013).

Previous research has identified numerous ways in which income uncertainty may negatively impact individuals' and their families' lives. Experiencing fluctuations in income has been associated with psychological depression (Prause et al., 2009), higher risk of mortality (Halliday, 2007), greater risk of divorce (Nunley & Seals, 2010), and increased probability of mortgage default (Diaz-Serrano, 2005). Much of this research has been based on 'inter-year' income volatility (income that varies between years), with research on 'intra-year' income volatility (income that varies within a year) being comparatively scarce.

Only recently has there been an increase in studies examining how financial and health outcomes are influenced by monthly fluctuations in income (e.g., Anvari-Clark & Ansong, 2022; Gladstone et al., 2020), many of which have been motivated by the growth of the gig economy (e.g., Sayre, 2022; Wang et al., 2022), as well as the severe disruption to income caused by the COVID-19 pandemic (Wang-Ly & Newell, 2022). However, given the challenges of manipulating income in real-world settings, these studies have typically been observational in nature. These observational studies are inherently limited in their ability to prove causal links between the experience of income volatility and outcomes of interest (Rosenbaum, 2015). The current study aims to fill this gap in the literature through an experimental investigation of the impact of income volatility on financial behaviour.

To our knowledge, only two studies have investigated the influence of intra-year income volatility via experimental methods, with both finding evidence suggesting that volatility drives individuals to behave in a more financially impatient manner. The first study involved providing income spikes to Kenyan participants via an unconditional cash transfer program (West et al., 2020). Financial impatience was measured by giving participants multiple hypothetical choices between a 'smaller-sooner' reward (e.g., KSH 500 today) and a 'larger-later' reward (e.g., KSH 525 in six months). Participants who received the income spikes tended to choose the smaller-sooner option more often

* Corresponding author.
E-mail address: nathanwangly@gmail.com (N. Wang-Ly).

https://doi.org/10.1016/j.jbef.2024.100941
Received 5 August 2023; Received in revised form 16 April 2024; Accepted 20 May 2024
Available online 27 May 2024

compared to those who did not receive the spikes, demonstrating a greater degree of financial impatience. In the second study, participants completed simulated work tasks for hypothetical income as part of a lab-based experiment (Peetz et al., 2021). Participants either received a stable income (the same income per work task) or a volatile income (income that varied between work tasks). At the end of the experiment, participants who received the volatile income were more likely to choose the smaller-sooner participation reward (USD 15 today vs USD 17 in two weeks) compared to those who had received the stable income.

Inspired by these two studies, we saw an opportunity to apply a similar experimental approach to better understand how income volatility influences key household financial decisions. In the present research, we chose to focus our attention on saving behaviour for two reasons. First, income uncertainty tends to be more severe for individuals in poorer financial circumstances (Bania & Leete, 2022). For these individuals, the ability to accumulate savings and create 'financial slack' is especially critical for meeting unexpected expenses and avoiding debt (Mullainathan & Shafir, 2009). Second, the literature offers conflicting perspectives on whether income volatility should increase or decrease saving behaviour, meaning our work could offer useful clarifying evidence.

On the one hand, the argument that experiencing income volatility increases saving behaviour seems intuitive. To illustrate, imagine that you wanted to accumulate a certain amount of savings for future needs (e.g., a holiday or to buffer against future emergencies). If you earned a steady and predictable income, a prudent strategy might involve setting aside a consistent amount of money such that you would achieve the goal within your desired timeframe. Now suppose that your income was instead volatile and unpredictable. The consistent saving strategy would no longer be feasible; instead, it might be necessary to save more to insure against a potential dip in future income. This is the basic premise underlying the traditional economic perspective of how income volatility should affect saving behaviour. Early theoretical work suggested that being uncertain about one's future earnings should boost precautionary saving motives (Drèze & Modigliani, 1972; Leland, 1968; Sandmo, 1970). This view has been supported by numerous empirical studies (for an overview, see Lugilde et al., 2019).[1] For example, one study involving US personal income data found that doubling the level of income uncertainty of an individual or household was associated with a 29 percent increase in their ratio of wealth to income (i.e., a greater saving rate) (Kazarosian, 1997).

On the other hand, there are several valid reasons why income volatility might instead be expected to reduce saving behaviour. First, if income volatility results in greater financial impatience (Peetz et al., 2021; West et al., 2020), individuals may be biased towards actions that offer immediate benefits (spending) over actions that offer benefits in the distant future (saving). This would be consistent with previous work showing that more present-focused time preferences are associated with lower retirement saving balances (Goda et al., 2019). Second, experiencing income volatility may induce feelings of stress or scarcity, as dips in income can create temporary shortfalls between available funds and upcoming expenses, forcing individuals to make difficult trade-offs between their immediate and future needs (Bania & Leete, 2022). Both stress and scarcity have been found to impair long-term and goal-directed decision making (Haushofer & Fehr, 2014; Mani et al., 2013, 2020; Schwabe & Wolf, 2009), providing additional reason to believe that income volatility might reduce one's propensity to save. Third, individuals may decide to save less as part of a considered choice to abandon (or at least heavily discount) their focus on the long-term. They may instead shift their focus to achieving short-term outcomes, where the associated rewards are typically better defined and allow for earlier gratification—an approach that has been argued to be a rational response to uncertainty (Fawcett et al., 2012; Pepper & Nettle, 2017). Consistent with this view, numerous studies have shown that saving rates are negatively impacted when individuals feel a lack of control over their future circumstances (Cobb-Clark et al., 2016; Perry & Morris, 2005; Shapiro & Wu, 2011).

2. Study overview

2.1. Research design

To assess the merit of these opposing perspectives, we conducted four online experiments which involved a novel (hypothetical) financial decision making task. Within this task, we manipulated the volatility of income participants earned for completing work tasks (similar to Peetz et al., 2021), and observed how this influenced their decisions to either spend their earnings or save for an impending emergency.

In Experiment 1, we tested three levels of income volatility (including no volatility) and observed the highest saving rates from participants whose income was most volatile. However, we also observed a discrepancy between participants' self-reported perceptions of how volatile their income was and their observed saving behaviour. Experiment 2 examined whether this was because participants' saving behaviour was instead being driven by other dimensions of the income sequence (e.g., the minimum income received). Experiment 3 provided evidence that the discrepancy was due to participants being asked to provide an absolute judgment of volatility; participants were able to distinguish between levels of volatility when instead making relative judgments. Finally, in Experiment 4, we tested a single condition in which we made participants' income even more volatile than the most volatile condition in Experiment 1. This established a clearer relationship between income volatility and increased saving behaviour while also suggesting the existence of a volatility threshold that needs to be exceeded for consumers to change their saving behaviour.

Throughout the study, we report on analyses that involve comparisons of participants' behaviour both within and between experiments. We recognise that these latter comparisons forego random assignment, which limits our ability to cleanly establish causality and introduces the possibility for alternative explanations outside of our experimental manipulations (Harris et al., 2006). However, it should be noted that our experiments were conducted using the same financial decision task and offered the same incentives. Additionally, as we demonstrate in analyses reported in Appendix Table A.3, our participants did not differ meaningfully across any of the demographic or financial characteristics we recorded (e.g., age, education level)—which was to be expected given we had drawn from the same participant pool (Prolific Academic) for all four of our experiments. Thus, while we acknowledge that purely experimental results do hold stronger weight, we also felt confident that even with our quasi-experimental analysis approach, we could attribute any effects observed to our manipulation of income volatility.

2.2. Paper's contributions

Our study seeks to make several contributions to the literature. First, we continue to expand the body of experimental work examining the effect of income volatility on financial decisions. We offer (quasi-experimental) evidence in support of the view that experiencing volatile income increases saving behaviour. Second, we examine the impact of different degrees of income volatility, allowing us to investigate not only the extensive margin of the income volatility effect, but also the intensive margin. Third, we offer a novel financial decision making task that has potential for reuse in future income volatility experiments or could be adapted to answer similar questions around influences on saving behaviour.

It is equally worth noting what our study does not seek to do. Our intention with the present research is not to form normative judgments

[1] However, see Fulford (2015) for evidence that income uncertainty is not an important motive for precautionary saving.

about how individuals should adjust their saving behaviour in response to income volatility. Our reasoning for this is two-fold. First, as highlighted in the Introduction, there are valid explanations that could justify either increases or decreases in saving. We therefore believe that focusing on providing a descriptive account, which documents how individuals actually respond when experiencing a volatile income, is the more pertinent knowledge gap to address in the current literature. Second, we believe there to be an inherent trade-off when designing a study between reasonably reflecting the decisions that individuals face in the real world and allowing for normative judgments to be made. Real-world financial decisions are plagued by unquantifiable risks and probabilities (or 'Knightian uncertainty'; Knight, 1921) that make it challenging (if not impossible) to determine what behaviour should be considered optimal. We return to this latter point in the General Discussion when we consider the many opportunities to expand upon our study.

3. Experiment 1

The purpose of Experiment 1 was to examine how participants' saving behaviour within our financial decision making task was influenced by the level of income volatility they experienced. We generated different sequences of income that varied in their coefficients of variation (CV). CV is a commonly referenced measure of income volatility (Bania & Leete, 2022; Farrell & Greig, 2016; Hannagan & Morduch, 2015) that can be calculated by dividing the standard deviation of a sequence by its mean. Therefore, the larger the CV value, the more volatile the income. One advantage of the CV metric is that it provides a standardised value that allows for comparison across incomes that may differ in magnitude, frequency, and timing.[2] For Experiment 1, we tested three income sequences with CV values of 0 (i.e., no volatility), 0.30, and 0.60.

3.1. Method

3.1.1. Participants

We recruited 241 participants[3] from Prolific Academic, an online participant recruitment platform. Our sample size was based on an a priori power analysis conducted using statistical software G*Power (Faul et al., 2007) based on detecting a small effect in a between-subjects study (δ =.80; γ =.20, α =.05). To be eligible for our sample, participants needed to be between the ages of 18 and 65, located in the UK, and fluent in English. Participants received £3.00 (USD 3.77) in exchange for their participation and were also eligible for an additional bonus of £3.00 (USD 3.77) based on their performance in the financial decision making task.

Participants' ages ranged between 19 and 65 years (M_{age} = 39.65, SD = 11.79), 59.8% of which were male, 39.4% female, and 0.8% who reported a different gender identity. For most participants, the highest education level attained was either through high school (36.1%) or an undergraduate degree (45.2%). Most participants were employed either full-time (61.0%) or part-time (16.2%), with the median personal income band being £20,000-£29,999. See Appendix Tables A.1 and A.2 for additional information about our samples and a comparison against nationally representative data.

3.1.2. Materials

The financial decision making task[4] used in Experiment 1 involved completing work tasks for hypothetical income and making hypothetical decisions around spending and saving for an impending emergency. In each round of the game, participants were asked to rearrange a six-character string of letters (e.g., "hksykjl") into alphabetical order as their work task. Upon completion, they received income and were asked how much they wanted to spend.[5] Money could be spent to purchase points (£1.00 = 1 point), which were later used to determine eligibility for the bonus performance-based payment.[6] Any money that was not spent remained as savings and was carried over into subsequent rounds, where it could again be spent or saved.

The income participants received for their work tasks varied depending on which of the three experimental conditions they were allocated to: No Volatility (n = 81), Low Volatility (n = 80), and High Volatility (n = 80). In all three conditions, participants received the same total income across the 15 task rounds: £15,000. However, the extent to which their income fluctuated differed. In the No Volatility condition, participants received £1,000 after each work task—corresponding to a CV of 0. In the Low Volatility condition, participants' income ranged between £431 and £1,340, with the overall sequence of incomes having a CV of 0.30.[7] In the High Volatility condition, incomes ranged between £173 and £2,088, with a CV of 0.60. For participants in the Low and High Volatility conditions, the order in which they received the different incomes was randomised. Details on the income sequences used across the experiments are shown in Table 1.

At the end of the task (after the 15th round), participants encountered a financial emergency: home repairs costing £4,500 that had resulted from storm damages. When this emergency occurred, if participants had adequate savings, they "won" the task, meaning they could keep the points they had purchased. However, if participants' savings were inadequate, they lost all their purchased points. Importantly, while participants were advised at the start of the task that there would be a financial emergency, they were not told when it would occur nor how much it would cost. Therefore, their objective was to balance the competing goals of maximising points purchased (to potentially earn the bonus payment) while ensuring they saved enough to withstand the financial emergency (to avoid losing their points).

Just prior to revealing the cost of the financial emergency, participants were asked to provide a series of responses. This timing was deliberate such that participants' responses would not be affected by the outcome of the experiment (i.e., whether they had saved adequately). Participants were asked to rate how volatile they perceived their income to have been during the task (0 = "Not volatile"; 10 = "Extremely volatile"). Participants were advised that they could think of volatility as a reflection of how well they could predict their income if there was an additional round[8] (Konovalova & Pachur, 2021). We similarly asked participants to rate how difficult it was to decide how much to spend or save each round (0 = "Not difficult"; 10 = "Extremely difficult").

[2] However, see Kvålseth (2017) for an alternative measure of volatility (the 'second-order coefficient of variation') that has additional useful properties such as being restricted between zero and one, and being less sensitive to outliers.

[3] In both Experiments 1 and 2, we encountered an issue during data collection which led to having one participant more than was intended.

[4] Select screenshots of the task and the experimental code are available in the Supplementary Materials (https://osf.io/pqgha/).

[5] Participants' speed of completion had no influence on their income earned. However, completion of the work task was required to proceed to the next round.

[6] Participants were eligible for a bonus payment (in addition to their participation payment) if their final score in the task was within the top 10 percent of their experimental condition.

[7] Previous work has estimated that the average household experiences a CV of between 0.30 and 0.40 (Farrell et al., 2019; Hannagan & Morduch, 2015); however, this varies substantially between lower- and higher-income households (Mills & Amick, 2010; Morris et al., 2015).

[8] Specifically, we gave participants the following explanation of volatility: *"You could think of volatility as reflecting how well you could predict the income you would receive if there was another round. The higher the volatility, the less likely you would be to predict this correctly."*

Table 1
Dimensions of income sequences used in Experiments 1–4.

Condition	Incomes (£)	CV	Min	Max	Range
No Volatility	1,000 1,000 1,000 1,000 1,000 1,000 1,000 1,000 1,000 1,000 1,000 1,000 1,000 1,000 1,000	0.00	1,000	1,000	0
Low Volatility	431 570 656 794 873 883 910 990 1,158 1,202 1,277 1,279 1,311 1,326 1,340	0.30	431	1,340	909
High Volatility	173 274 466 526 529 650 661 847 1,101 1,374 1,404 1,436 1,571 1,900 2,088	0.60	173	2,088	1,915
Same Range	173 881 911 929 946 952 956 981 989 996 1,022 1,041 1,063 1,072 2,088	0.37	173	2,088	1,915
Same Min	173 552 654 851 863 864 889 1,017 1,036 1,211 1,240 1,257 1,310 1,503 1,580	0.37	173	1,580	1,407
Same Max	681 689 746 767 772 782 785 816 983 994 1,029 1,242 1,275 1,351 2,088	0.37	681	2,088	1,407
Very High Volatility	44 152 184 299 372 536 763 767 788 1,101 1,300 1,618 1,654 1,965 3,457	0.90	44	3,457	3,413

Note. CV = Coefficient of variation. Min = Minimum income. Max = Maximum income.

Finally, we measured financial impatience using the 'fill-in-the-blank' approach (Smith & Hantula, 2008), which has been previously used in studies involving delay discounting (Chapman, 1996; Kirby & Maraković, 1995; Weatherly et al., 2010). Participants were instructed to imagine that they were given the choice between two bonus rewards for their involvement in the experiment. The first option was to receive £50 in two weeks' time. Alternatively, they could opt for an immediate reward. Participants were asked to indicate the lowest amount they would accept immediately ('smaller-sooner amount') that would make it preferable to the larger, delayed reward. The response was bounded between £0 and £50, with a lower smaller-sooner amount indicating a higher degree of financial impatience.

3.1.3. Procedure

Participants were advised that they would be playing a financial decision making game and received instructions outlining the work tasks and the objective of the game. This was followed by a practice stage which included three practice rounds (with income of £50 per work task) and a practice financial emergency (a £50 parking ticket fine).

After the practice stage, participants' savings and points were reset to zero and they began the 15 experiment rounds. Once the experiment rounds were complete, participants were informed that the financial emergency was about to occur and were asked to provide responses to the perceived volatility, perceived difficulty, and financial impatience measures described in the previous section. Following this, the cost of the emergency was revealed, and participants were advised of the outcome of the game. Finally, participants were asked to provide information about their demographics and financial situations before receiving a debrief of the experiment.

3.2. Results

For all four experiments, we conducted Bayesian statistical analyses (using default priors) and report the Bayes factors (BF_{10}) associated with our results. In cases where the Bayes factor was less than one (i.e., evidence favoured the null hypothesis), we report the inverse Bayes factor (BF_{01}) to assist with comparability of evidence strength. Interpretations of Bayes factors were based on guidelines suggested in Lee and Wagenmakers (2013). Unless indicated otherwise, analyses of variance (ANOVAs) were used for overall tests of differences across multiple conditions and t-tests were used for pairwise comparisons between conditions. Our experiment data and code have been made available in the Supplementary Materials (https://osf.io/pqgha/).

3.2.1. Perceived volatility ratings

We first examined whether there were differences between conditions in participants' ratings of how volatile their income was ('perceived volatility ratings'). As seen in Fig. 1, our analysis indicated that there was extreme evidence that perceived volatility ratings differed between conditions ($BF_{10} = 6.72 \times 10^{13}$), suggesting that our income volatility manipulation had worked. This difference was driven by substantially lower volatility ratings in the No Volatility condition ($M = 3.17$) compared to the Low Volatility ($M = 6.00$) and High Volatility ($M = 6.40$) conditions ($BF_{10} = 2.04 \times 10^8$ and 1.11×10^{11} respectively). However, there was moderate evidence suggesting that perceived volatility ratings did not differ between the Low and High Volatility conditions ($BF_{01} = 3.15$).

3.2.2. Final savings

We next analysed how much participants had saved at the end of the task ('final savings'). As seen in Fig. 2, we observed very strong evidence that final savings differed between conditions ($BF_{10} = 38.71$). Participants saved substantially more in the High Volatility condition ($M =$ £8,571.64) compared to the No Volatility ($M =$ £6,647.88) and Low Volatility ($M =$ £6,579.54) conditions ($BF_{10} = 33.67$ and 24.23 respectively). However, there was moderate evidence suggesting that final savings did not differ between the No and Low Volatility conditions ($BF_{01} = 5.84$).

3.2.3. Perceived difficulty and financial impatience

We then examined participants' ratings of how difficult they found it to decide how much to spend or save throughout the task ('perceived difficulty'). There was moderate evidence that these ratings did not differ between conditions ($BF_{01} = 6.44$). The mean rating provided across the conditions was 4.07 (out of 10). Likewise, there was moderate evidence that financial impatience did not vary between conditions, as measured via participants' smaller-sooner amounts ($BF_{01} = 7.67$). The mean amount participants were willing to accept immediately instead of a £50 reward in two weeks' time was £37.60.

3.3. Discussion

Our results indicated that participants saved substantially more in the High Volatility condition than the No and Low Volatility conditions—consistent with the view that income volatility increases saving motives. However, we did not observe the same boost in savings for the Low Volatility condition relative to the No Volatility condition. The simplest explanation would be that participants' income was volatile enough to impact their saving behaviour in the High Volatility condition but was not volatile enough to do so in the Low Volatility condition. This would imply that there exists a threshold between the CVs we used (0.30 for Low Volatility; 0.60 for High Volatility) which must be surpassed before a change in saving behaviour would be observed. However, complicating this explanation was the fact that participants in the Low and High Volatility conditions reported their income as being similarly volatile (based on their perceived volatility ratings). The question we thus needed to address was why participants in the Low Volatility condition did not save differently relative to the No Volatility condition despite perceiving their income to be as volatile as the High Volatile

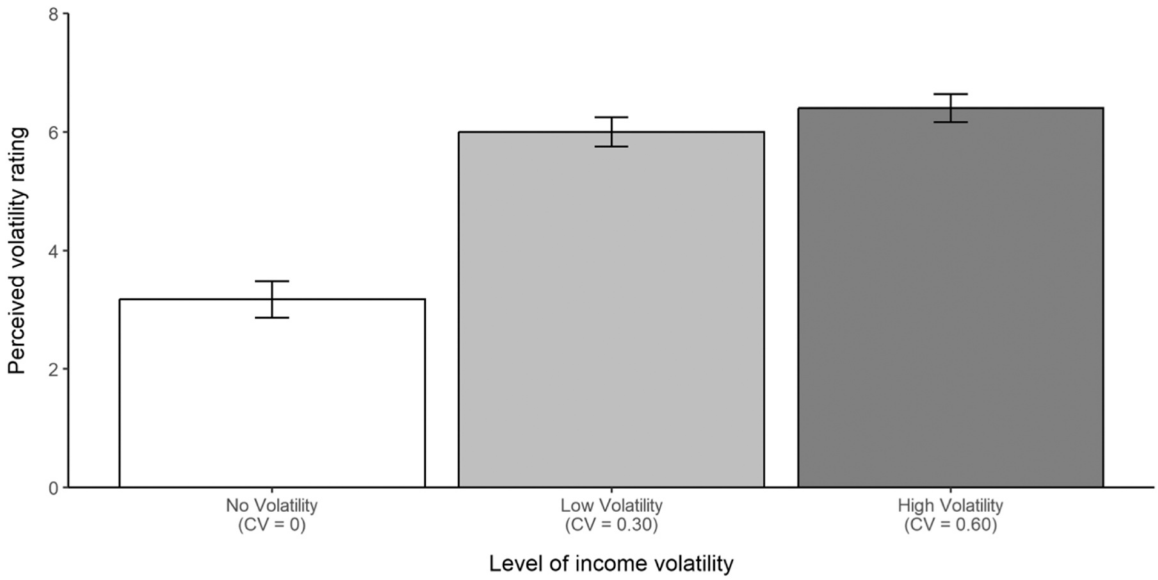

Fig. 1. Mean perceived volatility ratings by level of income volatility in Experiment 1. Standard error bars indicated.

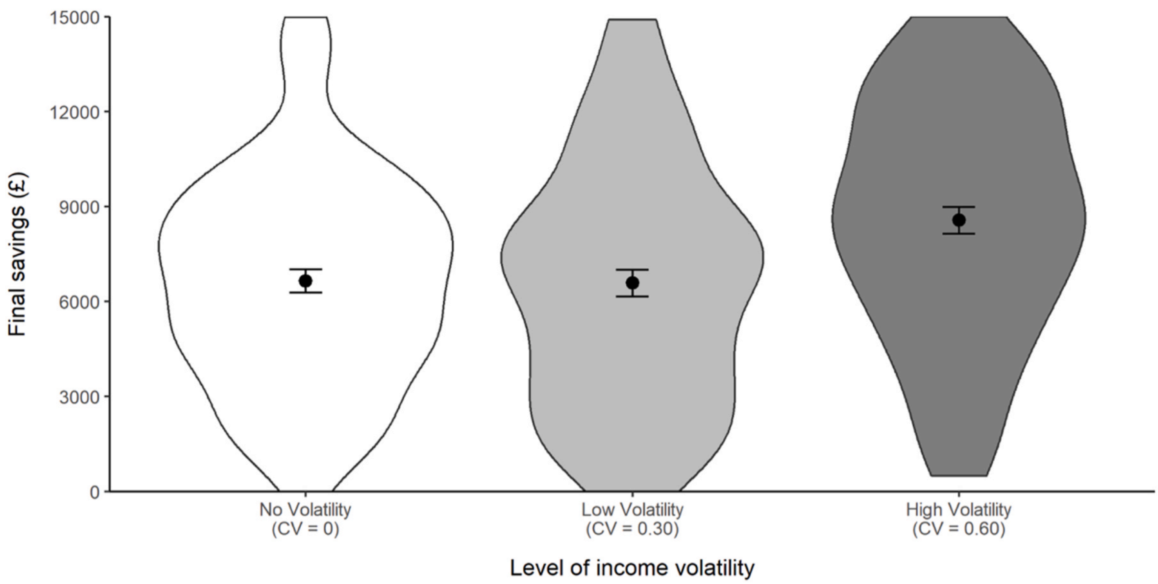

Fig. 2. Distribution of final savings by level of income volatility in Experiment 1. Means and standard error bars indicated.

condition.

One possibility we considered was whether the increased saving behaviour observed from participants in the High Volatility condition was being driven by something other than income volatility (as defined by CV). The idea was inspired by recent work which showed that individuals' perceptions of a number sequence's variance is not always perfectly related to the objective statistical variance of the sequence (Konovalova & Pachur, 2021). In their study, Konovalova and Pachur (2021) examined how other characteristics of a number sequence (e.g., mean, range, pairwise distance) contribute to perceptions of variance. The study found that a sequence's range often served as a better individual predictor of participants' perceptions of variance than statistical variance itself.

This led us to examine in Experiment 2 whether three other dimensions of the income sequences we gave participants could better account for their saving behaviour: the minimum income participants received during the task, the maximum income, and the range of incomes. As shown in Table 1, the income sequences used in Experiment 1 varied not only in their CV, but also these additional dimensions. As an example, it may have been the minimum income participants received that was driving their saving decisions—perhaps because the income shock associated with earning an unexpectedly low income forced participants to update their view of how uncertain the (experimental) world was, thereby prompting them to save more as a precaution (Chamon et al., 2013; Collins & Gjertson, 2013). If this were the case, it would explain why participants' perceived volatility ratings did not neatly map onto their saving behaviour.

While the main focus for this study is saving behaviour, it is worth noting that our income volatility manipulation did not appear to increase participants' financial impatience—contrary to prior findings (Peetz et al., 2021; West et al., 2020). One possible reason for our failure to observe an effect is that both our income volatility manipulation and our measure of financial impatience were hypothetical. While Peetz et al. (2021) also manipulated the volatility of hypothetical income, their measure of financial impatience involved participants choosing between two real payments. In the case of West et al., (2020), the measure of financial impatience involved hypothetical choices, but participants experienced volatility in their real income. Prior work has shown that the choices people make in hypothetical and real-world situations can differ (e.g., FeldmanHall et al., 2012; Green & Lawyer, 2014; Vlaev, 2012); for example, people tend to be more patient and altruistic when responding to hypothetical scenarios (Camerer & Mobbs, 2017).

A second possibility for the inconsistency in our findings could be due to our use of different methods of measuring financial impatience. While our study used a fill-in-the-blank approach, both prior studies used 'binary-choice' tasks, which give participants the choice between defined smaller-sooner and larger-later amounts. Previous work suggests that binary-choice tasks can produce steeper discounting rates compared to fill-in-the-blank task (Smith & Hantula, 2008). It is therefore possible that the effect of income volatility on financial impatience is small—one that our measure was unable to detect, but that more sensitive measures may have been able to.

A third possibility is that our findings did not corroborate previous work due to the differences in the participant samples we recruited. As shown in Appendix Tables A.1 and A.2, our sample of participants were broadly representative of the UK population. In contrast, the sample described in Peetz et al. (2021) consisted disproportionately of lower-income earners from the US. Prior research suggests there is a link between financial scarcity and cognitive functioning (e.g., Mani et al., 2013, 2020), which may explain the inconsistencies in our findings. Meanwhile, West et al. (2020)'s conclusions were based on a sample of Kenyan women—and there has similarly been substantial research documenting differences between Western and non-Western populations' decision processes (Henrich et al., 2010). Though exploring these potential explanations (or any others) was beyond the scope of our study, our null result suggests the need for further investigation to have greater confidence in the link between income volatility and financial impatience.

4. Experiment 2

For Experiment 2, we generated three new income sequences that would allow us to tease apart whether the minimum income, maximum income, or range of incomes would provide a better account for saving behaviour than CV. The new sequences were designed such that they could be compared against one another, as well as against the High Volatility condition from Experiment 1. Each sequence shared a different set of dimensions with the High Volatility condition: either the minimum income (but not the maximum or range), the maximum income (but not the minimum or range), or both the minimum and maximum incomes (and therefore also the range). Notably, however, all three conditions shared the same statistical variance as one another (CVs = 0.37), but were deliberately created to be less volatile than the High Volatility condition (CV = 0.60).

4.1. Method

4.1.1. Participants

We recruited 241 new participants from Prolific Academic using the same eligibility criteria described for Experiment 1. In addition to these criteria, we excluded anyone who had participants in the previous experiment. Participants against received a £3.00 (USD 3.77) participation payment and were eligible for a £3.00 (USD 3.77) bonus payment based on their task performance.

Participants' ages ranged between 18 and 65 years (M_{age} = 39.39, SD = 12.14), 44.8% of which were male, 53.9% female, and 0.8% who reported a different gender identity. For most participants, the highest education level attained was either through high school (41.1%) or an undergraduate degree (37.8%). Most participants were employed either full-time (52.7%) or part-time (20.3%), with the median personal income band being £20,000-£29,999.

4.1.2. Materials and procedure

Experiment 2 used the same financial decision making task and procedure as described for Experiment 1. However, participants were assigned to three new conditions which used different income sequences from the prior experiment: "Same Range" (n = 80), "Same Min" (n = 81), and "Same Max" (n = 80). The names of these conditions indicated what characteristic(s) their income sequence shared with the High Volatility condition from Experiment 1 (see Table 1).

In the Same Range condition, participants' income sequence shared the same minimum and maximum values as the High Volatility condition (£173 and £2,088 respectively), and therefore the same range (£1,915). However, the income sequence was constructed to have a lower coefficient of variation (CV = 0.37 vs 0.60 for the High Volatility condition). This would allow us to directly compare the two conditions to determine whether CV was driving the difference in saving behaviour or whether it was one of the other dimensions.

In the Same Min condition, participants' income ranged from £173 to £1,580 across rounds—sharing the same minimum income as the High Volatility condition but differing in the maximum income and range. Conversely, the Same Max condition shared the same maximum income as the High Volatility condition (£2,088) but had a different minimum income (£681) and range. Notably, however, the range of the Same Min and Same Max conditions were designed to be identical (£1,407). Both sequences were also constructed to share the same CV as the Same Range condition (CV = 0.37).

All three sequences therefore shared the same statistical variance (unlike in Experiment 1), but varied in their minimum income, maximum income, and range of incomes. As with the previous experiment, regardless of which income sequence participants received, their total income across the task was £15,000. By keeping this and all other task parameters the same as Experiment 1, it became possible to directly compare the High Volatility condition from the first experiment with the new conditions in Experiment 2.

4.2. Results

4.2.1. Comparisons between Same Range, Same Min, and Same Max conditions

We first analysed the data from our three Experiment 2 conditions. This allowed us to explore whether any differences in our measures were being driven by dimensions of participants' income sequences other than CV.

Across our four measures (perceived volatility, final savings, perceived difficulty, and financial impatience), we consistently observed no differences between our three conditions. The evidence for this was moderate for perceived volatility (BF_{01} = 3.91), strong for final savings (BF_{01} = 16.19), moderate for perceived difficulty (BF_{01} = 6.82), and strong for financial impatience (BF_{01} = 12.50).

4.2.2. Comparisons against High Volatility condition

We then compared the perceived volatility ratings and final savings of our three Experiment 2 conditions against the High Volatility condition from Experiment 1 (see Fig. 3).

There was anecdotal evidence that participants perceived their income to be more volatile in the High Volatility condition (M = 6.40) than the Same Range condition (M = 5.69) (BF_{10} = 1.52). However, there was moderate evidence that the perceived volatility ratings did not differ between the High Volatility condition and the Same Min (M = 6.31) or Same Max (M = 6.11) conditions (BF_{01} = 5.66 and 4.18 respectively).

In contrast, there was consistent evidence that participants saved more in the High Volatility condition (M = £8.571.64) compared to the other three conditions. The evidence was moderate for the Same Range (M = £7,096.65) and Same Max (M = £6,842.40) conditions (BF_{10} = 3.03 and 6.69 respectively), and anecdotal for the Same Min condition (M = £7,362.69) (BF_{10} = 1.16).

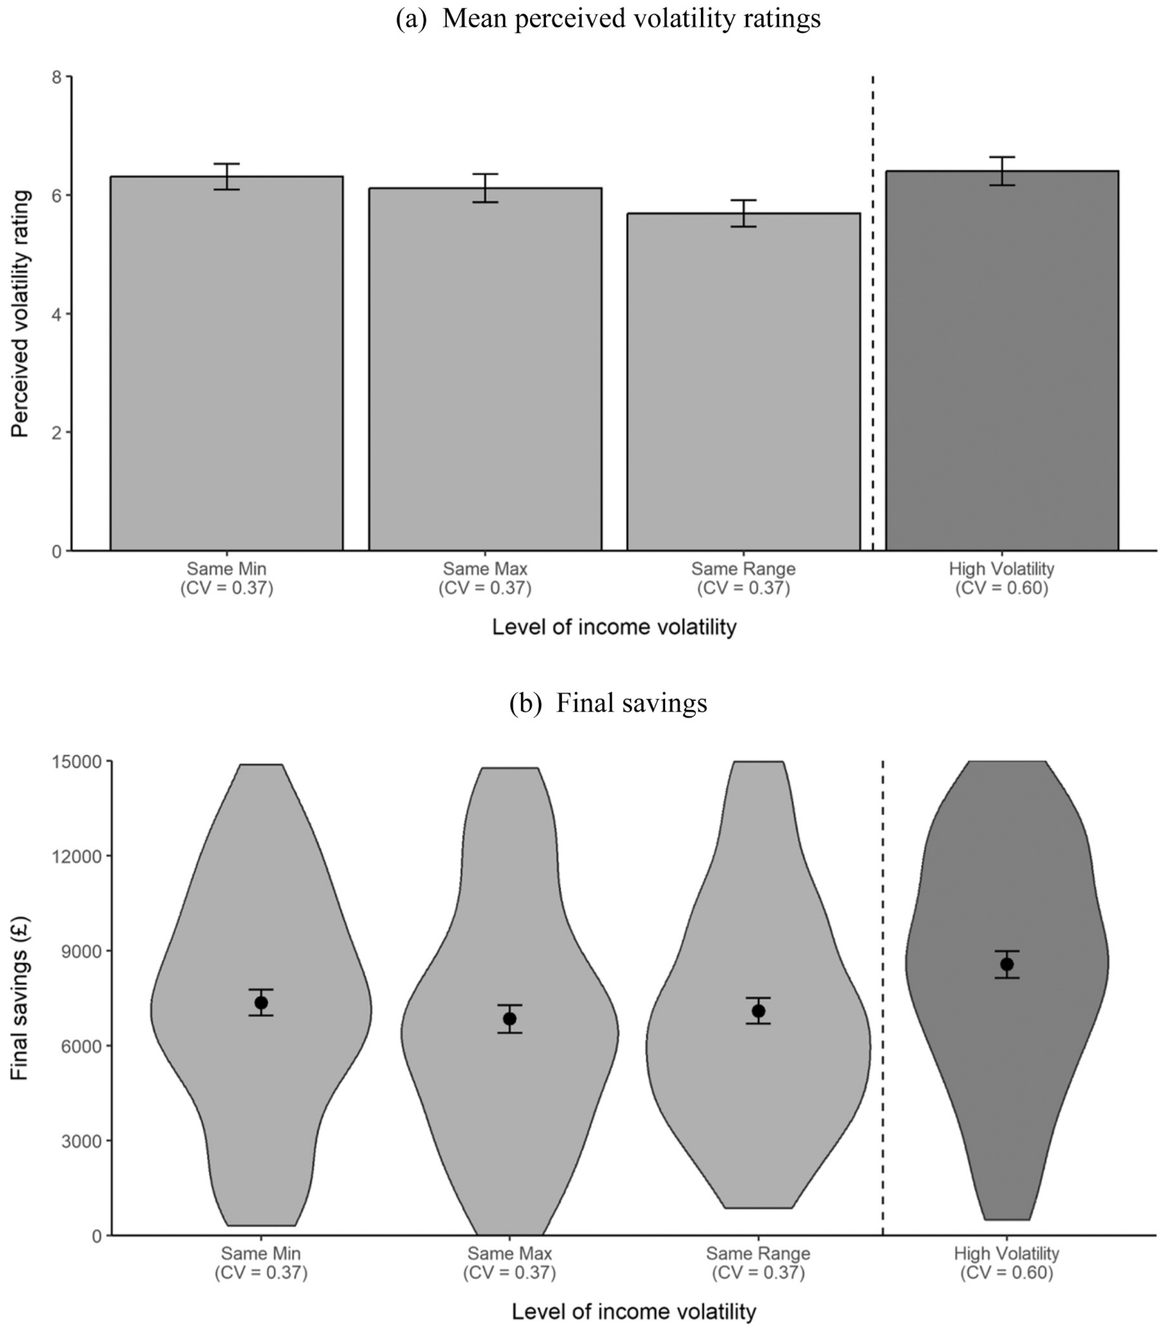

Fig. 3. Perceived volatility ratings and final savings by level of income volatility in Experiment 2. High Volatility condition from Experiment 1 shown on right of dotted line for comparison. Means and standard error bars indicated.

4.3. Discussion

Our results did not support the hypothesis that saving behaviour was driven by a dimension of the income sequence other than CV. Despite varying the minimum income, maximum income, and range of incomes between conditions, participants all exhibited similar saving behaviour. If anything, our findings provided greater reason to believe that CV was the driver of saving behaviour. All three conditions saved less than the High Volatility condition, with the consistent difference being their lower CV (CV = 0.37 vs 0.60). Thus, we saw further evidence that increased income volatility corresponded to higher savings within our task.

However, we continued to observe a discrepancy between participants' perceptions of their income's volatility and their saving behaviour. While participants saved more in the High Volatility condition, their perceived volatility ratings were similar to the ratings provided in our Experiment 2 conditions (with the exception of the Same Range condition, where the evidence of a difference was only anecdotal). This echoed what we had observed in Experiment 1, in which participants from the Low and High Volatility conditions had rated their income as similarly volatile despite exhibiting different saving behaviour.

This led us to consider a new hypothesis: that participants were giving similar responses on our perceived volatility rating scale because we were asking them to make an absolute judgment with only one objective point of reference (0 = "Not volatile at all"). Previous work suggests that judgments can differ when made from an absolute versus relative standpoint (e.g., Fox & Tversky, 1995; Goffin & Olson, 2011;

Stewart et al., 2006). It was therefore possible that participants in different conditions were perceiving the volatility of their income differently, but shared similar thought processes in how this should be reflected in their response on the scale—for example, thinking "there is some volatility but not an extreme amount, so I will respond with a seven". This could explain why participants who experienced some degree of income volatility (i.e., all conditions except for the No Volatility condition) provided similar ratings of perceived volatility while differing in their saving behaviour.

To examine this hypothesis, we needed participants to provide relative judgments of their income's volatility. To obtain these judgments, we adapted our procedure in Experiment 3 such that participants would have two attempts at the financial decision making task. These two attempts used income sequences with different levels of volatility, allowing us to compare participants' perceived volatility ratings between attempts.

5. Experiment 3

For Experiment 3, participants were given two attempts at the financial decision making task. In one attempt, participants received the same income as the Low Volatility condition from Experiment 1; in the other, they received the same income as the High Volatility condition. The ordering of the income sequences was counterbalanced.

Our goal was to examine whether participants would be sensitive to the different levels of volatility between their attempts. We preregistered[9] three key hypotheses based on this goal prior to conducting the experiment. First, participants would provide higher perceived volatility ratings during their High Volatility attempt compared to their Low Volatility attempt. Second, participants would provide similar ratings during their first attempt, irrespective of whether they received the Low Volatility or High Volatility income sequence. This would be consistent with what had been observed in Experiment 1. Third, perceived volatility ratings would differ when comparing second attempts at the task. We expected that since participants would be making relative judgments of volatility in their second attempts, they would provide higher ratings if this second attempt involved the High Volatility income sequence, whereas they would provide lower ratings if their second attempt instead involved the Low Volatility income sequence.

5.1. Method

5.1.1. Participants

We recruited 150 new participants from Prolific Academic using the same eligibility criteria as described in the previous two experiments. Sample size was based on a power analysis that allowed for detection of a small effect in both between- and within-subjects comparisons (δ =.80; γ =.20; α =.05). Participants received a £4.50 (USD 5.66) participation payment and were eligible or a £3.00 (USD 3.77) bonus payment based on their task performance.

Participants' ages ranged between 18 and 64 years ($M_{\text{age}} = 38.86$, SD = 11.67), 50.0% of which were male, 49.3% female, and 0.7% who reported a different gender identity. For most participants, the highest education level attained was either through high school (42.0%) or an undergraduate degree (40.7%). Most participants were employed either full-time (55.3%) or part-time (16.0%), with the median personal income band being £20,000-£29,999.

5.1.2. Materials and procedure

For Experiment 3, the procedure was adapted such that participants were given two attempts at the financial decision making task. In one attempt, participants received the same income sequence that was used in the Low Volatility condition from Experiment 1; in the other, they received the High Volatility income sequence. The order in which participants received these attempts was counterbalanced, meaning half of participants received the Low Volatility version first, while the other half received the High Volatility version first (n's = 75).

The inherent issue in allowing participants two attempts at the task is that their behaviour in the second attempt may be confounded by learnings from the first attempt. This would impair our ability to compare our results against the behaviour we had observed in Experiment 1. We sought to mitigate this concern through several design choices when adjusting our procedure from the previous experiments. First, at the end of participants' first attempt, we advised them that the emergency had occurred, but did not reveal how much it cost. This was only revealed at the end of the experiment after participants had completed both attempts. Second, we explicitly indicated to participants that the timing and cost of the emergency may differ between attempts. This sought to dissuade participants from assuming that the timing of the emergency would be the same for both attempts.[10] Finally, although participants' final score in the experiment was calculated by combining their performance in each attempt, we visually reset their score to zero in between attempts. This was intended to reiterate to participants that the two attempts were independent.

The overall procedure was as follows. Participants were advised that they would be given two attempts at a financial decision making game. After reading the instructions, participants underwent a practice stage, followed by their first attempt at the task. At the end of the first attempt, participants were informed that the financial emergency was about to occur and were asked to rate the volatility of their income. Participants were then told that the cost of the emergency would only be revealed after they had completed their second attempt, which they subsequently commenced. At the end of their second attempt, participants were again asked to rate the volatility of their income. Following this, the cost of both emergencies was revealed,[11] and participants were advised of the outcome of each attempt. Participants' total score across both attempts was used to determine eligibility for the bonus reward.[12] Finally, participants were asked to provide information about their demographic and financial situations, before receiving a debrief.

5.2. Results

5.2.1. Comparison of volatility ratings between attempts

Fig. 4 shows participants' mean perceived volatility ratings as a function of the level of volatility (Low or High) and whether it was their first or second attempt. As predicted, we observed moderate evidence that participants provided similar ratings during their first attempt, regardless of whether this was the Low Volatility or High Volatility version ($BF_{01} = 3.28$). However, contrary to our prediction, there was also anecdotal evidence that ratings did not differ on participants' second attempt either ($BF_{01} = 2.13$).

While we did not find evidence for our hypothesis when examining volatility ratings at the group level, we did observe evidence at the

[9] Our experiment was preregistered using the AsPredicted platform (https://aspredicted.org/bg6gi.pdf).

[10] We intentionally worded our instructions to say that the timing and cost of the emergency "may be different between attempts" to avoid any deception of participants (see Supplementary Materials for the complete set of instructions). In reality, the emergency was timed in both attempts to occur after the 15th round to allow for comparison with the income sequences used in Experiment 1.

[11] As with the previous experiments, the first emergency involved home repairs with a total cost of £4,500. The second emergency involved a car breakdown and cost £2,500.

[12] We recognise that we could have rewarded participants based on their performance on a randomly chosen attempt. By rewarding participants based on their cumulative performance, it is possible that they may have not treated the attempts as independent. For example, some participants may have hedged their bets by being risk averse in one attempt (saving lots) and risk seeking in the other (saving little).

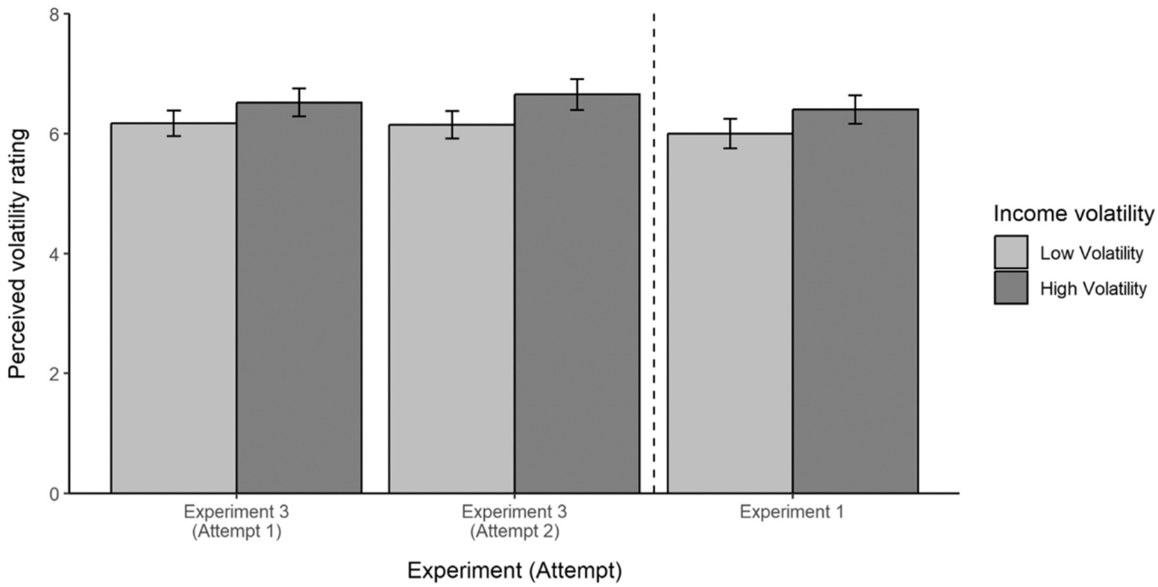

Fig. 4. Mean perceived volatility ratings by level of income volatility and task attempt in Experiment 3. Low and High Volatility conditions from Experiment 1 shown on right of dotted line for comparison. Standard error bars indicated.

individual participant level. We conducted a paired t-test comparing participants' ratings in their Low Volatility and High Volatility attempts. This indicated that there was anecdotal evidence of participants providing higher ratings during their High Volatility attempt compared to their Low Volatility attempt ($BF_{10} = 2.91$), with the mean difference between ratings being 0.43.

5.2.2. Additional pre-registered analyses

As part of our preregistration, we planned two additional analyses that were unrelated to our main hypotheses about perceived volatility ratings. First, we hypothesised that participants would save more in the High Volatility condition than in the Low Volatility condition (see Fig. 5). We examined this by conducting a Bayesian ANOVA where level of income volatility and task attempt were entered as variables. This provided Bayes factors for different combinations of predictive models, from which we could compute inclusion factors.[13] Our findings contradicted our hypothesis; there was moderate evidence against an income volatility level effect ($BF_{01} = 3.36$). Instead, we observed extreme evidence in favour of a task attempt effect ($BF_{10} = 5.01 \times 10^6$), with participants saving substantially more in their second attempt regardless of whether it involved the Low or High Volatility income sequence. There was also anecdotal evidence against an interaction between task attempt and income volatility level ($BF_{01} = 2.11$).

Our other additional hypothesis was that we should see similar perceived volatility ratings and final savings from participants completing their first attempts at the task as what was observed in Experiment 1. This would provide a signal of reliability for participants' behaviour within our task. We started by comparing the participants from Experiment 3 who received the Low Volatility income sequence first with participants in the Low Volatility condition from Experiment 1. There was moderate evidence to suggest that participants provided similar perceived volatility ratings (Fig. 4) and ended the task with similar final savings amounts (Fig. 5; $BF_{01} = 5.08$ and 4.51 respectively). When drawing the same comparison between the High Volatility counterparts, we again observed moderate evidence in favour of the perceived volatility ratings being similar ($BF_{01} = 5.44$). However, we found moderate evidence that participants' final savings differed ($BF_{10} = 5.50$); participants who received the High Volatility income sequence in their first task attempt in Experiment 3 appeared to have saved substantially less than those who were in the High Volatility condition from Experiment 1 (M = £6,834.69 vs £8,571.64).

5.3. Discussion

Experiment 3 provided some credence to our hypothesis that participants were sensitive to different levels of income volatility, but we were failing to capture this when asking for an absolute judgment on our 11-point scale (0 = "Not volatile"; 10 = "Extremely volatile"). This helped to explain why participants' ratings of income volatility across our previous experiments did not necessarily align with their observed saving behaviour. It also reintroduced the possibility of an income volatility 'threshold' that needed to be exceeded for participants to adjust their behaviour within the task. As a reminder, we observed in Experiment 1 that participants saved more in the High Volatility condition relative to those who had received a stable income, but that this difference was not observed for the Low Volatility condition.

However, before we could make any claims about a potential threshold, we first needed to resolve the inconsistency between the saving behaviour we had observed from participants in Experiment 1's High Volatility condition and the participants who received the High Volatility condition as their first attempt in Experiment 3. In theory, these participants should have exhibited similar saving behaviour; the only difference was that those in Experiment 3 would have known that they would be completing a second (independent) attempt at the task. The fact that we observed higher savings in Experiment 1 but not Experiment 3 called into question how reliable our original finding was, and by extension, whether there was an income volatility effect at all.

This led us to run a single additional experimental condition in Experiment 4 in which we further increased the income volatility that participants would experience ("Very High Volatility"; CV = 0.90). We expected that running this condition would help to build confidence in one of two conclusions. If these participants did not save any differently compared to participants who received stable incomes (in Experiment 1), we could be more confident that income volatility does not influence saving decisions, and that the increased saving effect we observed in Experiment 1 was likely a false positive. In contrast, if we did observe

[13] Inclusion factors are computed by comparing the posterior odds of models that include and exclude a predictor of interest (van Den Bergh et al., 2020). These help to quantify the level of evidence for whether the predictor should be included.

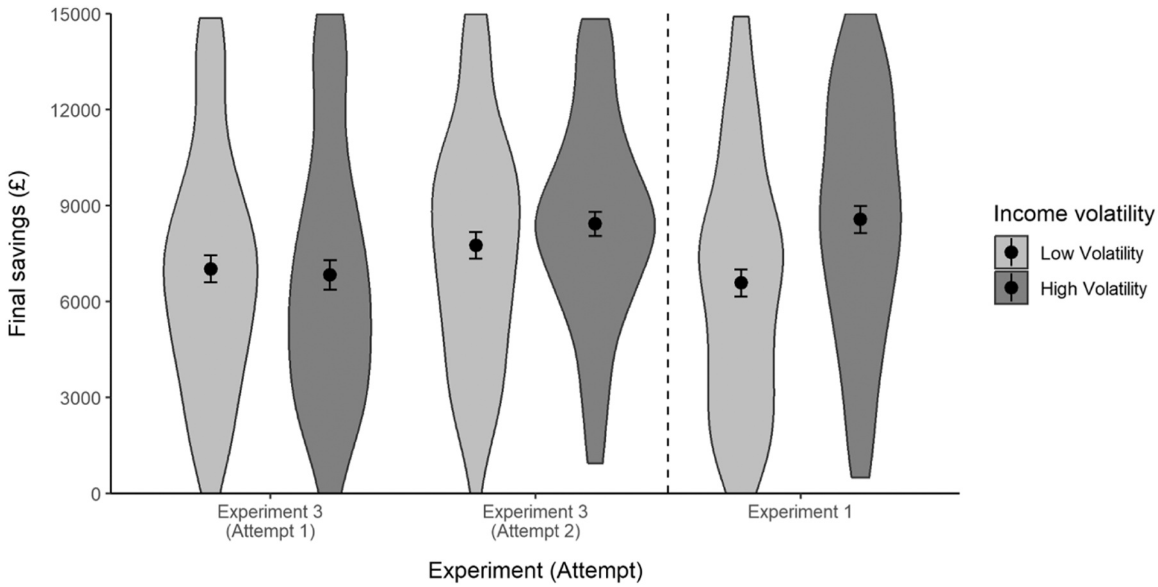

Fig. 5. Distribution of final savings by level of income volatility and task attempt in Experiment 3. Low and High Volatility conditions from Experiment 1 shown on right of dotted line for comparison. Means and standard error bars indicated.

higher savings in this Very High Volatility condition, we would instead have greater support for an effect of income volatility on saving behaviour and consider our results from Experiment 3 to be a false negative.

6. Experiment 4

In Experiment 4, we generated a new income sequence that was more volatile than any of our previous conditions. This condition was known as the Very High Volatility condition and had a CV of 0.90. All participants were placed in this condition and completed only one attempt at the financial decision making task.

6.1. Method

6.1.1. Participants

We recruited 80 new participants from Prolific Academic using the same eligibility criteria as described in the previous experiments. Our sample size was based on having the same number of participants per condition as Experiment 1. Participants received £3.00 (USD 3.77) in exchange for their participation and were also eligible for an additional bonus of £3.00 (USD 3.77) based on their performance in the financial decision making task.

Participants' ages ranged between 19 and 65 years ($M_{\text{age}} = 36.63$, SD $= 11.04$), 42.5% of which were male and 57.5% female. In terms of education level, nearly all participants were either high school graduates (35.0%), held an undergraduate degree (31.2%), or held a postgraduate degree (30.0%). Most participants were employed either full-time (48.8%) or part-time (21.2%), with the median personal income band being £20,000-£29,999.

6.1.2. Materials and procedure

Experiment 4 used the same financial decision making task and followed the same procedure as Experiment 1. All participants—henceforth referred to as the Very High Volatility condition—received a newly generated income sequence that ranged from £44 to £3,457, with the CV of the sequence being 0.90 (see Table 1). Consistent with all other income sequences, the total income received across the task remained at £15,000.

6.2. Results

In this section, we first compare perceived volatility ratings and final savings between our Very High Volatility condition and the No Volatility condition in Experiment 1. We then report on analyses using our combined data across all four experiments, where we grouped conditions with the same level of income volatility. As part of these analyses, we pooled our conditions from Experiment 2 into a group labelled "Medium Volatility"; these participants all received an income sequence with a CV of 0.37—between the CVs of Experiment 1's Low Volatility condition (CV = 0.30) and High Volatility condition (CV = 0.60). We then pooled participants' first attempts from Experiment 3 with their respective Experiment 1 conditions (i.e., Low or High Volatility). Finally, we excluded data from participants' second attempts in Experiment 3; our prior results suggested that participants tended to save more on their second attempt (regardless of level of income volatility), rendering these data incomparable to our other conditions, where participants were completing their first or only attempt.

6.2.1. Comparison of Very High Volatility to No Volatility

We observed extreme evidence that participants perceived their income to be substantially more volatile in the Very High Volatility condition ($M = 6.79$) compared to the No Volatility condition ($M = 3.17$) ($BF_{10} = 1.03 \times 10^{13}$). There was also strong evidence that participants saved more in the Very High Volatility condition than the No Volatility condition ($M =$ £8,342.80 vs £6,647.88) ($BF_{10} = 10.97$).

6.2.2. Combined analysis on income volatility effect

A Bayesian ANOVA indicated that there was extreme evidence that perceived volatility ratings differed between income volatility levels ($BF_{10} = 4.62 \times 10^{25}$). However, as evident from Fig. 6, this effect appears to be entirely driven by lower volatility ratings in Experiment 1's No Volatility condition. When excluding this condition from the analysis, there was anecdotal evidence that the remaining conditions did not differ in their volatility ratings ($BF_{01} = 1.41$).

A similar analysis of participants' final savings yielded inconclusive evidence as to whether they differed as a function of income volatility level ($BF_{01} = 1.01$) (see Fig. 7). We thus chose to exploit the fact that each income volatility level was quantifiable based on its CV and that this could be treated as a continuous measure. When instead analysing final savings as a function of CV, we observed very strong evidence that a

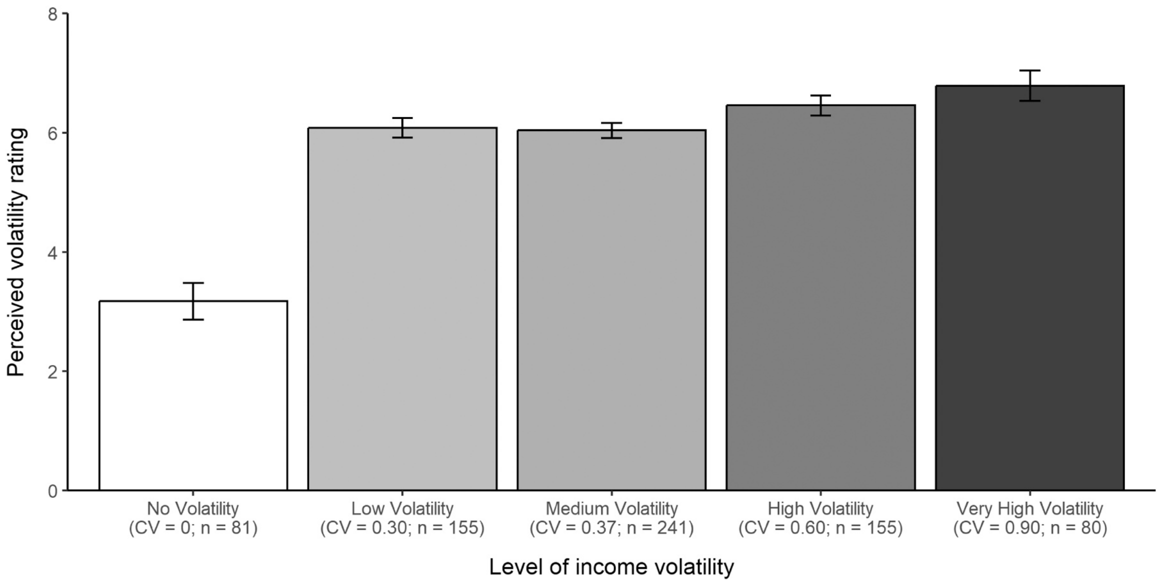

Fig. 6. Mean perceived volatility ratings by level of income volatility across Experiments 1–4. Standard error bars indicated.

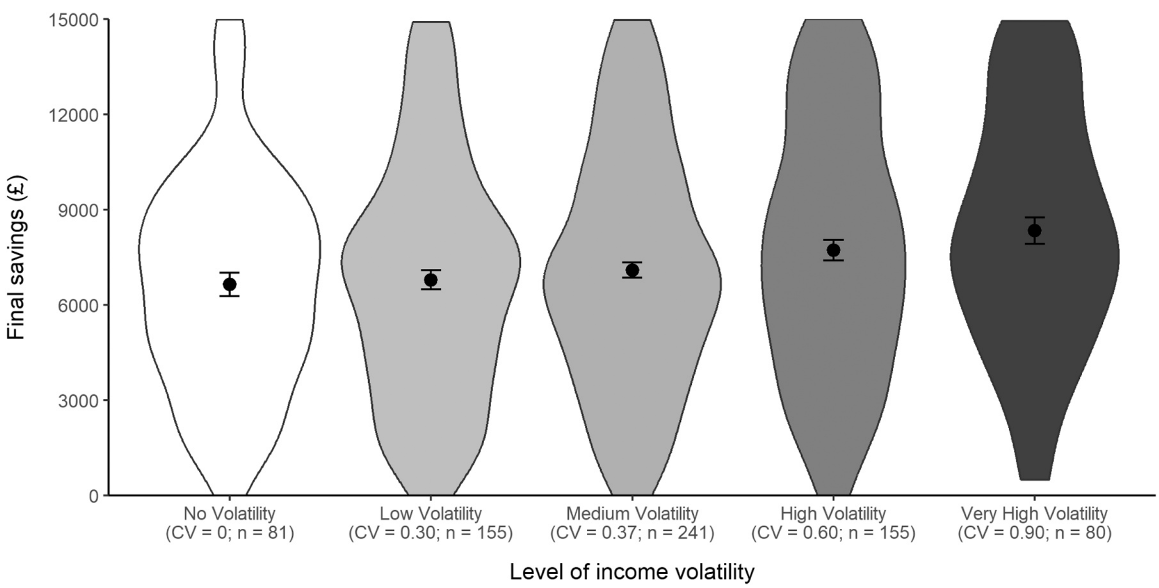

Fig. 7. Distribution of final savings by level of income volatility across Experiments 1–4. Means and standard error bars indicated.

higher CV was associated with more savings ($BF_{10} = 45.89$). A Bayesian regression approach yielded a median posterior estimate for the CV effect to be £2,134.69. This suggested that a 0.1 increase in CV corresponded to an average increase in savings of approximately £213 within our task (95% HDI: [£92, £328])—equivalent to a 1.42pp increase in savings (when considering the maximum possible savings amount of £15,000).

7. General discussion

7.1. Income volatility and saving behaviour

Across four experiments, our collective results provide preliminary evidence that people save more when faced with fluctuations in income and are consistent with the view that uncertainty drives precautionary savings motives (Carroll & Kimball, 2018). However, this interpretation comes with two caveats. First, it should be reiterated that our primary finding is based on quasi-experimental evidence. Our experimental analyses of the relationship between income volatility and saving behaviour yielded inconsistent results. Experiment 1 saw an increase in saving behaviour for the High Volatility condition relative to the Low Volatility condition, but this was not replicated in Experiment 3. Only when comparing data across experiments did we observe more reliable effects, which should naturally raise questions about the robustness of our findings. Second, leaving aside the quasi-experimental nature of our evidence, the effect sizes we observed could be considered economically small. Our results imply that individuals experiencing a typical amount of income volatility (CV = 0.30) would save 4.26 percent more than individuals experiencing no volatility (CV = 0).[14]

Both caveats suggest that it may be more appropriate to conclude that substantial levels of income volatility are required to meaningfully

[14] While we have focused on mean differences in saving behaviour throughout our study, we also conducted exploratory analyses examining the behaviour of individual participants (see Appendix B).

shift saving behaviour. This alternative interpretation would be consistent with prior research in other contexts that has found individuals to be surprisingly insensitive to volatility. For example, Ehm et al. (2014) observed a 'volatility inadaptability bias' when asking participants to indicate how much they would invest in an asset given its expected returns and volatility. Participants chose to invest similar proportions of their allotted funds regardless of whether their asset had a low or high risk-return profile, resulting in the latter group having portfolios that were over twice as volatile as the former. Along similar lines, Parker (2017) found little evidence of households increasing their spending when government stimulus payments turned out to be larger than anticipated. Their analysis suggests that contrary to theoretical predictions, consumers were not adjusting their consumption in response to receiving a windfall—which can be otherwise thought of as a form of (positive) income volatility.

It could also be the case that the effect of income volatility on saving behaviour is not linear; rather, there is a threshold of volatility that must be exceeded before individuals feel the need to save more. This idea of a threshold aligns with the concept of 'just noticeable differences' that is well-established in research on perception (e.g., Bradley et al., 1999; Quené, 2007; Zhang et al., 2008). Our results suggest that if such a threshold did exist, it would lie around a CV level of 0.60, which would be about twice the amount of volatility experienced by a typical household. Below this level (i.e., in the Low Volatility condition of Experiment 1 and the three Experiment 2 conditions), saving behaviour did not differ relative to participants who received steady income. At this level (i.e., our High Volatility condition), we observed inconsistent results—increased saving behaviour in Experiment 1 but not in Experiment 3. Above this level (i.e., in the Very High Volatility condition of Experiment 4), we observed a clear increase in savings.

7.2. Perceptions of income volatility

Our findings also highlight the importance of carefully considering how one chooses to define and measure income volatility. Our analysis indicated that there was a relationship between participants' saving behaviour and the objective level of income volatility experienced (as defined by CV). However, we would not have observed a relationship had we relied on participants' self-reported ratings of income volatility (see Appendix C for this supplementary analysis). Given the clear discrepancies we observed between perceived and objective income volatility throughout our experiments (participants provided similar ratings regardless of whether their income had a CV of 0.30 or 0.90), this was hardly surprising. It does, however, illustrate the potential pitfalls in relying on self-reported measures of income volatility when conducting financial decision making research. This is especially true for studies which rely on coarse measures. For example, when surveying households on how much their income fluctuates, the Consumer Financial Protection Bureau (2023) only allow for three possible responses: "About the same each month", "Varies somewhat from month to month", and "Varies a lot from month to month". An alternative measurement approach, while not necessarily perfect, could be to instead ask respondents to estimate their monthly income over a prior period and to then calculate an 'objective' metric.

This is not to say that self-reported measures of income volatility are invalid. Subjective ratings can provide useful insight into how individuals perceive their situation and, in the context of income volatility, what level of instability they deem tolerable. This is important to research as there is bound to be heterogeneity in individuals' preferences—just as there are people who report having low financial wellbeing despite appearing objectively well-off and vice versa (Comerton-Forde et al., 2018). The point we instead seek to emphasise is that perceived and objective income volatility need to be viewed as related, but ultimately separate constructs.

Recognising this distinction suggests several directions that future researchers might choose to extend upon our work. First, future work may seek to compare CV against other measures of objective income volatility, such as the second-order coefficient of volatility proposed by Kvålseth (2017). There may also be opportunities to build upon our second experiment and further investigate whether other characteristics of income sequences better explain either people's perceptions of their volatility or people's subsequent saving decisions. While we examined the role of the minimum and maximum incomes received based on prior work by Konovalova and Pachur (2021), there remain many other characteristics that could be worth investigating—for example, the skewness of the income sequence (Holzmeister et al., 2020), or whether the sequence is generally increasing or decreasing (Loewenstein & Sicherman, 1991). Finally, it may be interesting to explore the extent to which predictability moderates any income volatility effects (C. Y. Zhang & Sussman, 2023). For some individuals, income may fluctuate between pay periods, but in a predictable manner. This might occur if income is correlated with the time of year (e.g., seasonal work or annual bonuses) or if individuals have control over how much they work and therefore their earnings (Peetz et al., 2021). It may turn out to be the case that the degree of volatility in one's income is unimportant, and that what really matters is its level of predictability.

7.3. Practical implications

Our findings also suggest that individuals who experience substantial fluctuations in their income may be motivated to save more. This presents an opportunity for financial services institutions to develop interventions that help consumers follow through with these intentions. Banks, for example, have visibility over their customers' income and can therefore objectively identify which of their customers have highly volatile incomes. For these customers, the intervention could be as simple as deploying a timely message when a large spike in income is detected (relative to the customers' usual earnings). Such 'just-in-time' interventions can be highly effective (often used in health settings; e.g., Forman et al., 2019; van Der Laan & Orcholska, 2022); in this case, potentially by combating the natural tendency to spend more of money that is considered a bonus or windfall (Arkes et al., 1994; Epley & Gneezy, 2007).

7.4. Study limitations

As described in the Introduction, our study was motivated by the observation that there had been little prior experimental research examining the impact of intra-year income volatility on saving behaviour. On the one hand, we believe that the experimental findings we have reported make a useful contribution to this otherwise sparse literature. At the same time, we recognise there are several opportunities to improve upon our approach—as is bound to be the case when conducting research in areas without well-established experimental paradigms or procedures. We highlight below three key limitations in our study and provide suggestions for how they could be overcome in future research.

7.4.1. Establishing expected or optimal saving behaviour

First, the nature of the financial decision making task we have created does not allow for analyses of what behaviour should be expected or considered optimal. Unlike other saving tasks used in previous research (e.g., Brown et al., 2009), there are too many uncertainties involved—the distribution of possible incomes, emergency costs, and emergency timings—that render the required calculations intractable. As noted earlier in the paper, this was by design; our focus was on understanding how income volatility influenced saving decisions (not whether this was the 'right' or 'wrong' thing to do) and we wanted the task to resemble the uncertainties of real world financial decisions to the extent that was possible. However, because of our design choice, we are constrained in our ability to draw further conclusions about the behaviour we observed across our experiments. We cannot, for example, try to account for participants' saving decisions using existing models of behaviour, either normative (e.g., Expected Utility Theory) or behavioural (e.g., Prospect Theory). We likewise cannot make claims about whether participants are behaving rationally or in a systematically biased manner.

Future studies could overcome this limitation by reducing the degree of uncertainty within the task slightly: enough so that it becomes possible to predict how participants should behave under different models, but not so much that it no longer can be considered representative of the real-world decisions that consumers make. For example, before making any saving decisions within the task, participants could be shown a representative sample of incomes that corresponds to their experimental condition's level of volatility. In doing so, participants would be able to make informed predictions about how much income they should receive throughout the task (as is the case in real life), even if the exact income on each round remains uncertain. Along the same lines, participants could be given more information about the timing of the financial emergency. For instance, participants could be told that there is a fixed probability of the emergency occurring on each round (e.g., 5 percent)—akin to the fact that in the real world, a financial emergency could emerge at any time. Another option would be to provide participants with a simple distribution of possible emergency timings (e.g., it being equally likely to occur between rounds 10 and 20)—which would be more analogous to saving for retirement, where consumers may know roughly, but not precisely, when they will retire.

7.4.2. Explaining the income volatility effect

A second limitation is our inability to identify the exact mechanism that drove the differences in saving behaviour between our experimental conditions. Evidently, participants were sensitive to the volatility of their income; they tended to save more when we manipulated their income to be more volatile. However, these saving decisions did not appear to be linked to participants' self-reported ratings of how volatile they perceived their income to be (see Appendix C). We are therefore able to conclude that participants' saving behaviour is influenced by their income's volatility but are unable to conclude why this is the case.

Future work could consider collecting additional information about participants' subjective beliefs surrounding the task. For example, participants could be asked to make estimates about the minimum, average, or maximum income they expect to receive. They could also be asked to estimate the riskiness of their income (e.g., how likely their income would fall below or exceed a given amount). Similar predictions could be elicited around the financial emergency. These self-reported measures could then be used to support or refute alternative explanations for what is driving the effect of income volatility on saving behaviour.

7.4.3. Incentivisation and external validity

A final limitation relates to the way we incentivised participants across the experiment. In hindsight, there were two potential issues with our approach that may impact the external validity of our findings. First, the points system we implemented could be considered unrealistic because of how severely participants were punished for failing to save adequately for the financial emergency. In our task, participants who had not saved enough lost all the points they had previously accrued. This is of course not reflective of real life, where even the extreme case of bankruptcy does not result in losing all the benefits of prior consumption. Consequently, participants' saving decisions within our task may not be representative of the decisions they would make under normal circumstances; they may have felt pressured to save more to avoid our extreme penalty. While our intention behind this penalty was to provide a clear deterrent from spending all of one's income, we recognise that this could have been achieved through a more realistic points system. For example, rather than losing all their points for failing to save adequately for the emergency, a future study might instead deduct a flat number of points (e.g., 5,000 points). This would more accurately represent the real world, in which the utility consumers gain from consumption is not retroactively affected by their ability to meet future saving obligations.

The other potential issue is our use of tournament incentives, whereby we rewarded participants with a bonus payment (on top of their fixed payment) if they were within the top 10 percent of scores. Our primary motivation for this incentive scheme was to have greater control over the payouts for the experiment. As this was a novel task and we did not have a reference point for how much participants would score; we were concerned that rewarding participants proportionally to their accumulated points might lead us to pay participants inadequately (if they scored less than anticipated) or lead us to exceed our available experiment budget (if scores were higher than anticipated). It is possible that this incentive scheme encouraged participants to engage in greater risk-taking behaviour than is typical (by spending more), as they may have assumed that it was the only way to end up in the top 10 percent of scores. Future work may seek to avoid this issue by reverting to a simple linear incentive scheme (e.g., £1.00 per 1,000 points earned). If doing so, it may be beneficial to start with a small-scale pilot study to gauge how much participants typically save and adjust the payment rate as necessary.

8. Conclusion

Recent economic trends, such as the rapid rise of the gig economy, have made it increasingly common for individuals to have incomes that are unstable and fluctuate within a year. Our study set out to investigate how experiencing this income volatility affects saving decisions—and to our knowledge, is one of the first to explore this experimentally. Using a novel financial decision making task, we provide (quasi-experimental) evidence that higher income volatility results in increased saving behaviour. The present research suggests many fruitful avenues for further work that will improve our understanding of the relationship between income volatility and saving and allow us to better tailor financial products and services to individuals who experience fluctuations in their income.

Funding organisations

This work was supported by the Australian Government Research Training Program (RTP) scholarship.

Declarations of interest

N. Wang-Ly: None.
B. R. Newell: None.
Institutional Review Board Approval:
This study received University of New South Wales Human Ethics Committee approval HC3693. All experimental subjects gave informed consent.

CRediT authorship contribution statement

Nathan Wang-Ly: Writing – review & editing, Writing – original draft, Methodology, Formal analysis, Conceptualization. **Ben Rhodri Newell:** Writing – review & editing, Methodology, Formal analysis, Conceptualization.

Appendix A. Demographic and financial characteristics of our samples

Table A1
Demographic characteristics of samples in Experiments 1–4.

	Exp 1	Exp 2	Exp 3	Exp 4	Overall	Pop
N	241	241	150	80	712	-
Age (%)						
18–24 years old	11.6	10.8	9.3	17.5	11.5	13.5
25–34 years old	27.4	27.4	34.7	25.0	28.7	21.8
35–44 years old	26.1	28.2	24.0	31.2	27.0	21.4
45–54 years old	20.3	20.3	20.0	17.5	19.9	20.9
55–64 years old	13.7	12.4	12.0	6.3	12.1	20.7
65 years or older	0.8	0.8	0.0	1.3	0.7	1.7
Prefer not to say	0.0	0.0	0.0	1.3	0.1	-
Gender (%)						
Male	59.8	44.8	50.0	42.5	50.7	45.4
Female	39.4	53.9	49.3	57.5	48.5	48.5
Non-binary/Gender fluid	0.4	0.8	0.7	0.0	0.6	0.1
Different identity	0.4	0.0	0.0	0.0	0.1	0.0
Prefer not to say	0.0	0.4	0.0	0.0	0.1	6.0
Education level (%)						
Did not complete high school	0.8	1.2	1.3	2.5	1.3	18.6
High school graduate	36.1	41.1	42.0	35.0	38.9	30.1
Undergraduate degree (e.g., Bachelor's degree)	45.2	37.8	40.7	31.2	40.2	35.5
Postgraduate degree (e.g., Master's or Doctoral degree)	17.4	19.1	16.0	30.0	19.1	15.8
Prefer not to say	0.4	0.8	0.0	1.3	0.6	-

Note. Pop = Population. Sources for population statistics: OECD (2023) and Office for National Statistics (2023b, 2023c).

Table A2
Financial characteristics of samples in Experiments 1–4.

	Exp 1	Exp 2	Exp 3	Exp 4	Overall	Pop
N	241	241	150	80	712	-
Employment status (%)						
Full-time	61.0	52.7	55.3	48.8	55.6	50.8
Part-time	16.2	20.3	16.0	21.2	18.1	20.1
Contract/Temporary	2.5	1.7	0.0	1.3	1.5	-
Student	4.2	6.2	3.3	10.0	5.3	10.1
Unemployed	5.0	8.3	13.3	6.3	8.0	9.7
Unable to work	5.0	2.5	4.7	6.3	4.2	5.3
Other	5.0	7.8	7.3	5.0	6.3	4.0
Prefer not to say	1.7	0.4	0.0	1.3	0.8	-
Pay frequency (%)[a]						
Weekly	10.0	11.6	8.0	7.5	9.8	11.6
Fortnightly	5.0	2.1	2.7	1.3	3.1	1.5
Monthly	72.2	68.5	68.0	62.5	69.0	86.9
Other	4.2	5.8	5.3	3.8	4.9	-
Not Applicable	7.1	10.8	15.3	18.8	11.4	-
Prefer not to say	1.7	1.2	0.7	6.3	1.8	-
Annual personal income (%)[b]						
Negative or nil income	1.7	4.6	4.0	6.3	3.7	-
£0-£9,999	16.6	15.8	18.0	17.5	16.7	-
£10,000-£19,999	16.2	14.1	16.7	11.2	15.0	28.0
£20,000-£29,999	20.3	25.7	25.3	20.0	23.2	29.0
£30,000-£39,999	16.2	13.3	14.0	15.0	14.6	15.0
£40,000-£49,999	8.7	5.8	7.3	10.0	7.6	9.0
£50,000-£59,999	5.8	4.2	4.7	1.3	4.5	4.0
£60,000-£69,999	3.3	4.2	2.7	5.0	3.7	2.0
£70,000 or more	5.4	4.6	1.3	7.5	4.5	13.0
Prefer not to say	5.8	7.9	6.0	6.3	6.6	-

Note. Pop = Population. Sources for population statistics: HM Revenue & Customs (2023) and Office for National Statistics (2022, 2023a).
[a]Population data for pay frequency combines Weekly and Other frequencies. Four-weekly pay frequency has been treated as Monthly in our calculations. [b] Population data for annual personal income exclude individuals with no income tax liability.

Table A3
Comparison of participant sample characteristics between Experiments 1–4.

Characteristic	Inverse Bayes factor (BF_{01})			
	Exp 1 vs Exp 2	Exp 1 vs Exp 3	Exp 1 vs Exp 4	All Exps
Age	1.71×10^5	2.38×10^5	291.76	7.95×10^{11}
Gender	324.08	1.30×10^5	812.89	5.75×10^{11}

(*continued on next page*)

Table A3 (*continued*)

Characteristic	Inverse Bayes factor (BF_{01})			
	Exp 1 vs Exp 2	Exp 1 vs Exp 3	Exp 1 vs Exp 4	All Exps
Education level	2,335.28	3,509.91	19.42	3.70×10^{8}
Employment status	1,589.83	129.49	1,425.54	9.77×10^{9}
Pay frequency	1,681.36	265.96	3.70	4.60×10^{7}
Annual personal income	6,847.00	4,777.58	369.42	2.93×10^{11}

Note. For each characteristic (e.g., age), we performed a Bayesian test to compare the proportion of participants in each experiment who fell under each response category (e.g., 18–24 years old)—analogous to the chi-square test used in frequentist statistics. We report the inverse Bayes factor for each test, which indicates the relative likelihood of our observed data if there were no differences between our samples compared to if there were differences.

Appendix B. Individual-level saving behaviour

While our main results focus on the average saving behaviour of our experimental conditions, here we share exploratory analyses examining the saving behaviour of individual participants.[15] We used a K-means clustering approach to investigate whether there were consistent patterns in how participants were choosing to approach the task.

We started by determining the set of data features that would be used to cluster participants. While we did not have strong hypotheses about what clusters to expect, there were two saving strategies we considered plausible. First, participants might aim to save the same absolute amount per round where permitted by their income (e.g., £300). Second, participants might instead aim to save the same relative amount of their income each round (e.g., 20 percent). Thus, our set of features included the mean amount participants saved per three-round block in the task. We also included the relative amount saved per block—that is, what proportion of the income earned during that block did participants save. As added measures, we included the fit of participants' savings trajectories to both linear and quadratic models. All features were normalised before clustering.

We then examined the incremental variance in participants' saving behaviour that was explained by increasing the number of clusters—starting from one and ending with ten. As shown in Figure B.1, there appeared to be diminishing returns beyond having four clusters. Thus, we proceeded with the remainder of our analysis on the assumption that participants' behaviour could be grouped into four clusters.

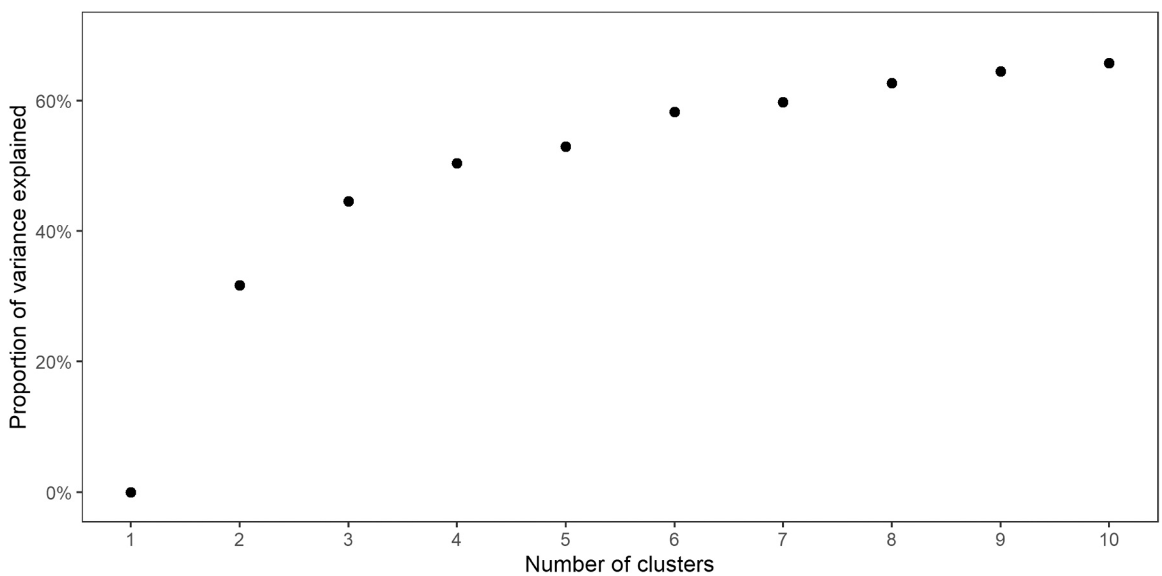

Figure B1. Proportion of variance in features explained by number of clusters specified.

Table B.1 shows the number of participants assigned to each cluster and the median feature values per cluster (also see Figure B.2 for example savings trajectories). Consistent with our predictions, the clusters appeared to reflect groups of participants who tended to save similar relative amounts (Cluster 1: 'Proportion Savers') and similar absolute amounts (Cluster 3: 'Amount Savers') of their income each round. Of the two remaining clusters, one group of participants appeared to save during the earlier rounds of the task, but then noticeably reduce their saving rate in the later rounds (Cluster 2: 'Early Savers'). These were potentially the participants who were aiming towards saving a set amount to buffer against the task's financial emergency and so began to feel comfortable once they reached or were close to this target. The final group appeared to save little overall (Cluster 4: 'Low Savers')—perhaps because they were either underestimating how much the emergency would require or were overestimating how much time they would have to save.

[15] As described at the start of Section 6.2, we combined our data across all four experiments for this analysis.

Table B1
Median saving behaviour data features per cluster.

	Cluster 1:Proportion Savers	Cluster 2:Early Savers	Cluster 3: Amount Savers	Cluster 4: Low Savers
N	180	94	330	108
Mean savings				
Block 1	800.00	767.33	465.00	407.00
Block 2	782.17	677.17	388.83	140.67
Block 3	803.83	441.00	400.00	0.00
Block 4	799.33	333.33	365.33	0.00
Block 5	816.33	178.33	450.83	39.67
Relative savings				
Block 1	0.86	0.76	0.49	0.43
Block 2	0.84	0.61	0.42	0.13
Block 3	0.84	0.49	0.44	0.00
Block 4	0.83	0.40	0.41	0.00
Block 5	0.80	0.25	0.45	0.04
Linear model fit				
Slope	808.81	479.46	419.35	82.32
Intercept	8.59	1,125.34	114.59	955.86
R^2	0.99	0.92	0.97	0.53
Quadratic model fit				
Slope (quadratic)	1.42	-27.15	-0.64	-7.28
Slope (linear)	761.35	915.57	425.98	180.51
Intercept	4.51	-80.05	131.61	491.45
R^2	0.99	0.98	0.98	0.68

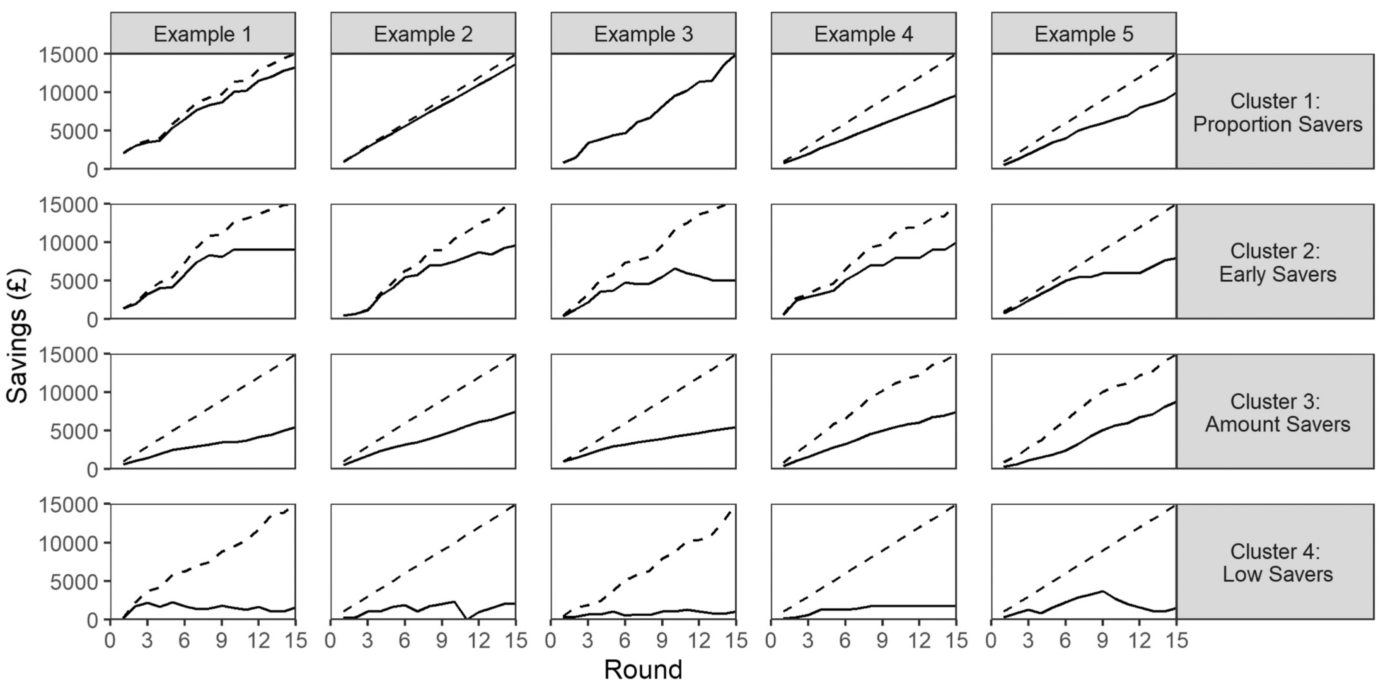

Figure B2. Example savings trajectories for each cluster. Solid line indicates actual savings across rounds. Dashed line indicates cumulative income received across rounds.

Finally, we compared the proportion of participants assigned to each cluster as a function of the level of income volatility experienced (defined by CV, the coefficient of variation) (see Figure B.3). We note two potentially interesting observations that may be worth exploring further in future research. First, that participants appeared to be less likely to save consistent amounts (i.e., to be 'Amount Savers') as income volatility increased. This is somewhat to be expected, as large variations in income (whether increases or decreases) naturally make it difficult to stick to saving the same amount. Second, that as incomes became more volatile, participants appeared to be more likely to save consistent proportions of their income (i.e., to be 'Proportion Savers'). This could be because participants were employing a rule of thumb (e.g., always save 20 percent of income) to reduce the difficulty of knowing how much to save in the face of income uncertainty.

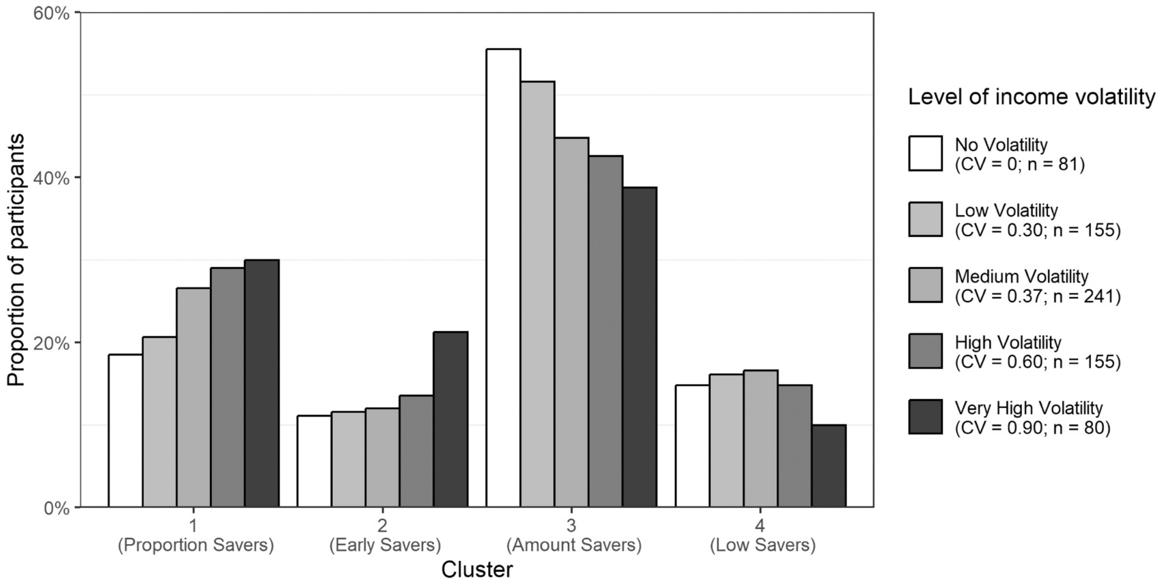

Figure B3. Proportion of participants assigned to each cluster as a function of level of income volatility experienced.

Appendix C. Subjective and objective income volatility as predictors of saving behaviour

As described in Section 6.2.2, we conducted a Bayesian regression analysis to analyse the relationship between participants' final savings (across all four experiments) and their objective level of income volatility (as defined by CV, the coefficient of variance). This yielded very strong evidence that a higher CV was associated with more savings ($BF_{10} = 45.89$).

We estimate the size of this effect using the *rstanarm* package (v2.21.4; Goodrich et al., 2023) in *R 4.2.0* (R Core Team, 2022). We apply the inbuilt Markov chain Monte Carlo (MCMC) sampling approach (four chains with 2,000 iterations each) using default priors to produce the posterior distribution of our regression coefficients. Table C.1 presents the output of this analysis. It shows the median estimate for the CV coefficient (2,134.69), which implies that an increase of CV of 0.1 corresponded to an average increase in participants' final savings of about £213 (95% HDI: [£92, £328]). Figure C.1 presents this information visually.

Table C1
Summary of posterior distributions of regression with savings and CV.

Parameter	Median	95% CI	PD	ROPE (%)	R-hat	ESS
(Intercept)	6,360.17	[5,783.66, 6,902.41]	100%	0%	1.000	4,113
CV	2,134.69	[920.05, 3,276.66]	100%	0%	1.000	4,170

Note. PD = Probability of direction. ROPE = Region of Practical Equivalence (set by default at a range of ±0.1 standard deviations; Kruschke & Liddell, 2018). ESS = Effective Sample Size.

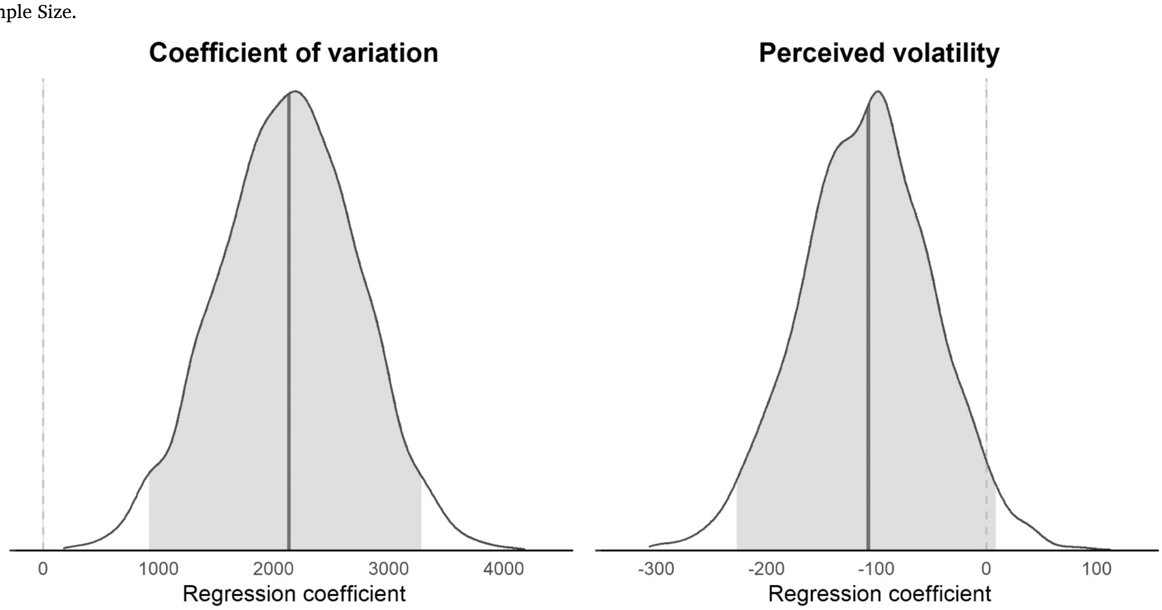

Figure C1. Posterior distributions of CV and perceived volatility coefficients. Thick, dark grey lines indicate median estimates. Light grey shaded areas indicate 95% confidence intervals.

We applied the same approach to examine the relationship between participants' subjective ratings of their income's volatility ('perceived volatility') and their final savings. However, contrary to our previous result, we observed anecdotal evidence of there being no relationship between the two variables ($BF_{01} = 2.32$). Table C.2 presents the output of this analysis, with the posterior distribution being plotted in Figure C.1.

Table C2
Summary of posterior distributions of regression with savings and perceived volatility.

Parameter	Median	95% CI	PD	ROPE (%)	R-hat	ESS
(Intercept)	7,906.16	[7,114.11, 8,605.00]	100%	0%	1.001	3,952
Perceived volatility	-107.61	[-225.23, 8.91]	96.7%	0%	1.001	4,008

Note. PD = Probability of direction. ROPE = Region of Practical Equivalence (set by default at a range of ±0.1 standard deviations; Kruschke & Liddell, 2018). ESS = Effective Sample Size.

Appendix D. Supporting information

Supplementary data associated with this article can be found in the online version at doi:10.1016/j.jbef.2024.100941.

References

Anvari-Clark, J., Ansong, D., 2022. Predicting financial well-being using the financial capability perspective: the roles of financial shocks, income volatility, financial products, and savings behaviors. J. Fam. Econ. Issues *43*, 730–743. https://doi.org/10.1007/s10834-022-09849-w.

Arkes, H.R., Joyner, C.A., Pezzo, M.V., Nash, J.G., Siegel-Jacobs, K., Stone, E., 1994. The psychology of windfall gains. Organ. Behav. Hum. Decis. Process. *59* (3), 331–347. https://doi.org/10.1006/obhd.1994.1063.

Bania, N., Leete, L., 2022. Monthly income volatility and health outcomes. Contemp. Econ. Policy *40* (4), 636–658. https://doi.org/10.1111/coep.12580.

Bradley, J.S., Reich, R., Norcross, S.G., 1999. A just noticeable difference in C50 for speech. Appl. Acoust. *58* (2), 99–108. https://doi.org/10.1016/S0003-682X(98)00075-9.

Brown, A.L., Chua, Z.E., Camerer, C.F., 2009. Learning and visceral temptation in dynamic saving experiments. Q. J. Econ. *124* (1), 197–231. https://doi.org/10.1162/qjec.2009.124.1.197.

Camerer, C., Mobbs, D., 2017. Differences in behavior and brain activity during hypothetical and real choices. Trends Cogn. Sci. *21* (1), 46–56. https://doi.org/10.1016/j.tics.2016.11.001.

Carroll, C.D., Kimball, M.S., 2018. Precautionary saving and precautionary wealth. In: Macmillan Publishers Ltd (Ed.), The new Palgrave dictionary of economics. Palgrave Macmillan UK, pp. 10591–10599. https://doi.org/10.1057/978-1-349-95189-5_2079.

Chamon, M., Liu, K., Prasad, E., 2013. Income uncertainty and household savings in China. J. Dev. Econ. *105*, 164–177. https://doi.org/10.1016/j.jdeveco.2013.07.014.

Chapman, G.B., 1996. Temporal discounting and utility for health and money. Article 3. J. Exp. Psychol.: Learn., Mem., Cogn. *22* (3) https://doi.org/10.1037/0278-7393.22.3.771.

Cobb-Clark, D.A., Kassenboehmer, S.C., Sinning, M.G., 2016. Locus of control and savings. J. Bank. Financ. *73*, 113–130. https://doi.org/10.1016/j.jbankfin.2016.06.013.

Collins, J.M., Gjertson, L., 2013. Emergency savings for low-income consumers. Focus *30* (1), 12–17.

Comerton-Forde, C., Ip, E., Ribar, D.C., Ross, J., Salamanca, N., Tsiaplias, S., 2018. Using survey and banking data to measure financial wellbeing. No. 1; Commonwealth Bank of Australia and Melbourne Institute Financial Wellbeing Scales Technical Report, Issue No. 1. ⟨https://fbe.unimelb.edu.au/_data/assets/pdf_file/0010/2839429/CBA_MI_Tech_Report_No_1.pdf⟩.

Consumer Financial Protection Bureau, 2023. Making ends meet in 2023: Insights from the Making Ends Meet survey. 2023–8; CFPB Office of Research Publication. ⟨https://files.consumerfinance.gov/f/documents/cfpb_making-ends-meet-in-2023_report_2023-12.pdf⟩.

Diaz-Serrano, L., 2005. Income volatility and residential mortgage delinquency across the EU. J. Hous. Econ. *14* (3), 153–177. https://doi.org/10.1016/j.jhe.2005.07.003.

Drèze, J.H., Modigliani, F., 1972. Consumption decisions under uncertainty. J. Econ. Theory *5* (3), 308–335.

Ehm, C., Kaufmann, C., Weber, M., 2014. Volatility inadaptability: Investors care about risk, but cannot cope with volatility. Rev. Financ. *18* (4), 1387–1423. https://doi.org/10.1093/rof/rft032.

Epley, N., Gneezy, A., 2007. The framing of financial windfalls and implications for public policy. J. Socio-Econ. *36* (1), 36–47. https://doi.org/10.1016/j.socec.2005.12.012.

Farrell, D., Greig, F., 2016. JPMorgan Chase Institute. Paychecks, paydays, Online Platf. Econ. ⟨https://www.jpmorganchase.com/content/dam/jpmc/jpmorgan-chase-and-co/institute/pdf/jpmc-institute-volatility-2-report.pdf⟩.

Farrell, D., Greig, F., Yu, C., 2019. Weathering volatility 2.0: A monthly stress test to guide savings. JPMorgan Chase Institute. ⟨https://www.jpmorganchase.com/content/dam/jpmc/jpmorgan-chase-and-co/institute/pdf/institute-volatility-cash-buffer-report.pdf⟩.

Faul, F., Erdfelder, E., Lang, A.-G., Buchner, A., 2007. G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. Behav. Res. Methods *39* (2), 175–191. https://doi.org/10.3758/BF03193146.

Fawcett, T.W., McNamara, J.M., Houston, A.I., 2012. When is it adaptive to be patient? A general framework for evaluating delayed rewards. Behav. Process. *89* (2), 128–136. https://doi.org/10.1016/j.beproc.2011.08.015.

FeldmanHall, O., Mobbs, D., Evans, D., Hiscox, L., Navrady, L., Dalgleish, T., 2012. What we say and what we do: The relationship between real and hypothetical moral choices. Cognition *123* (3), 434–441. https://doi.org/10.1016/j.cognition.2012.02.001.

Forman, E.M., Goldstein, S.P., Crochiere, R.J., Butryn, M.L., Juarascio, A.S., Zhang, F., Foster, G.D., 2019. Randomized controlled trial of OnTrack, a just-in-time adaptive intervention designed to enhance weight loss. Transl. Behav. Med. *9* (6), 989–1001. https://doi.org/10.1093/tbm/ibz137.

Fox, C.R., Tversky, A., 1995. Ambiguity aversion and comparative ignorance. Q. J. Econ. *110* (3), 585–603. https://doi.org/10.2307/2946693.

Fulford, S.L., 2015. The surprisingly low importance of income uncertainty for precaution. Eur. Econ. Rev. *79*, 151–171. https://doi.org/10.1016/j.euroecorev.2015.07.016.

Gladstone, J., Nieboer, J., & Raghavan, K. (2020). *Understanding consumer financial wellbeing through banking data* (58; FCA Occasional Paper). Financial Conduct Authority. ⟨https://www.fca.org.uk/publication/occasional-papers/occasional-paper-58.pdf⟩.

Goda, G.S., Levy, M., Manchester, C.F., Sojourner, A., Tasoff, J., 2019. Predicting retirement savings using survey measures of exponential-growth bias and present bias. Econ. Inq. *57* (3), 1636–1658. https://doi.org/10.1111/ecin.12792.

Goffin, R.D., Olson, J.M., 2011. Is it all relative? Comparative judgments and the possible improvement of self-ratings and ratings of others. Perspect. Psychol. Sci. *6* (1), 48–60. https://doi.org/10.1177/1745691610393521.

Goodrich, B., Gabry, J., Ali, I., & Brilleman, S. (2023). *rstanarm: Bayesian applied regression modeling via Stan* (2.21.4) [Computer software]. ⟨https://mc-stan.org/rstanarm⟩.

Green, R.M., Lawyer, S.R., 2014. Steeper delay and probability discounting of potentially real versus hypothetical cigarettes (but not money) among smokers. Behav. Process. *108*, 50–56. https://doi.org/10.1016/j.beproc.2014.09.008.

Halliday, T.J., 2007. *Income volatility and health* (3234; IZA Discussion Paper Series). The Institute for the Study of Labor. ⟨https://docs.iza.org/dp3234.pdf⟩.

Hannagan, A., Morduch, J., 2015. Income gains and month-to-month income volatility: Household evidence from the US Financial Diaries. 2659883; NYU Wagner Research Paper. ⟨https://www.financialaccess.org/publications-index/2015/usfd⟩.

Harris, A.D., McGregor, J.C., Perencevich, E.N., Furuno, J.P., Zhu, J., Peterson, D.E., Finkelstein, J., 2006. The use and interpretation of quasi-experimental studies in medical informatics. J. Am. Med. Inform. Assoc. *13* (1), 16–23. https://doi.org/10.1197/jamia.M1749.

Haushofer, J., Fehr, E., 2014. On the psychology of poverty. Science *344* (6186), 862–867. https://doi.org/10.1126/science.1232491.

Henrich, J., Heine, S.J., Norenzayan, A., 2010. The weirdest people in the world? Article 2–3. Behav. Brain Sci. *33* (2–3) https://doi.org/10.1017/S0140525X0999152X.

Hipp, L., Bernhardt, J., Allmendinger, J., 2015. Institutions and the prevalence of nonstandard employment. Socio-Econ. Rev. *13* (2), 351–377. https://doi.org/10.1093/ser/mwv002.

HM Revenue & Customs. (2023). *Percentile points from 1 to 99 for total income before and after tax* [dataset]. ⟨https://www.gov.uk/government/statistics/percentile-points-from-1-to-99-for-total-income-before-and-after-tax⟩.

Holzmeister, F., Huber, J., Kirchler, M., Lindner, F., Weitzel, U., Zeisberger, S., 2020. What drives risk perception? A global survey with financial professionals and laypeople. Manag. Sci. *66* (9), 3977–4002. https://doi.org/10.1287/mnsc.2019.3526.

Kalleberg, A.L., Hewison, K., 2013. Precarious work and the challenge for Asia. Am. Behav. Sci. *57* (3), 271–288. https://doi.org/10.1177/0002764212466238.

Kazarosian, M., 1997. Precautionary savings—A panel study. Rev. Econ. Stat. *79* (2), 241–247. https://doi.org/10.1162/003465397556593.

Kirby, K.N., Maraković, N.N., 1995. Modeling myopic decisions: Evidence for hyperbolic delay-discounting within subjects and amounts. Organ. Behav. Hum. Decis. Process. *64* (1), 22–30. https://doi.org/10.1006/obhd.1995.1086.

Knight, F.H., 1921. Risk, uncertainty and profit. Houghton Mifflin.

Konovalova, E., Pachur, T., 2021. The intuitive conceptualization and perception of variance. Cognition *217*, 104906. https://doi.org/10.1016/j.cognition.2021.104906.

Kruschke, J.K., Liddell, T.M., 2018. The Bayesian New Statistics: Hypothesis testing, estimation, meta-analysis, and power analysis from a Bayesian perspective. Psychon. Bull. Rev. *25* (1), 178–206. https://doi.org/10.3758/s13423-016-1221-4.

Kvålseth, T.O., 2017. Coefficient of variation: The second-order alternative. J. Appl. Stat. *44* (3), 402–415. https://doi.org/10.1080/02664763.2016.1174195.

Lee, M.D., Wagenmakers, E.-J., 2013. Bayesian cognitive modeling: A practical course. Cambridge University Press.

Leland, H.E., 1968. Saving and uncertainty: The precautionary demand for saving. Q. J. Econ. *82* (3), 465. https://doi.org/10.2307/1879518.

Loewenstein, G., Sicherman, N., 1991. Do workers prefer increasing wage profiles? J. Labor Econ. 9 (1), 67–84.

Lugilde, A., Bande, R., Riveiro, D., 2019. Precautionary saving: a review of the empirical literature. J. Econ. Surv. *33* (2), 481–515. https://doi.org/10.1111/joes.12284.

Mani, A., Mullainathan, S., Shafir, E., Zhao, J., 2013. Poverty impedes cognitive function. Science *341* (6149), 976–980. https://doi.org/10.1126/science.1238041.

Mani, A., Mullainathan, S., Shafir, E., Zhao, J., 2020. Scarcity and cognitive function around payday: a conceptual and empirical analysis. J. Assoc. Consum. Res. *5* (4), 365–376. https://doi.org/10.1086/709885.

Mills, G.B., Amick, J., 2010. *Can savings help overcome income instability?* (Brief 18; Perspectives on Low-Income Working Families). The Urban Institute. ⟨https://www.urban.org/sites/default/files/publication/32771/412290-Can-Savings-Help-Overcome-Income-Instability-.PDF⟩.

Morris, P., A., Hill, H., D., Gennetian, L., A., Rodrigues, C., Wolf, S., 2015. *Income volatility in U.S. households with children: Another growing disparity between the rich and the poor?* (No. 1429-15; IRP Discussion Paper). Institute for Research on Poverty. ⟨https://www.irp.wisc.edu/wp/wp-content/uploads/2018/05/dp142915.pdf⟩.

Mullainathan, S., Shafir, E., 2009. Savings policy and decisionmaking in low-income households. In: Barr, M., Blank, R. (Eds.), Insufficient funds: Savings, assets, credit and banking among low-income households. Russell Sage Foundation Press, pp. 121–145.

Nunley, J.M., Seals, A., 2010. The effects of household income volatility on divorce. Am. J. Econ. Sociol. *69* (3), 983–1010. https://doi.org/10.1111/j.1536-7150.2010.00731.x.

OECD, 2023. OECD indicators. In: *Education at a glance 2023*. OECD Publishing. https://doi.org/10.1787/e13bef63-en.

Office for National Statistics. (2022). Annual Survey of Hours and Earnings (ASHE): Proportion of all employee jobs with weekly, fortnightly, four-weekly and monthly pay periods, UK, April 2021 [dataset]. ⟨https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/earningsandworkinghours/adhocs/14237annualsurveyofhoursandearningsasheproportionofallemployeejobswithweeklyfortnightlyfourweeklyandmonthlypayperiodsukapril2021⟩.

Office for National Statistics. (2023a). *Economic activity status* [dataset]. ⟨https://www.ons.gov.uk/datasets/TS066/editions/2021/versions/6/filter-outputs/e8f48b47-901b-47ab-b602-c044ca18519e⟩.

Office for National Statistics. (2023b). *Estimates of the population for England and Wales* (Mid-2022: 2023 local authority boundaries edition of this dataset) [dataset]. ⟨https://www.ons.gov.uk/peoplepopulationandcommunity/populationandmigration/populationestimates/datasets/estimatesofthepopulationforenglandandwales⟩.

Office for National Statistics. (2023c). Gender identity (eight categories) by sex, England and Wales: Census 2021 (Census 2021 edition of this dataset) [dataset]. ⟨https://www.ons.gov.uk/peoplepopulationandcommunity/culturalidentity/genderidentity/datasets/genderidentityeightcategoriesbysexenglandandwalescensus2021⟩.

Parker, J.A., 2017. Why don't households smooth consumption? evidence from a $25 million experiment. Am. Econ. J.: Macroecon. *9* (4), 153–183. https://doi.org/10.1257/mac.20150331.

Peetz, J., Robson, J., Xuereb, S., 2021. The role of income volatility and perceived locus of control in financial planning decisions. Front. Psychol. *12*, 638043 https://doi.org/10.3389/fpsyg.2021.638043.

Pepper, G.V., Nettle, D., 2017. The behavioural constellation of deprivation: causes and consequences. Behav. Brain Sci. *40*, e314 https://doi.org/10.1017/S0140525X1600234X.

Perry, V.G., Morris, M.D., 2005. Who is in control? The role of self-perception, knowledge, and income in explaining consumer financial behavior. J. Consum. Aff. *39* (2), 299–313. https://doi.org/10.1111/j.1745-6606.2005.00016.x.

Prause, J., Dooley, D., Huh, J., 2009. Income volatility and psychological depression. Am. J. Community Psychol. *43* (1–2), 57–70. https://doi.org/10.1007/s10464-008-9219-3.

Quené, H., 2007. On the just noticeable difference for tempo in speech. J. Phon. *35* (3), 353–362. https://doi.org/10.1016/j.wocn.2006.09.001.

R Core Team, 2022. *R: A language and environment for statistical computing* (4.2.0) [Computer software]. R Foundation for Statistical Computing. ⟨https://www.R-project.org/⟩.

Rosenbaum, P.R., 2015. How to see more in observational studies: Some new quasi-experimental devices. Annu. Rev. Stat. Its Appl. *2* (1), 21–48. https://doi.org/10.1146/annurev-statistics-010814-020201.

Sandmo, A., 1970. The effect of uncertainty on saving decisions. Rev. Econ. Stud. *37* (3), 353. https://doi.org/10.2307/2296725.

Sayre, G.M., 2022. The costs of insecurity: Pay volatility and health outcomes. Journal of Applied Psychology. Advance online publication. https://doi.org/10.1037/apl0001062.

Schwabe, L., Wolf, O.T., 2009. Stress prompts habit behavior in humans. J. Neurosci. *29* (22), 7191–7198. https://doi.org/10.1523/JNEUROSCI.0979-09.2009.

Senate Select Committee on Job Security, 2022. The job insecurity report. ⟨https://parlinfo.aph.gov.au/parlInfo/download/committees/reportsen/024780/toc_pdf/The jobinsecurityreport.pdf;fileType%3Dapplication%2Fpdf⟩.

Shapiro, J., Wu, S., 2011. Fatalism and savings. J. Socio-Econ. *40* (5), 645–651. https://doi.org/10.1016/j.socec.2011.05.003.

Smith, C.L., Hantula, D.A., 2008. Methodological considerations in the study of delay discounting in intertemporal choice: A comparison of tasks and modes. Behav. Res. Methods *40* (4), 940–953. https://doi.org/10.3758/BRM.40.4.940.

Stewart, N., Chater, N., Brown, G.D.A., 2006. Decision by sampling. Cogn. Psychol. *53* (1), 1–26. https://doi.org/10.1016/j.cogpsych.2005.10.003.

van Den Bergh, D., van Doorn, J., Marsman, M., Draws, T., van Kesteren, E.-J., Derks, K., Dablander, F., Gronau, Q.F., Kucharský, Š., Gupta, A.R.K.N., Sarafoglou, A., Voelkel, J.G., Stefan, A., Ly, A., Hinne, M., Matzke, D., Wagenmakers, E.-J., 2020. A tutorial on conducting and intepreting a Bayesian ANOVA in JASP. L'Ann. ée Psychol. *120* (1), 73–96. https://doi.org/10.3917/anpsy1.201.0073.

van Der Laan, L.N., Orcholska, O., 2022. Effects of digital just-in-time nudges on healthy food choice – A field experiment. Food Qual. Prefer. *98*, 104535 https://doi.org/10.1016/j.foodqual.2022.104535.

Vlaev, I., 2012. How different are real and hypothetical decisions? Overestimation, contrast and assimilation in social interaction. J. Econ. Psychol. *33* (5), 963–972. https://doi.org/10.1016/j.joep.2012.05.005.

Wang, S., Li, L.Z., Coutts, A., 2022. National survey of mental health and life satisfaction of gig workers: The role of loneliness and financial precarity. BMJ Open *12* (12), e066389. https://doi.org/10.1136/bmjopen-2022-066389.

Wang-Ly, N., Newell, B.R., 2022. Allowing early access to retirement savings: Lessons from Australia. Econ. Anal. Policy *75*, 716–733. https://doi.org/10.1016/j.eap.2022.07.002.

Weatherly, J.N., Terrell, H.K., Derenne, A., 2010. Delay discounting of different commodities. J. Gen. Psychol. *137* (3), 273–286. https://doi.org/10.1080/00221309.2010.484449.

West, C., Whillans, A., DeVoe, S., 2020. *Income volatility increases financial impatience* (21–053; Working Paper). Harvard Business School. ⟨https://www.hbs.edu/ris/Publication%20Files/21-053_54843a1e-7599-4981-a829-bac3234830cc.pdf⟩.

Zhang, C.Y., Sussman, A.B., 2023. Income predictability and budgeting. SSRN Electron. J. https://doi.org/10.2139/ssrn.4492101.

Zhang, X., Lin, W., Xue, P., 2008. Just-noticeable difference estimation with pixels in images. J. Vis. Commun. Image Represent. *19* (1), 30–41. https://doi.org/10.1016/j.jvcir.2007.06.001.